Documentation for Confluence 8.4

Contents

- Confluence Data Center and Server documentation . . . 11
  - Get started . . . 12
    - Tutorial: Navigate Confluence . . . 13
      - The dashboard . . . 14
      - The space directory . . . 16
      - The space sidebar . . . 17
      - Keyboard shortcuts . . . 20
      - Complete your mission . . . 22
    - Tutorial: Space ace . . . 23
      - Create a project space . . . 24
      - Create your personal space . . . 26
      - Create the team's PR space . . . 27
      - Delete and archive spaces . . . 29
  - Spaces . . . 30
    - Create a Space . . . 33
      - Create a Space From a Template . . . 36
      - Space Keys . . . 38
    - Navigate Spaces . . . 39
    - Space Permissions Overview . . . 41
      - Assign Space Permissions . . . 44
      - Make a Space Public . . . 47
      - Give Access to Unlicensed Users from Jira Service Management . . . 49
    - Organize your Space . . . 50
      - Set up a Space Home Page . . . 55
      - Use Labels to Categorize Spaces . . . 59
    - Customize your Space . . . 61
      - Configure the Sidebar . . . 63
      - Edit a Space's Color Scheme . . . 65
      - Apply a Theme to a Space . . . 67
        - Documentation theme migration FAQ . . . 68
      - Customize Space Layouts . . . 74
      - Change the Space Logo . . . 76
    - Copy a Space . . . 77
    - Archive a Space . . . 87
    - Delete a Space . . . 89
    - Export Content to Word, PDF, HTML and XML . . . 90
      - Customize Exports to PDF . . . 93
        - Advanced PDF Export Customizations . . . 100
        - Create a PDF in Another Language . . . 103
  - Pages and blogs . . . 104
    - Create and Edit Pages . . . 105
    - Blog Posts . . . 108
    - The Editor . . . 110
      - Symbols, emojis and special characters . . . 115
      - Collaborative editing . . . 120
    - Move and Reorder Pages . . . 123
    - Copy a Page . . . 125
    - Delete or Restore a Page . . . 128
    - Add, Remove and Search for Labels . . . 130
      - Display Pages with Label Macros . . . 133
    - Drafts . . . 134
      - Concurrent Editing and Merging Changes . . . 136
    - Page Restrictions . . . 137
    - Links . . . 142
    - Anchors . . . 147
    - Tables . . . 150
    - Add, Assign, and View Tasks . . . 153
    - Autocomplete for links, files, macros, mentions and emojis . . . 156

Page Layouts, Columns and Sections . . . 161
Create Beautiful and Dynamic Pages . . . 164
Page Templates . . . 168
Create a Template . . . 170
Create a Page from a Template . . . 174
Blueprints . . . 176
Decisions Blueprint . . . 179
File List Blueprint . . . 181
Meeting Notes Blueprint . . . 184
Product Requirements Blueprint . . . 186
Shared Links Blueprint . . . 188
Jira Report Blueprint . . . 190
Retrospective Blueprint . . . 193
How-To Article Blueprint . . . 195
Troubleshooting Article Blueprint . . . 197
Create a Blueprint-Style Report . . . 199
Import Content Into Confluence . . . 206
Import a Word Document into Confluence . . . 207
Undefined Page Links . . . 209
View Page Information . . . 210
Page History and Page Comparison Views . . . 211
Confluence Markup . . . 214
Confluence Storage Format . . . 215
Confluence Wiki Markup . . . 227
Files . . . 236
Upload Files . . . 237
Display Files and Images . . . 239
Manage Files . . . 243
Share and Comment on Files . . . 246
Edit Files . . . 248
Edit in Office using the Office Connector . . . 252
Install Atlassian Companion . . . 256
Confluence Mobile . . . 259
Using the Confluence Data Center and Server mobile app . . . 263
Using Confluence via your mobile browser . . . 267
Mobile Device Management (MDM) . . . 270
Macros . . . 274
Advanced Roadmaps Macro . . . 278
Anchor Macro . . . 280
Attachments Macro . . . 283
Blog Posts Macro . . . 287
Change History Macro . . . 291
Chart Macro . . . 294
Wiki Markup Examples for Chart Macro . . . 314
Cheese Macro . . . 317
Children Display Macro . . . 318
Code Block Macro . . . 322
Column Macro . . . 327
Content by Label Macro . . . 330
Content by User Macro . . . 334
Content Report Table Macro . . . 336
Contributors Macro . . . 339
Contributors Summary Macro . . . 343
Create from Template Macro . . . 348
Create Space Button Macro . . . 351
Excerpt Include Macro . . . 353
Excerpt Macro . . . 355
Expand Macro . . . 358
Favorite Pages Macro . . . 361
Gallery Macro . . . 363
Global Reports Macro . . . 367
HTML Include Macro . . . 369
HTML Macro . . . 371
IM Presence Macro . . . 373

Include Page Macro . . . 374
Info, Tip, Note, and Warning Macros . . . 377
Jira Chart Macro . . . 379
Jira Issues Macro . . . 383
Labels List Macro . . . 391
Livesearch Macro . . . 394
Loremipsum Macro . . . 397
Multimedia Macro . . . 399
Navigation Map Macro . . . 402
Network Macro . . . 405
Noformat Macro . . . 407
Office Excel Macro . . . 409
Office PowerPoint Macro . . . 413
Office Word Macro . . . 416
Page Index Macro . . . 419
Page Properties Macro . . . 421
Page Properties Report Macro . . . 426
Page Tree Macro . . . 430
Page Tree Search Macro . . . 434
Panel Macro . . . 436
PDF Macro . . . 439
Popular Labels Macro . . . 442
Profile Picture Macro . . . 445
Recently Updated Dashboard Macro . . . 448
Recently Updated Macro . . . 451
Recently Used Labels Macro . . . 456
Related Labels Macro . . . 459
Roadmap Planner Macro . . . 462
RSS Feed Macro . . . 464
Search Results Macro . . . 467
Section Macro . . . 469
Space Attachments Macro . . . 472
Space Details Macro . . . 475
Spaces List Macro . . . 476
Status Macro . . . 479
Table of Contents Macro . . . 482
Table of Content Zone Macro . . . 488
Task Report Macro . . . 493
Team Calendar Macro . . . 496
User List Macro . . . 499
User Profile Macro . . . 502
View File Macro . . . 505
Widget Connector Macro . . . 506
Your profile and settings . . . 513
Your User Profile . . . 514
Change Your Password . . . 517
Edit Your User Settings . . . 518
Set Your Profile Picture . . . 520
Choose Your Home Page . . . 522
Save for later . . . 523
View and Revoke OAuth Access Tokens . . . 524
Collaboration . . . 526
Network Overview . . . 527
Likes and Popular Content . . . 529
Mentions . . . 530
Share a Page or Blog Post . . . 532
Comment on pages and blog posts . . . 533
Watch Pages, Spaces and Blogs . . . 536
Manage Watchers . . . 539
Email Notifications . . . 540
Subscribe to RSS Feeds within Confluence . . . 544
Subscribe to pre-specified RSS feeds . . . 545
The RSS Feed Builder . . . 546
Subscribe to a Network RSS Feed . . . 548

Workbox Notifications . . . 550
Analytics . . . 553
View insights on your site . . . 555
View insights on spaces . . . 557
View insights on pages . . . 559
Search . . . 561
Confluence Search Syntax . . . 565
Confluence Search Fields . . . 570
Search the People Directory . . . 575
Recently Viewed Pages and Blog Posts . . . 577
Ranking of Search Results . . . 578
Permissions and restrictions . . . 581
Confluence Groups . . . 583
Check who can view a page . . . 584
Inspect permissions . . . 587
Permissions best practices . . . 599
Team Calendars . . . 604
Team Calendars Quick Tour . . . 605
Create, Add, and Edit Calendars . . . 607
Restrict a Calendar . . . 610
Add Events . . . 612
Add Jira Events . . . 614
Event Types . . . 618
Custom Event Types . . . 620
Reminders . . . 622
Embed Calendars on Confluence Pages . . . 624
Watch a Calendar . . . 625
Share Calendars . . . 627
Subscribe to Team Calendars from Third-Party Calendars . . . 629
Subscribe to Team Calendars from Microsoft Outlook . . . 632
Subscribe to Team Calendars from Outlook on the web . . . 636
Subscribe to Team Calendars from Apple Calendar . . . 637
Subscribe to Team Calendars from Apple iOS Calendar . . . 639
Subscribe to Team Calendars from Google Calendar . . . 641
Subscribe to Team Calendars from Android Calendar . . . 643
Subscribe to Team Calendars from Thunderbird . . . 645
Subscribe to Third-Party Calendars from Team Calendars . . . 647
Subscribe to Outlook Calendar from Team Calendars . . . 648
Subscribe to Apple Calendar from Team Calendars . . . 649
Subscribe to Google Calendar from Team Calendars . . . 650
Subscribe to Opsgenie Calendars from Team Calendars . . . 651
Subscribe to PagerDuty Schedules from Team Calendars . . . 652
Subscribe to Teamup Calendars from Team Calendars . . . 653
Export Team Calendars Content to Other Calendars . . . 654
Delete or Remove a Calendar . . . 656
Team Calendars FAQ . . . 658
Add-ons and integrations . . . 659
Use Jira applications and Confluence together . . . 660
Request Marketplace Apps . . . 666
Use a WebDAV Client to Work with Pages . . . 668
Mail Archives . . . 670
Add a Mail Account . . . 671
Delete and Restore Mail . . . 673
Import Mail from an mbox . . . 674
Create content with email . . . 675
Gadgets . . . 678
Activity Stream Gadget . . . 680
Confluence Page Gadget . . . 681
Confluence QuickNav Gadget . . . 684
Confluence use-cases . . . 685
Develop Technical Documentation in Confluence . . . 686
Use Confluence as a Knowledge Base . . . 692
Use Confluence as your Intranet . . . 694
Confluence for Software Teams . . . 696

Confluence administrator's guide . . . 697
Getting Started as Confluence Administrator . . . 700
Manage Users . . . 703
Add and Invite Users . . . 705
Delete or Disable Users . . . 710
Restore Passwords To Recover Admin User Rights . . . 716
Edit User Details . . . 717
Change a Username . . . 719
Managing Site-Wide Permissions and Groups . . . 721
Confluence Groups for Administrators . . . 722
Adding or Removing Users in Groups . . . 724
Global Permissions Overview . . . 726
Setting Up Public Access . . . 731
Revoke access for unlicensed users from Jira Service Management . . . 732
Configuring User Directories . . . 733
Configuring the Internal Directory . . . 735
Connecting to an LDAP Directory . . . 736
Connecting to an Internal Directory with LDAP Authentication . . . 759
Connecting to Crowd or Jira for User Management . . . 766
Managing Multiple Directories . . . 779
Managing Nested Groups . . . 783
Synchronizing Data from External Directories . . . 786
Diagrams of Possible Configurations for User Management . . . 789
User Management Limitations and Recommendations . . . 795
Requesting Support for External User Management . . . 799
Disabling the Built-In User Management . . . 801
Single sign-on for Confluence Data Center . . . 802
Managing System and Marketplace Apps . . . 803
Writing User Macros . . . 804
User Macro Template Syntax . . . 812
Customizing your Confluence Site . . . 817
Changing the Look and Feel of Confluence . . . 818
Customizing the Confluence Dashboard . . . 819
Changing the Site Logo . . . 820
Customizing Color Schemes . . . 822
Styling Confluence with CSS . . . 823
Working with Themes . . . 828
Customizing Site and Space Layouts . . . 831
Customizing a Specific Page . . . 839
Customizing the Login Page . . . 840
Modify Confluence Interface Text . . . 841
Customizing Email Templates . . . 842
Changing the Default Behavior and Content in Confluence . . . 843
Administering Site Templates . . . 844
Changing the Site Title . . . 845
Choosing a Default Language . . . 846
Configuring the Administrator Contact Page . . . 848
Configuring the Site Home Page . . . 850
Customizing Default Space Content . . . 852
Editing the Site Welcome Message . . . 853
Integrating Confluence with Other Applications . . . 855
Linking to Another Application . . . 856
Configure an outgoing link . . . 858
Configure an incoming link . . . 860
OAuth 2.0 scopes for incoming links . . . 861
Configuring Workbox Notifications . . . 862
Integrating Jira and Confluence . . . 867
Registering External Gadgets . . . 868
Configuring the Office Connector . . . 872
Managing Webhooks . . . 875
Managing your Confluence License . . . 879
Managing Confluence Data . . . 883
Database Configuration . . . 884
Database JDBC Drivers . . . 885

Database Setup for Oracle . . . 887
Database Setup for PostgreSQL . . . 891
Database Setup for SQL Server . . . 894
Database Setup For MySQL . . . 897
Embedded H2 Database . . . 902
Migrating to Another Database . . . 904
Configuring Database Character Encoding . . . 909
Configuring database query timeout . . . 911
Surviving Database Connection Closures . . . 912
Configuring a datasource connection . . . 914
Configuring Confluence Data Center to work with Amazon Aurora . . . 920
Encrypt database password . . . 924
Backup and Restore . . . 935
Production Backup Strategy . . . 938
Scheduling a Backup . . . 940
Back up a Site . . . 945
Back up a Space or multiple Spaces . . . 947
Restore a Site . . . 949
Restore a Space or multiple Spaces . . . 953
Restore a Test Instance from Production . . . 957
Restore data from other backups . . . 958
Retrieving file attachments from a Backup . . . 959
Troubleshooting failed XML site backups . . . 962
Import a space from Confluence Cloud . . . 969
Attachment Storage Configuration . . . 972
Hierarchical File System Attachment Storage . . . 975
Configuring Attachment Size . . . 980
Configuring S3 object storage . . . 981
Confluence Data Model . . . 992
Finding Unused Spaces or Pages . . . 1002
Data Import and Export . . . 1003
Import a Text File . . . 1004
Auditing in Confluence . . . 1005
Audit Log Events in Confluence . . . 1009
Audit Log Integrations in Confluence . . . 1017
Set retention rules to delete unwanted data . . . 1020
Data pipeline . . . 1029
Data pipeline export schema . . . 1036
Configuring Confluence . . . 1042
Viewing System Information . . . 1043
Tracking Customizations Made to your Confluence Installation . . . 1044
Viewing System Properties . . . 1045
Configuring the Server Base URL . . . 1046
Configuring the Confluence Search and Index . . . 1048
Configuring Indexing Language . . . 1049
Configuring Search . . . 1050
Content Index Administration . . . 1051
Enabling OpenSearch . . . 1057
Rebuilding the Ancestor Table . . . 1058
Setting Up Confluence to Index External Sites . . . 1059
Setting Up an External Search Tool to Index Confluence . . . 1060
Configuring Mail . . . 1061
Configuring a Server for Outgoing Mail . . . 1062
Configuring a Server for Incoming Mail . . . 1064
Setting Up a Mail Session for the Confluence Distribution . . . 1066
Configuring the Recommended Updates Email Notification . . . 1067
The Mail Queue . . . 1068
Configuring Character Encoding . . . 1069
Troubleshooting Character Encodings . . . 1070
Other Settings . . . 1074
Configuring a WebDAV client for Confluence . . . 1075
Configuring HTTP Timeout Settings . . . 1081
Configuring Number Formats . . . 1082
Configuring Shortcut Links . . . 1083

Configuring Time and Date Formats . . . 1085
Enabling the Remote API . . . 1088
Enabling Threaded Comments . . . 1089
Installing a Language Pack . . . 1090
Installing Patched Class Files . . . 1092
Configuring System Properties . . . 1093
Recognized System Properties . . . 1096
Working with Confluence Logs . . . 1122
Configuring Logging . . . 1126
log4j Logging Levels . . . 1129
Troubleshooting SQL Exceptions . . . 1130
Configure access logs . . . 1131
Scheduled Jobs . . . 1133
Configuring the Allowlist . . . 1140
Configuring the Time Interval at which Drafts are Saved . . . 1142
Configuring Confluence Security . . . 1143
Confluence Security Overview and Advisories . . . 1144
Proxy and HTTPS setup for Confluence . . . 1147
Configuring Web Proxy Support for Confluence . . . 1149
Connecting to LDAP or Jira applications or Other Services via SSL . . . 1151
Using Apache with mod_proxy . . . 1153
Running Confluence behind NGINX with SSL . . . 1159
Running Confluence Over SSL or HTTPS . . . 1163
Using Apache to limit access to the Confluence administration interface . . . 1169
Using Apache with mod_jk . . . 1171
Using mod_rewrite to Modify Confluence URLs . . . 1172
Configuring Secure Administrator Sessions . . . 1173
Confluence Cookies . . . 1174
Using Fail2Ban to limit login attempts . . . 1177
Securing Confluence with Apache . . . 1179
Best Practices for Configuring Confluence Security . . . 1180
Hiding the People Directory . . . 1183
Configuring Captcha for Spam Prevention . . . 1184
Hiding External Links From Search Engines . . . 1185
Configuring Captcha for Failed Logins . . . 1186
Configuring XSRF Protection . . . 1188
User Email Visibility . . . 1189
Anonymous Access to Remote API . . . 1190
Configuring RSS Feeds . . . 1191
Preventing and Cleaning Up Spam . . . 1192
Configuring a Confluence Environment . . . 1194
Confluence Home and other important directories . . . 1195
Application Server Configuration . . . 1199
Managing Application Server Memory Settings . . . 1200
Starting Confluence Automatically on System Startup . . . 1201
Start Confluence Automatically on Linux . . . 1202
Start Confluence Automatically on Windows as a Service . . . 1206
Performance Tuning . . . 1209
Cache Performance Tuning . . . 1212
Cache Statistics . . . 1214
Memory Usage and Requirements . . . 1217
Requesting Performance Support . . . 1219
Compressing an HTTP Response within Confluence . . . 1223
Garbage Collector Performance Issues . . . 1224
Troubleshooting Slow Performance Using Page Request Profiling . . . 1225
Confluence Diagnostics . . . 1228
Faster permissions service . . . 1230
Confluence guardrails . . . 1232
Data Collection Policy . . . 1236
Managing emojis . . . 1237
Administering Collaborative Editing . . . 1238
Possible Confluence and Synchrony Configurations . . . 1243
Configuring Synchrony . . . 1249
Set up a Synchrony cluster for Confluence Data Center . . . 1252

Migrate from a standalone Synchrony cluster to managed Synchrony 1260
Troubleshooting Collaborative Editing 1262
Using read-only mode for site maintenance 1267
Administering the Atlassian Companion App 1270
Notifications from Atlassian 1273
Administer analytics 1274
Monitor application performance 1279
Monitor Confluence with Prometheus and Grafana 1282
App metrics reference 1284
Live Monitoring Using the JMX Interface 1293
Confluence installation and upgrade guide 1310
System Requirements 1311
Server Hardware Requirements Guide 1313
Running Confluence in a Virtualized Environment 1317
Confluence Installation Guide 1318
Installing Confluence 1319
Get a Confluence Data Center trial license 1320
Install a Confluence Data Center trial 1321
Installing Confluence on Windows 1323
Installing Confluence on Linux 1337
Unattended installation 1352
Change listen port for Confluence 1355
Start and Stop Confluence 1357
Installing Confluence Data Center 1360
Upgrading Confluence Data Center 1365
Adding and Removing Data Center Nodes 1367
Change Node Discovery from Multicast to TCP/IP or AWS 1368
Running Confluence Data Center in AWS 1371
Getting started with Confluence Data Center on Azure 1373
Running Confluence Data Center on a Kubernetes cluster 1378
Installing Java for Confluence 1380
Setting the JAVA_HOME Variable in Windows 1381
Change the Java vendor or version Confluence uses 1383
Creating a Dedicated User Account on the Operating System to Run Confluence . . 1387
Confluence Setup Guide 1389
Configuring Jira Integration in the Setup Wizard 1394
Upgrading Confluence 1398
Upgrading Beyond Current Licensed Period 1406
Confluence Post-Upgrade Checks 1408
Migration from Wiki Markup to XHTML-Based Storage Format 1409
Migration of Templates from Wiki Markup to XHTML-Based Storage Format 1413
Upgrading Confluence Manually 1415
Create a staging environment for upgrading Confluence 1421
Upgrade Confluence without downtime 1425
Upgrade a Confluence cluster manually without downtime 1428
Upgrade a Confluence cluster on AWS without downtime 1433
Upgrade a Confluence cluster through the API without downtime 1438
Roll back a rolling upgrade 1441
Upgrade task troubleshooting 1443
Supported Platforms 1445
End of Support Announcements for Confluence 1449
Bundled Tomcat and Java versions 1474
Supported Platforms FAQ 1491
Migrate your Confluence site 1492
Upgrade from Confluence Server to Data Center 1493
Migrate from Confluence Cloud to Server 1495
Migrating Confluence Between Servers 1500
Move to a non-clustered installation 1501
From Confluence Evaluation through to Production Installation 1503
Cloud Migration Assistant for Confluence 1506
Confluence Data Center documentation 1517
Getting Started with Confluence Data Center 1518
Confluence Server and Data Center feature comparison 1520
Clustering with Confluence Data Center 1523

External Process Pool for Confluence Data Center 1532
Document conversion for Confluence Data Center 1534
PDF export in Confluence Data Center 1536
Restricted Functions in Confluence Data Center 1538
Set up a Confluence Data Center cluster 1540
Confluence Data Center Performance 1546
Confluence Data Center disaster recovery 1552
Data Center Troubleshooting 1557
Troubleshooting a Data Center cluster outage 1561
Use a CDN with Atlassian Data Center applications 1563
Configure your CDN for Confluence Data Center 1567
Improving instance stability with rate limiting 1571
Adjusting your code for rate limiting 1578
Running Confluence Data Center on a single node 1581

Confluence Data Center and Server documentation

Confluence is where you create, organize, and discuss work with your team.

Get started

Kick start your Confluence journey with our navigation and space tutorials.

View guide

Whats new in 8.4

Time to upgrade? Get the low down on the latest and greatest changes in Confluence.

View release notes

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 organization 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 colonization 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

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**Saved for later**.
- **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 whole right hand column is ready and waiting for you to customize.

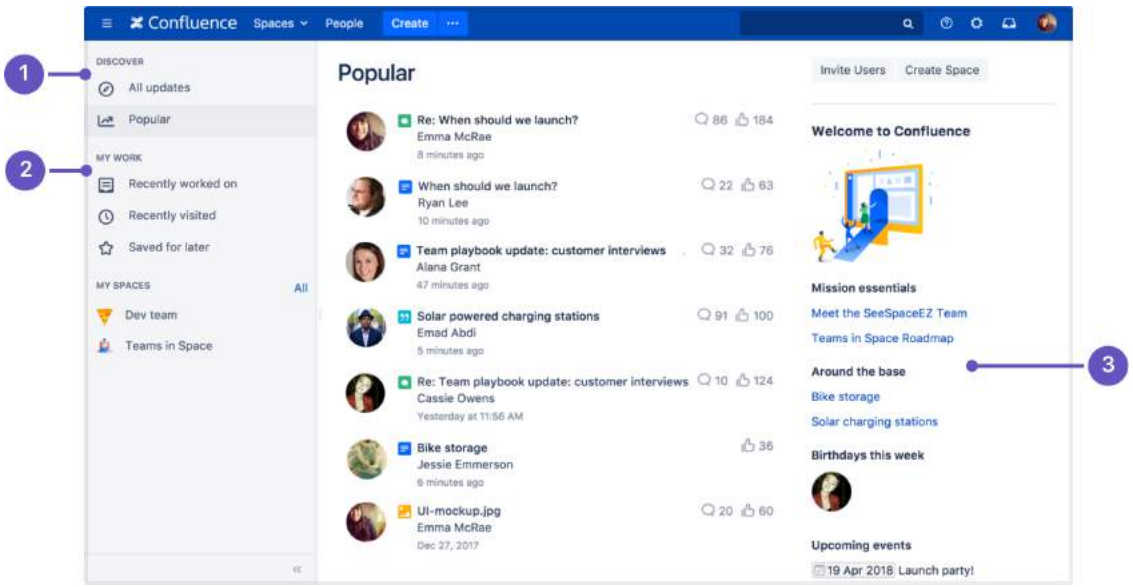

1. **Discover:** see what's happening in your site.
2. **Your work:** get work done with recent and useful pages at your fingertips.
3. **Customize:** admins can add useful content to welcome people to the site.

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 cancustomize the global dashboardthat all users see

Next

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

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.

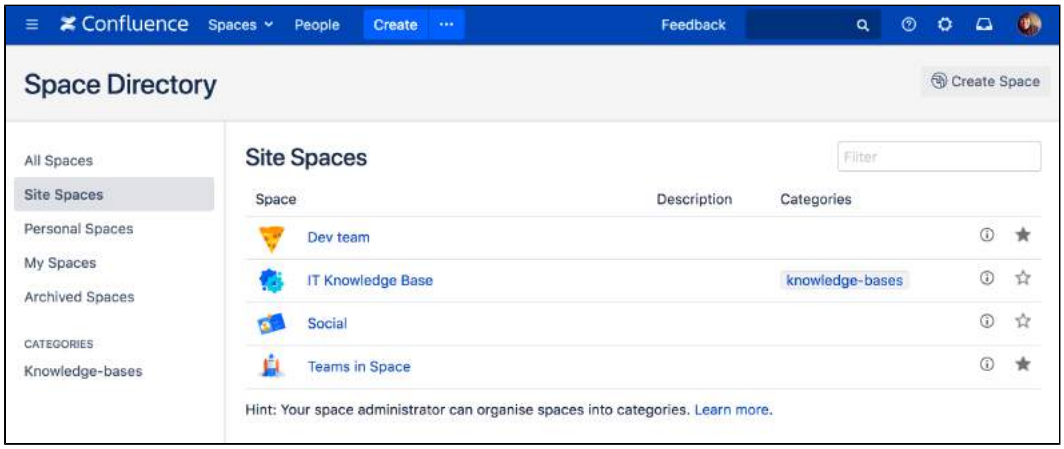

You can also categorize spaces with labels you create yourself. We're not going to cover that here, but, if you'd like to know more, you can check outUse Labels to Categorize Spaces.

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 sidebarisa feature of every Confluence space; it's where you'll find the page tree (a hierarchical list of pages in the space), customizable 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.

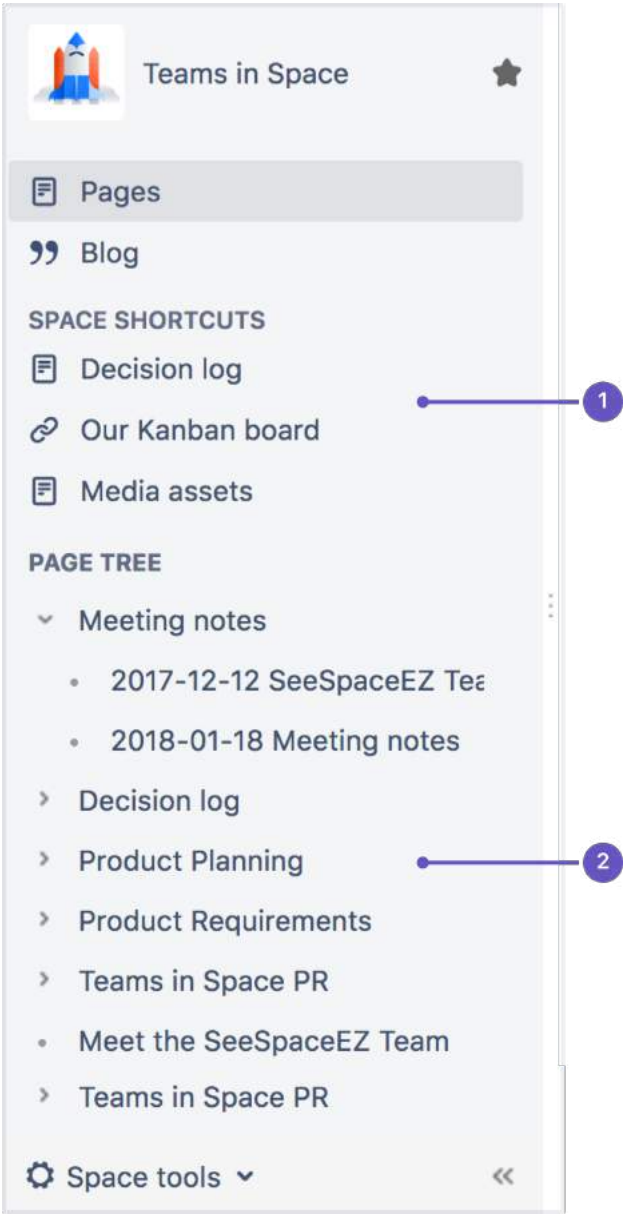

1. **Space shortcuts**: link to Confluence pages or other pages on the web
2. **Page tree**: hierarchical view of pages in this space.

The page tree in the sidebar shows the 200 pages closest to where you are. Hit Show all pages, if you want to see all the pages in a space.

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.

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:

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

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 toConfluence pages or spaces, or to any other content on the web. Try linking tothis 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

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

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*withinConfluence*.

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 thehelp iconat 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 for your specific operating system and browser.

Screenshot: Keyboard shortcuts menu for Google Chrome on macOS.

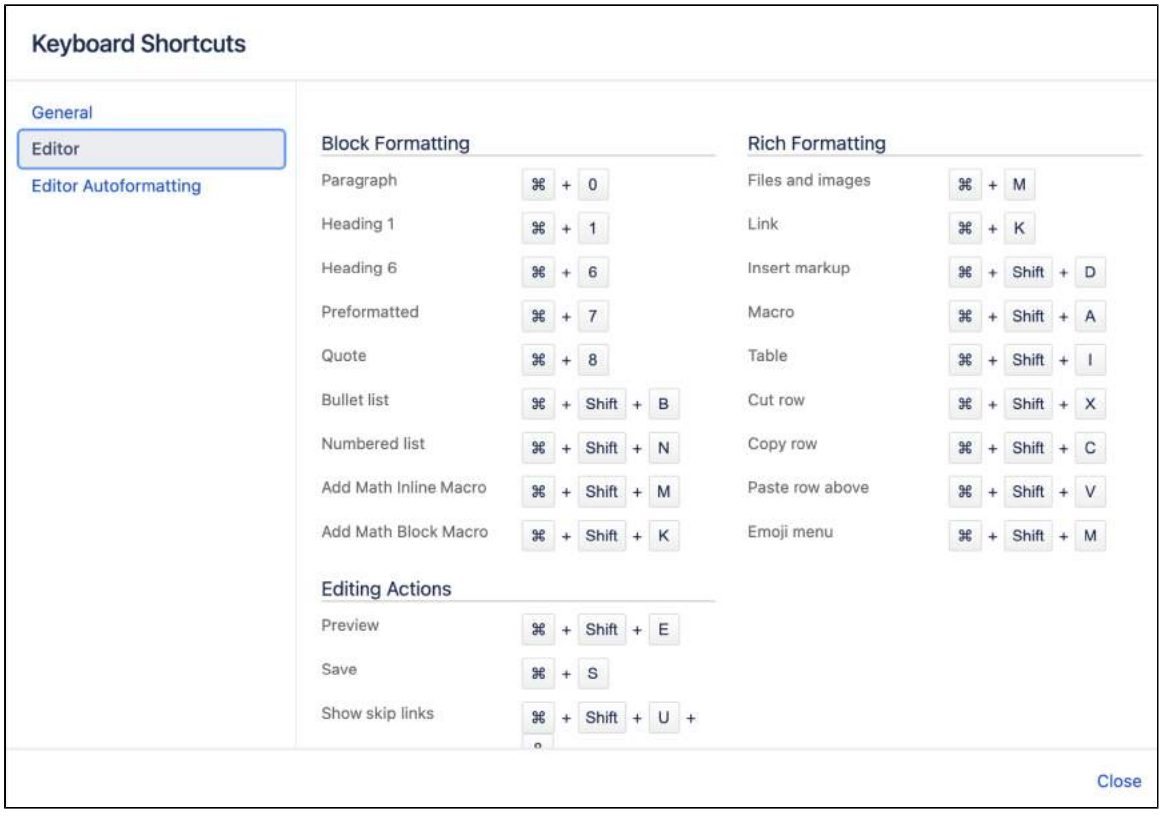

The keyboard shortcuts are broken up into 3 categories:

- **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

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

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 favorited a space
- Taken a look at and customized 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 aspace ace.

Tutorial: Space ace

This tutorial will take you on a journey through Confluence to create and customize 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 ofresponsibility. You need the power to go with it.

Mission brief

You're in charge of organizing 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 organize 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
- Createyour own space to keep yourself organized
- 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 Confluencespaces 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 colonization 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 organization. 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 colonization 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 colonization 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 organization, but not to the general public.

Every space has a default home page, which you can customize 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.

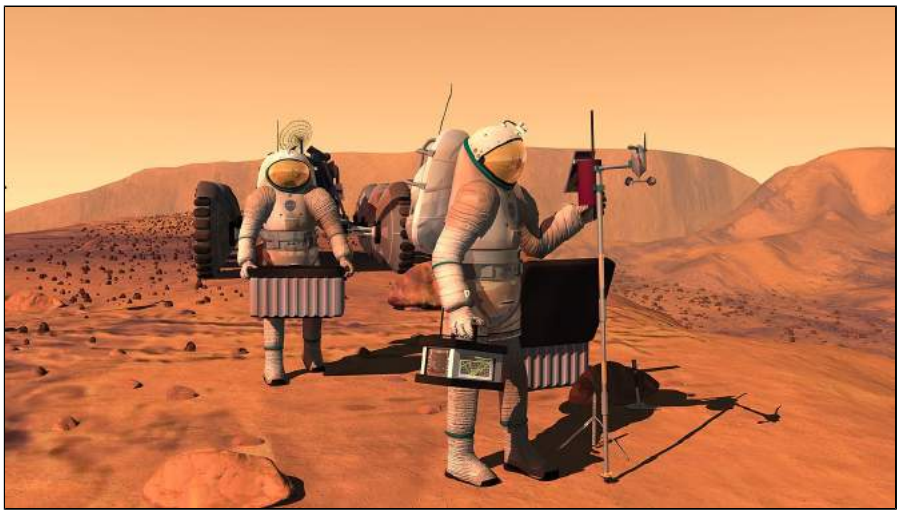

A**human mission to Mars**has been the subject ofscience fiction,engineering, and scientific proposals throughout the 20th century and into the 21st century. The plans comprise proposals to land onMars, eventuallysettling onandterraforming the planet, while exploiting its moons,Phobosand Deimos.

Exploration of Marshas 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. Thelist of manned Mars mission plans in the 20th centuryshows the various mission proposals that have been put forth by multiple organizations andspace agenciesin this field ofspace exploration.

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

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

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 ofthe 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.

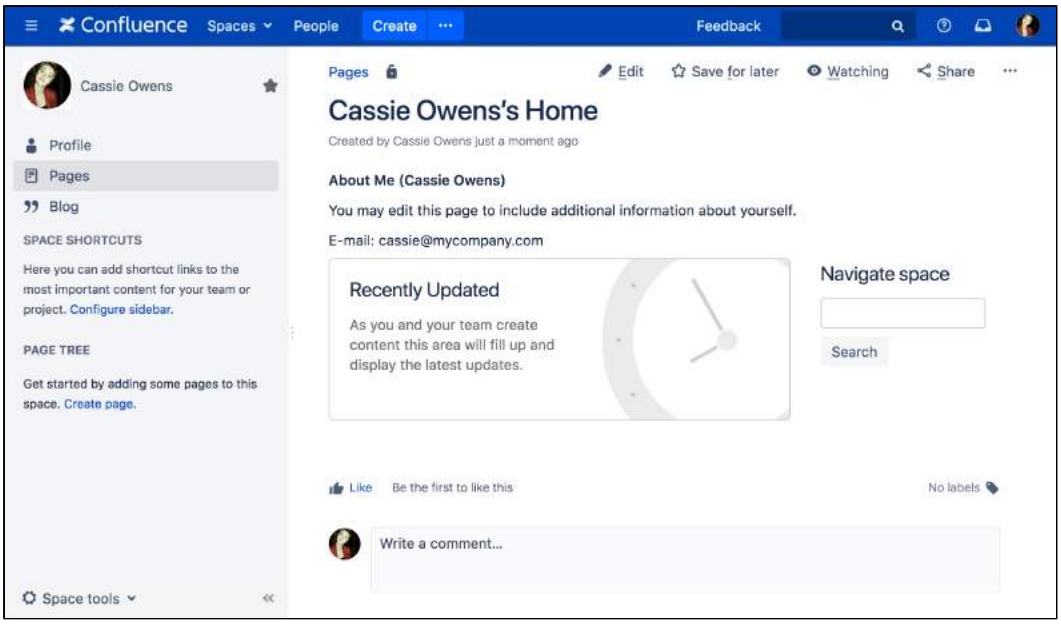

4. Choose**Space tools**>**Permissions**from the bottom of the siderbar
5. Hit**Edit Permissions**(enter your password if prompted)

   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 colonize 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.

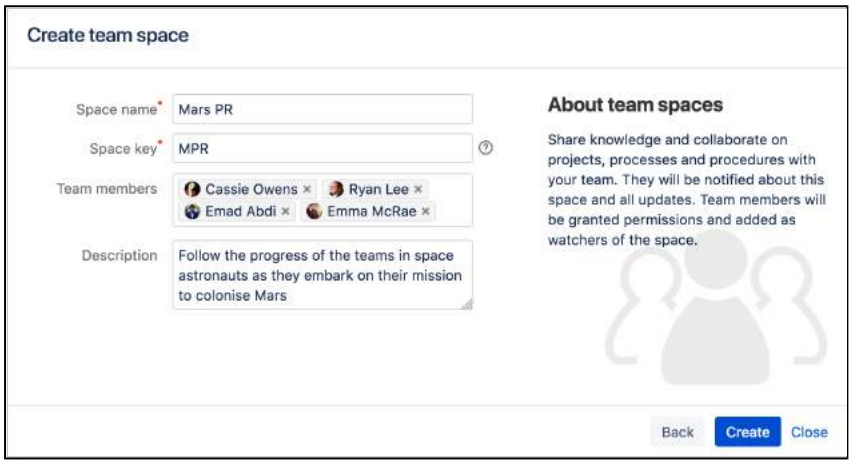

Each space you create also has its own blog, so your social media team will be able to create posts in this space and speak directly to all those adoring fans. But none of those fans can see this space. Yet.

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 license for Confluence.

Next

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

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.

Spaces

What is a space?

Spaces are Confluence's way of organizing content into meaningful categories. Think of it like having different folders into which you can put your work.

Spaces come in two main varieties:

- **Site spaces** These spaces are found in the **Space Directory** and are the areas where you create content and collaborate with others. They are sometimes called global spaces.
- **Personal spaces** Every Confluence user can set up a personal space which they can keep private or make public so others can view and edit. Personal spaces are listed in the **People Directory** and found under your personal profile.

On this page:

- What is a space?
- How do I use a space?

Related pages:

- Create a Space
- Space Permissions Overview
- Navigate Spaces
- Organize your Space
- Customize your Space
- Archive a Space
- Export Content to Word, PDF, HTML and XML
- Delete a Space

How do I use a space?

Create as many spaces as you need to get things done:

- **Team spaces** Give each team (QA, HR, Engineering, Support, ...) their own space so they can focus and make their information easier for everyone to find.
- **Project spaces** Put all the information related to your project in one place. This allows everyone to work together in Confluence instead of emailing back and forth.
- **Personal space** Store everything you're working on individually, keep your to-do lists, and polish any content before you move it into another shared space.

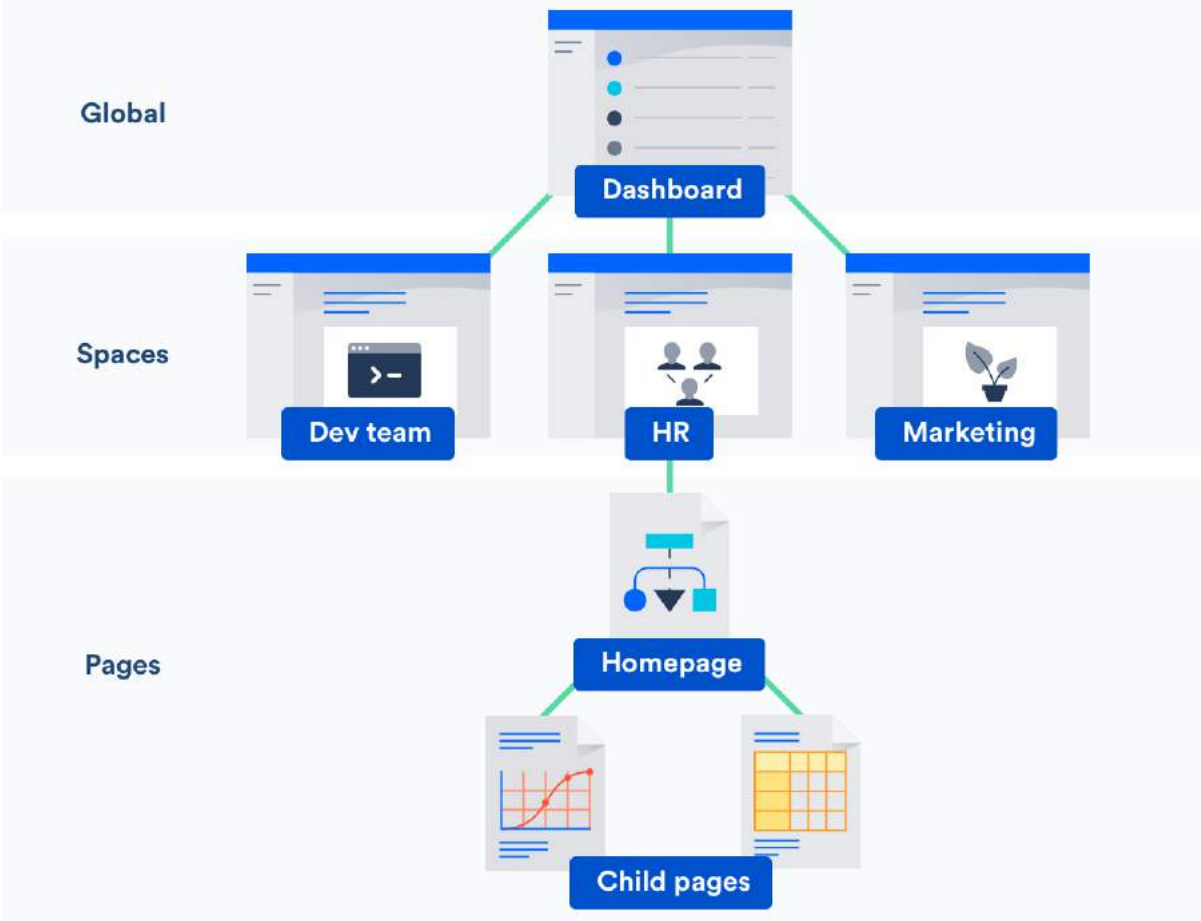

Stay up to date with spaces

- Watch a space to stay up-to-date with any changes.
- Add important spaces to your dashboard so that they're easy to find again.

Administer spaces
If you have admin permissions for a space you can:

- Change the colors, logo, sidebar and homepage of that space.
- Control who can view and edit your space.
- Make your space public to share it with people who don't have access to Confluence.

Space Permissions
Some things we should make clear about space admin permissions:

- The person who creates a space automatically has admin permissions for that space.
- Space admins can grant admin permissions to others.
- Space admins don't have to be Confluence admins and can have special permissions for a single space. For example, you are the admin for your personal space, no matter what kind of access you have anywhere else.

Want more ideas for using spaces? Check out our kickass guides on how to:

- Use Confluence as a Knowledge Base
- Develop Technical Documentation in Confluence
- Use Confluence as your Intranet
- Use Confluence for Software Teams

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Create a Space

There's no limit to the number of Spaces you can create on Confluence. You can choose to set up a space for each team, project, or a mix of both depending on your needs.

Each space in Confluence functions autonomously, which means that each space:

- Has its own homepage, blog, pages, comments, files, and RSS feeds.
- Can be customized with different color schemes, logo and sidebar.
- Has its own set of permissions, as set by the space admin.

For example, an IT team can create one overarching space with all their roadmaps, details of sub-teams, and a list of all the people and roles within that team. They could then create a new space for each sub-team, such as Quality Assurance, Developers, and Documentation, with guidelines, long term plans, and knowledge articles within them. Each project that these teams work on could also have its own space, which could be linked to the team spaces using labels.

On this page:

- Create a personal space
- Create a site space
- Space permissions
- Linking related spaces
- Tips

Related pages:

- Spaces
- Space Keys
- Create a Space From a Template

Create a personal space

Your personal space is always owned by you, and you can use it to store your individual work, keep track of tasks, blog about what you've been working on, or just use it to polish your pages before you move them into a site space.

1. Choose your **profile picture** at top right of the screen, then choose **Add Personal Space..**
2. Choose **Create**.

You can change the permissions for 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.

To create a personal space you need the 'Personal Space' global permission.

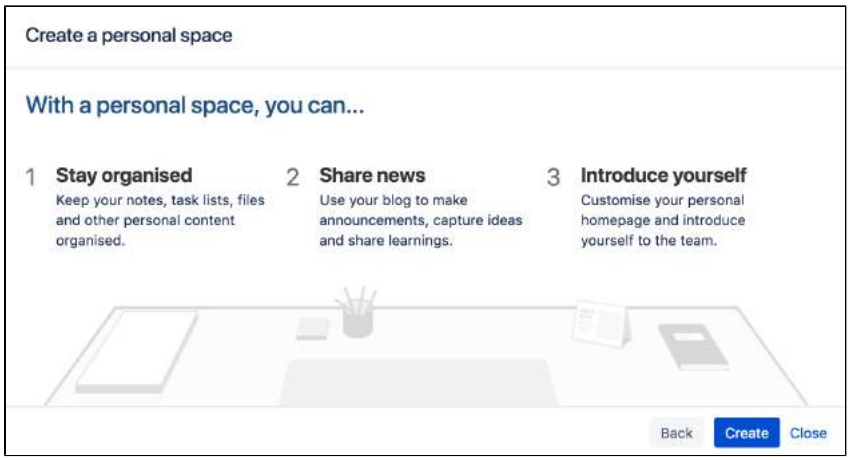

Create a site space

You can create a site space for any team or project that would benefit from having a place where people can work together and store related files. You can create these as blank spaces, or use templates, called space blueprints , to help you create Team spaces, Knowledge Bases spaces, or Documentation spaces quickly and easily.

1. Hit **Spaces** > **Create space** in the header.
2. Pick a type of space.
3. Enter the required details and create your space.

⚠Choose your space key carefully as you can't change this later.

Each space you create will automatically have a home page that you can customize to display relevant information for people viewing the space. If you use a space blueprint when creating a space, it will customize your home page for you.

ℹTo create a site space you need the 'Create Space' global permission.

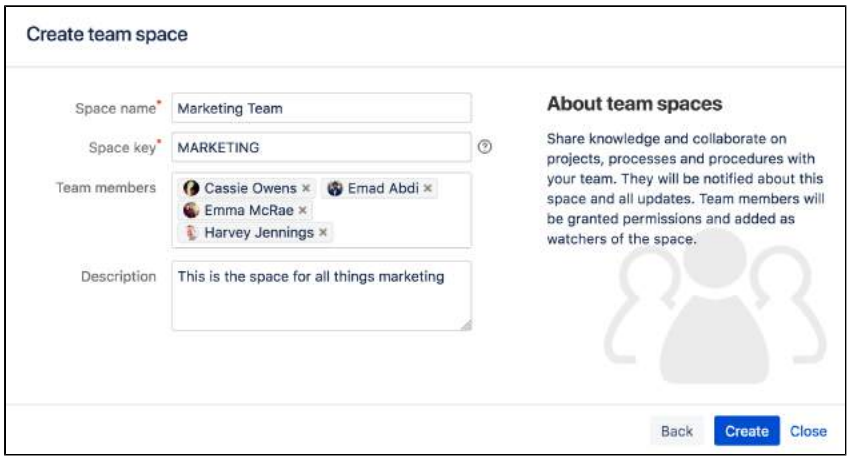

Space permissions

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, which means that they can then grant permissions to other users and groups. See Space Permissions Overview for more information.

ℹSystem Administrators can edit the permissions of spaces in their Confluence site at any time.

Linking related spaces

You can link related spaces together using labels. This will create categories in the space directory for each label, grouping all spaces with that label together.

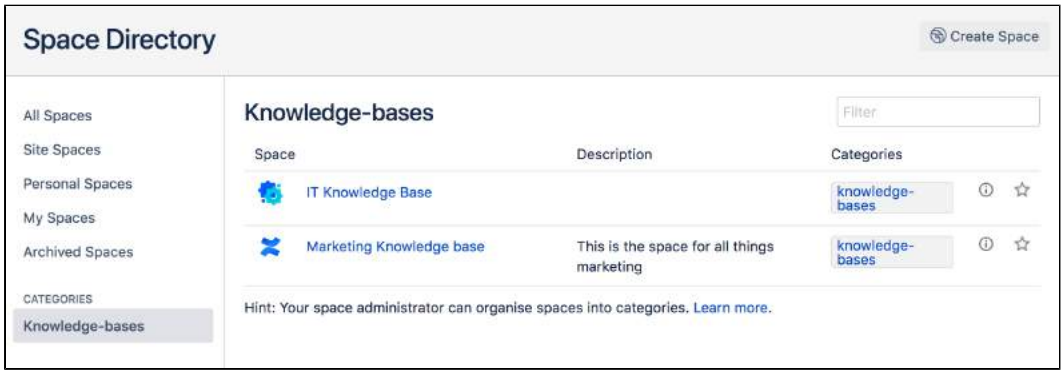

- You can also add a space description to make it easier for visitors to find the right space within each category.
- To help navigate between related spaces, you can use the Spaces List Macro on a page and filter by category. This will let you insert a list of all the other spaces in a certain category into your space. You can use this, for example, to keep a list in your team space of all the project spaces your team is working on.
- If you want to link to only certain pages of related content, rather than whole spaces, you can use the Content Report Table Macro. You can use this, for example, in a space that functions as a workplace directory, to create a list of all the team pages with everyone's roles and contact details across your organization.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Tips

- If your needs change, or your spaces grow too big, it's easy to copy or move content from one space to another.
- If the content or purpose of your space changes, you can update the space name, logo, colors and description to reflect those changes.
- If you no longer need a space, such as when a project has been completed, you can archive it, which makes it less visible but retains the content on your site so that you can still refer back to it later.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Create a Space From a Template

Confluence comes with a number of templates, known as **space blueprints**, which have a customized homepage and sidebar, and may contain page blueprints or sample content to help you get started.

Related pages:

- Create a Space

This page is about space blueprints. You can also use Blueprints to create individual pages.

Your Confluence administrator can customize or disable particular blueprint templates, so they may be different to the examples shown here.

Types of space blueprints

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 grants 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 customizable. The Knowledge Base space blueprint also Use Jira applications and Confluence together.

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.

Software project space

This space is designed to help you organize your software project. The purpose-built space home page lets you view and edit your roadmap, see team members, and Use Jira applications and Confluence together#JIRA Software. Create pages in this space for requirements, meeting notes, decisions, retros, and more.

The software project space blueprint will only appear if you have linked Confluence to your Jira Software instance.

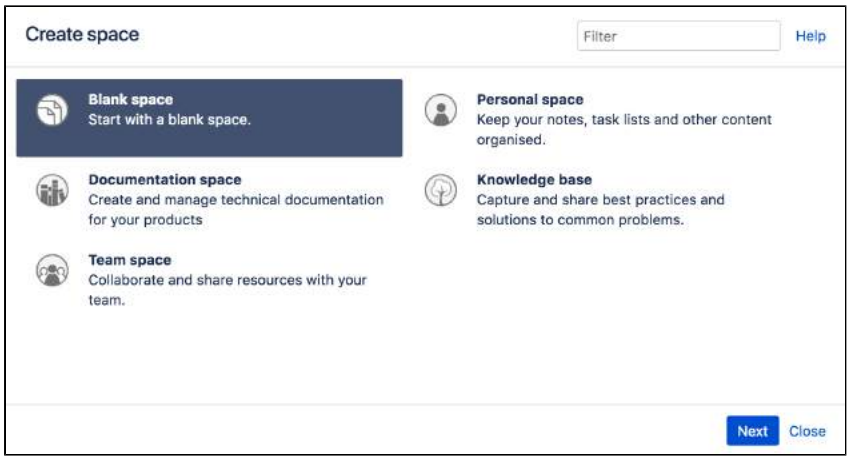

Check out our guides for some more tips on how to:

- Use Confluence as a Knowledge Base
- Develop Technical Documentation in Confluence
- Use Confluence as your Intranet
- Confluence for Software Teams

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Space Keys

Each Confluence space has a**space key**,whichis a short,unique identifier that formspart of the URL for that space.

Related pages:

- Create a Space
- Create a Space From a Template

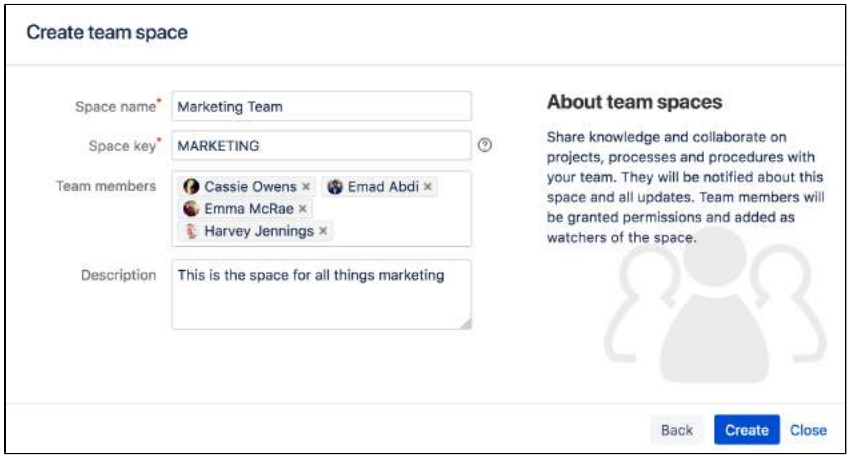

When you create a site space, Confluencewill suggesta space keyor you can enter your own key to make it more memorable.

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://<yoursite>/display/marketing`

ℹPersonal spaces always use your username as the space key.

Choosing a space key

Each space key:

- Must be unique.
- Can contain any alphanumeric character (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!

Navigate Spaces

How is content arranged in spaces?

Think of a space as the container that holds all the important stuff a team, group, or project needs to work. These are autonomous that means that each space has its own pages, blogs, files, comments and RSS feeds.

Each space is automatically created with a homepage - the first page you'll see when you navigate to the space. You can edit your homepage and your sidebar to help people navigate their way around your space.

Spaces can't be nested you can't have a space within a space but you can Use Labels to Categorize Spaces. Spaces with the same label will appear together in the the space directory and in the recent activity area of the dashboard.

On this page:

- How is content arranged in spaces?
- View all spaces in Confluence

Related pages:

- Organize your Space
- Watch Pages, Spaces and Blogs
- Pages and blogs
- Search

Inside the space, you can nest your pages, and you can create as many levels of hierarchy as you need. Each space also has its own blog, which lets you 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, and the pages and blogs within it, usingSpace Permissions Overview.

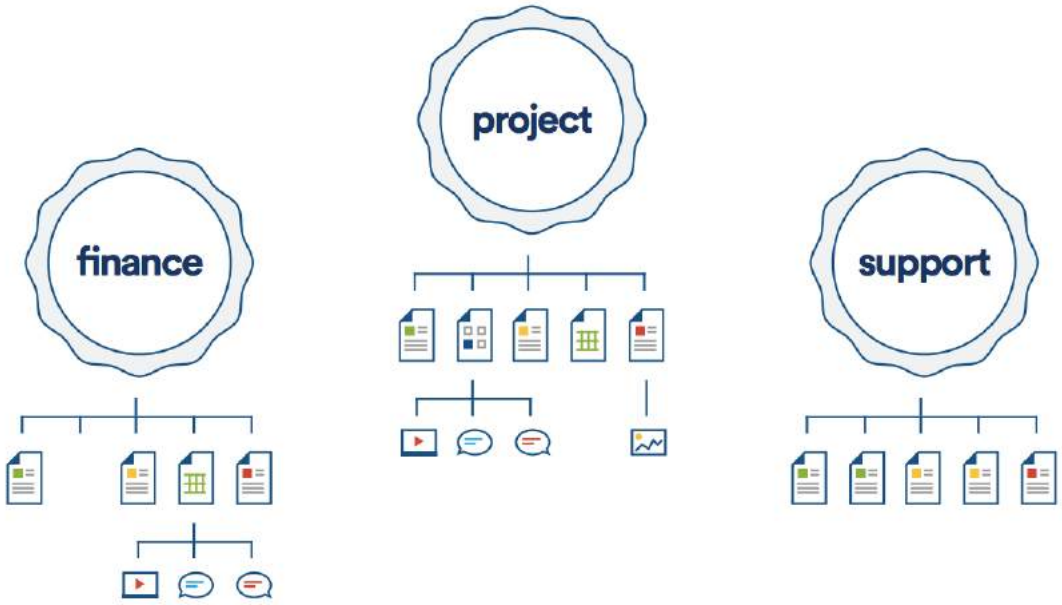

View all spaces in Confluence

There are two main ways to view spaces in Confluence:

- **The space directory**choose**Spaces > Space directory**in the Confluence headerfor 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**.

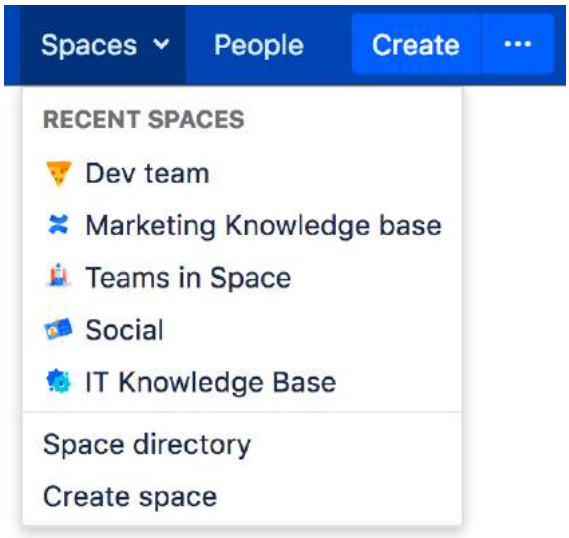

The **spaces menu** in the header also displays a list of your recently viewed spaces, allowing you to quickly navigate to the things you view most often.

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

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Space Permissions Overview

Every Confluence space has its own set of permissions which determine what people can do in the space.

Space permissions are set by the space administrator. The user who created the space is automatically a space administrator, and other users can also be granted Space Admin permission.

Related pages:

- Confluence Groups
- Assign Space Permissions
- Page Restrictions
- Global Permissions Overview

This page is about **Space Permissions**. Confluence also lets you set Page Restrictions.

Edit space permissions

To change permissions for a space, choose**Space tools**>**Permissions**from the bottom of the sidebar, then choose**Edit Permissions**to change permission settings.

SeeAssign Space Permissions for more details.

Permissions Summary

The following permissions can be assigned in a space:

Category	Permission
All	**View**gives you permission to access the content in this space, and see it in the space directory and other places like the dashboard. **Delete own**gives you permission to delete any pages, blogs, attachments and comments you've created in this space (regardless of whether other users have subsequently edited the content).
Pages	**Add page** gives you permission to create new pages and edit existing pages in this space (assuming the page is not restricted for editing). **Delete page** gives you permission to delete any page in the space.
Blog	**Add blog** gives you permission to create new blog posts and edit existing blog posts in this space (assuming the blog post is not restricted for editing). **Delete page** gives you permission to delete any blog post in the space. Delete permission is also required to move a page or blog to a different space.
Attachments	**Add attachment** gives you permission to upload (attach) files to pages and blog posts in this space, and to edit attached files using the Companion app. **Delete attachment** gives you permission to remove attached files from pages or blog posts in the space. People with only Add page or blog permissions can still insert existing attached files in the editor, and remove files from the editor, so they're not displayed on the page or blog post. They can't however upload a new file, a new version of an existing file, edit an existing file using the Companion app, or delete the attached file itself.
Comments	**Add comments** gives you permission to add comments to a page, blog post or attached file. **Delete comments** gives you permission to delete any comment on a page, blog post or attached file.

Restrictions	**Add restrictions** gives you permission to apply page-level restrictions to a page or blog post. You can restrict a page for viewing, or just for editing. **Delete restrictions** gives you permission to remove restrictions from any page or blog post.
Mail	**Delete mail** gives you permission to delete mail items that have been archived in this space. This is not a commonly used feature.
Space	**Export space** gives you permission to export all the contents of the space to PDF, HTML or XML. This is different to single page exports - anyone who can view a page can also export it. **Admin** gives you permission to access all space administration tools, including things like permissions, templates, look and feel, and the ability to delete the whole space.

Here's how it looks on the **Permissions** screen:

How space permissions work

Space permissions are additive. If a user is granted permissions as an individual or as a member of one or more groups, Confluence will combine these permissions together. This is sometimes known as their effective permissions.

> Sasha is a member of the `confluence-users` group and the `developers` group. The `confluence-users` group has 'export' permission, but does not have 'restrict' permission. The `developers` group has 'restrict' permission but does not have 'export' permission.
>
> By being a member of these two groups, Sasha can restrict *and* export content. The permissions do not conflict, they combine to determine what Sasha is allowed to do in this space.

If you have Confluence Data Center,Inspect permissions provides space admins and Confluence administrators a great way to view someone's effective permissions.

Who is the space admin?

The user who created the space is automatically a space administrator, and other users can also be granted Space Admin permission.

To find out who is an administrators in your space, either:

- Go to**Space Tools**>**Overview** in the space.
- Go to **Spaces** > **Space Directory** on the header, then choose the**Space Details**icon beside the space.

If you accidentally deny all admin access to a space, so that nobody has access to administer the space any more, you can ask someone with Confluence Administrator global permission to recoverSpace Permissions.

Space admin superpowers

Space administrators can do a lot of things in the space such as:

- grant permissions to users and groups (and themselves)
- create templates

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

- change the look and feel
- delete the space
- manually remove page restrictions (including on pages they can't see)
- manage watchers, to change who is watching a page
- inspect permissions to see what users can do in the space (Data Center only)

Confluence administrators aren't necessarily space administrators.If they don't have the Space Admin permission (as an individual or member of a group), they can recover permissions to the space, which will grant them space admin permission.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Assign Space Permissions

If you are the administrator of a space, you control the permissions for it. You can choose to assign/revoke permissions on either an individual user basis, or usingConfluence Groups.

> Want to know the best way to set permissions for your team's needs? Check out our Permissions best practices guide.

To view 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**.

The Edit Space Permissions page is divided into the following sections:

On this page:

- Grant space permission
- Set default space permission for all new spaces
- Revoke space permissions
- Inspect permissions
- Manage and recover space admin permissions

Related pages:

- Space Permissions Overview
- Global Permissions Overview
- Make a Space Public
- Give Access to Unlicensed Users from Jira Service Management

- **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). Note: allowing anonymous access in a space will allow all logged in users to see that space, even if your site is notPublic.

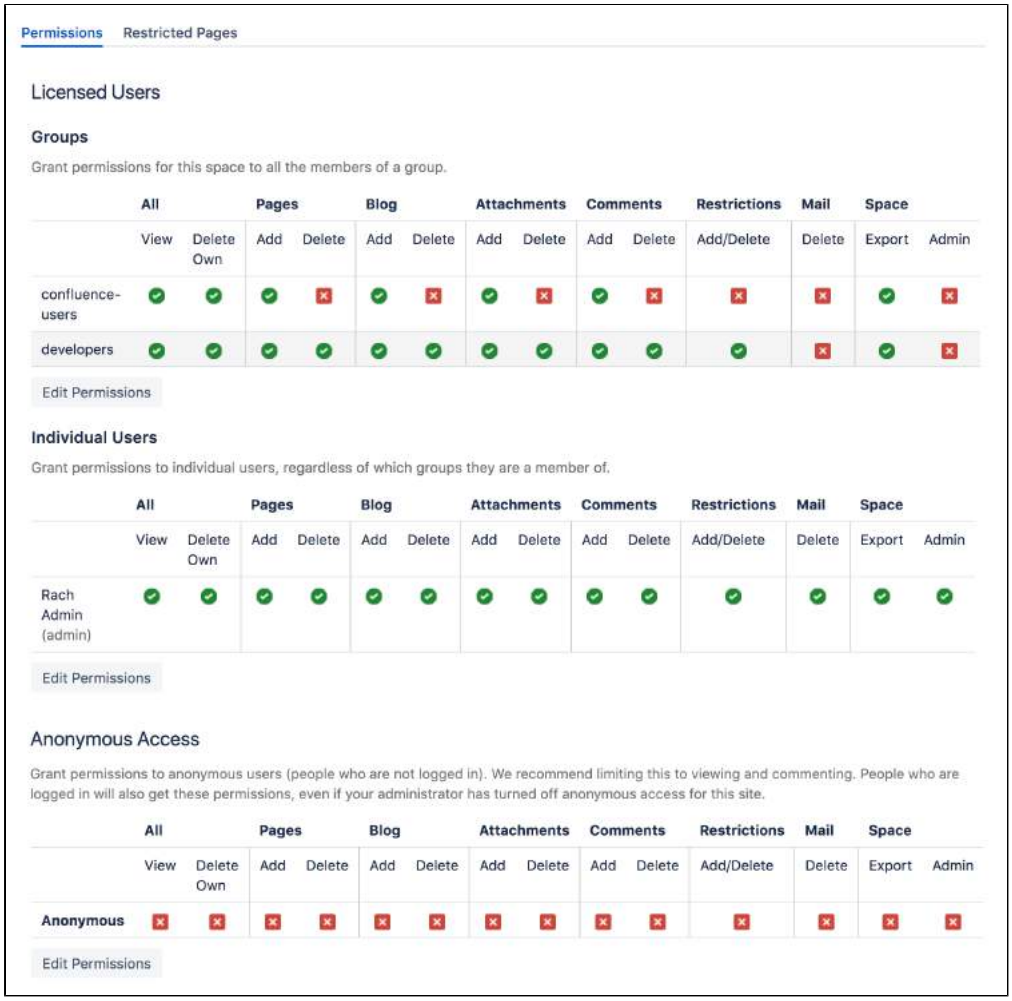

Grant space permission

To add a new user or group to the permissions list:

1. Search for either a group or userin their respective sections and choose**Add**.The group or user will appear in the list.
2. Select the specific permissions you'd like to apply then choose**Save all**.

You can bulk assign or revoke permissions by selecting**Select All**or**Deselect All**.

Permissions are managed on a space by space basis. Your Confluence Administrator is able to set default space permissions, which will apply to any new spaces.

Set default space permission for all new spaces

If you regularly need to grant permissions to the same groups each time a new space is created, you should consider updating the default space permissions.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

To set the default space permissions:

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

Default permissions are configurable forgroupsonly, not for individual or anonymous users.

Revoke 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.

Inspect permissions

If you need to troubleshoot why someone can or can't do something in your space, and you have a Data Center license, you can inspect permissions. SeeInspect permissions for more information.

Manage and recover space admin permissions

If you're aConfluence Administrator you can recover permissions to a space. This is useful when the only person with Space Admin permissions to a space leaves your organisation, for example.

To recover Space Admin permissions:

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

You can then choose**Manage Permissions**, and add any other appropriate permissions to the space. Requests to recover permissions are recorded in the Confluence audit log.

People with System Administrator permissions are able to manage permissions for all spaces, they do not need to first recover permissions.

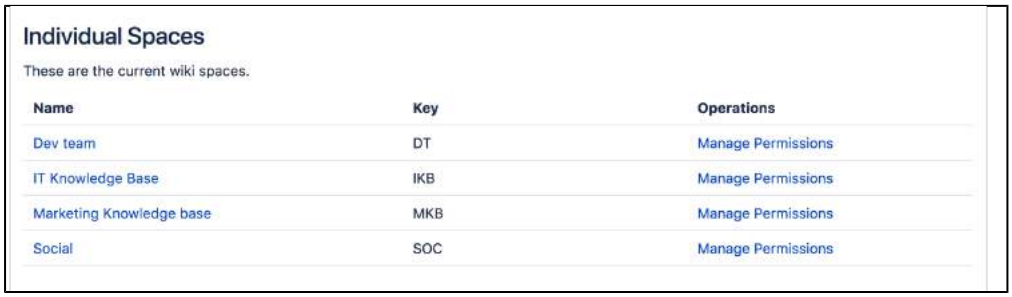

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Make a Space Public

If your site is public, and you want to share your space with people who are not logged in to Confluence (anonymous users), you can make your space public.

Making a space public does not let you choose who you want to share it with a public space can be viewed by anyone inside or outside of your organization.

Related pages:

- Give Access to Unlicensed Users from Jira Service Management

In order to make a space public, your administrator must first turn on the global permission for anonymous access.

This permission doesn't automatically grant anonymous users permission to see any of the spaces on your site, that is done on a space by space basis.

To make your space public:

1. Go to the space and choose**Space tools**>**Permissions**from the bottom of the sidebar.
2. Choose**Edit Permissions**.
3. Scroll down to the**Anonymous Access**section and select the specific permissions you'd like anonymous users to have.
4. **Save All**to apply the changes.

You can't grant space administration or page restriction rights to anonymous users. You can grant Delete Own, but it will have no effect, as we have no way of knowing who an anonymous user is.

What happens when the site is not public?

If your Confluence administrator turns off anonymous access to your site, users who are not logged in will no longer be able to see any spaces. However, all logged in users (regardless of their group membership) will be able to see all spaces that have granted space permissions to anonymous users.

Auditing considerations

There are some additional things to be aware of if you grant the**Add**page permissionto anonymous users.

You won't be alerted, when closing the editor or publishing a page, if the only unpublished changes on the page were made by anonymous users. This means a logged in user may inadvertently publish changes they were not aware had been made to the page.

The changes themselves are visible in the page, but the usual warning dialog will not appear if the only people to have made changes were not logged in.

If there are unpublished changes from both logged in users and anonymous users, the warning dialog will appear, but only the logged in users will be listed in the dialog. Changes made by all users (including anonymous) will be included if you view the changes from that dialog.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Give Access to Unlicensed Users from Jira Service Management

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

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

Related pages:

- Make a Space Public

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).

You canedit this permissionat any time to revoke access to a space, but it can only be re-enabled from Jira Service Management.

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.

SeeUse Jira applications and Confluence together for more information about Jira Service Management integration.

Organize your Space

Here's a few tips that'll help you organize your space so that everyone can find what they're looking for and stay on top of what's important to them.

On this page:

- How do I organize content within my space?
  - Pages and blogs
  - Configure the sidebar
  - Using labels
- How do I keep my space tidy?
  - Create a set of guidelines
  - Use page blueprints
  - Create from template macro
  - Create your own page templates
- How do I help my team stay on top of what's important?
  - My Spaces
  - Save for later
  - Watch a page, blog or space
  - @mentions

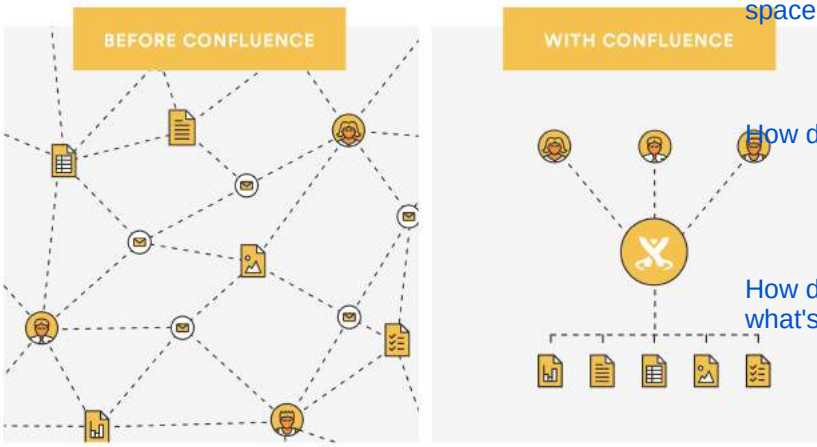

How do I organize content within my space?

Pages and blogs

Everything you create in Confluence, from meeting notes to retrospectives and everything in between, takes the form of either pages or blogs.

- Your homepage will be the first thing that visitors to your site see, so to help them find relevant content, start by curating your homepage with useful macros and including information about what is in your space. SeeSet up a Space Home Pagefor more information.
- If you're creating content that is specific mainly to the current timeframe, and isn't going to change over time, create it as a blog post. Your blog displays as an infinite scroll, so it surfaces the latest news and visitors just need to scroll down if they're interested in older content.
- If you're creating content that you want to last, and possibly evolve over time, then create it as a page. Pages nest, so every page can have its own child pages, which lets you organize your content into categories and subcategories.

Configure the sidebar

You can Configure the Sidebarto make it easier to navigate through your space.

The space shortcuts section of the sidebar lets you link to important content. You can use this to highlight pertinent pages within your space, related content from other spaces, or to external content that is relevant to your space.

The navigation display lists the pages in your space in either a page tree or child pages format. If you only want some content to be visible in the sidebar, you can hide the navigation display and put the pages you want to remain visible under Space shortcuts instead.

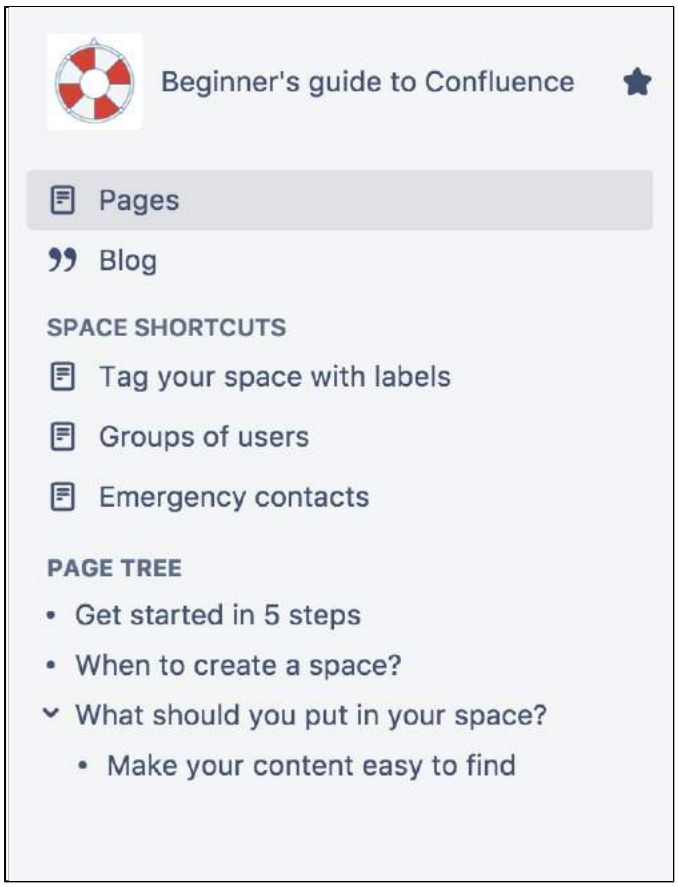

The page tree in the sidebar shows the 200 pages closest to where you are. Hit Show all pages, if you want to see all the pages in a space.

Using labels

Labels are keywords or tags that you can add to pages, blog posts, and attachments.

- Define your own labels and use them to categorize, identify, or bookmark content in Confluence. For example, if you assign the label 'accounting' to all accounts related pages on your site, you'll then be able to:
  - Browse all pages with that label in a single space or across the site.
  - Display a list of pages with that label.
  - Search based on that label.
- Use theContent by Label Macroto create a table of contents for your space that is organized by label categories.
- Labels aren't exclusive, so you can have as many labels as you want on a page. The page will then appear under each of those categories. See Use Labels to Categorize Spacesfor more information.

How do I keep my space tidy?

If you have lots of people creating in the same space, things can get messy fast. You can prevent this by taking a few easy steps.

Create a set of guidelines

- Let your collaborators know about what parent pages to create their child pages under, so no content gets lost or misplaced.
- Decide on standard labels to add to pages, blogs, and attachments so all content gets neatly categorized.
- Add a link to this in the **Space Shortcuts** section of the sidebar so that it's easy to find.

Use page blueprints

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Blueprintsare templates that come with formatting, macros and sample content. You can customize these Blueprintsfor each space. Everything created from a Blueprint will have it's own index in the sidebar, so for example if you use the Meeting Notes Blueprint, you can select 'Meeting Notes' in the sidebar to see a list of all the meeting notes in your space.

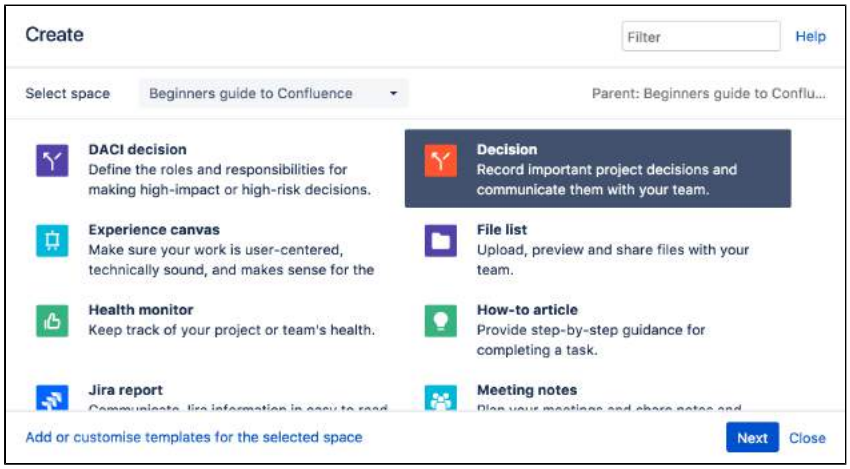

Create from template macro

Make things simpler for other contributors by using the Create from Template Macro. The Create from Template Macro lets you put a button on a page that links to a specific template of your choice. When the button is clicked, the macro opens the editor, ready to add a new page, and adds content to that page based on the given template.

Create your own page templates

Createyour owntemplatesfor any content that you want formatted the same way every time. For example, if you have to create a regular report tracking the same criteria, create a template with headings, variable dates, tables, and spaces for any graphics, so that each time all you have to do is input the new data instead of creating the whole report from scratch.

How do I help my team stay on top of what's important?

If you've got a lot of content on Confluence, staying on top of everything may seem a little daunting but these features will help your team save and track all the content they care about.

My Spaces

Add any spaces that you want to be able to navigate to easily to your list of 'My Spaces'. This list can be found under your dashboard and in the Space Directory, and you can also use theSpaces List Macroto display it on a page or blog.

To add a space to your list of 'My Spaces', either navigate to that space or find it under the Space Directory, and select the star icon next to the Space Name. To remove it from the list, just select the star icon again.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

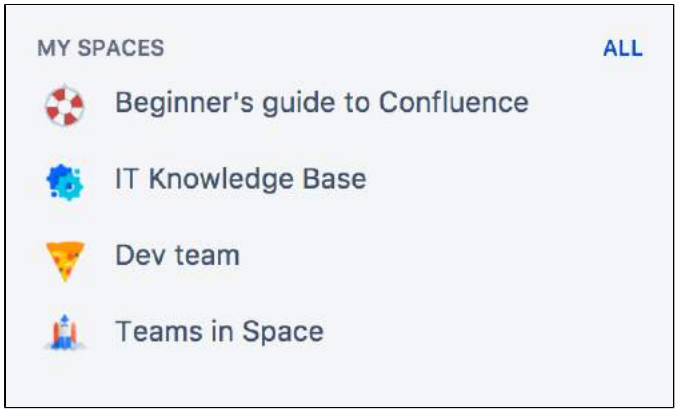

Save for later

If you only want links to certain pages or blogs rather than a whole space, you can choose Save for later and these will appear on your dashboard and under your profile. You can use the Favorite Pages Macroto display a list of all of everything you've saved for later.

Watch a page, blog or space

If you want to keep track of all the changes made to a page, blog, or space, you can also watchthem. Watching any content means that you will receive email notifications for all edits, deletions, attachments or comments made to that content.

- To watch a page, navigate to the page you want to watch, then choose **Watch** > **Watch page** , or if you want to watch the whole space, select **Watch all content in this space**.
- To watch a blog, navigate to that blog and choose **Watch this blog**.
- To stop watching something, deselect the relevant checkbox.

You can also manage watchers for your own space. This is useful when, for example, you're creating a new project, and want the team members on that project to stay notified of its progress. Go to any page in that space and choose **Watch** > **Manage Watchers**, then add or delete any names under 'Watching this space'.

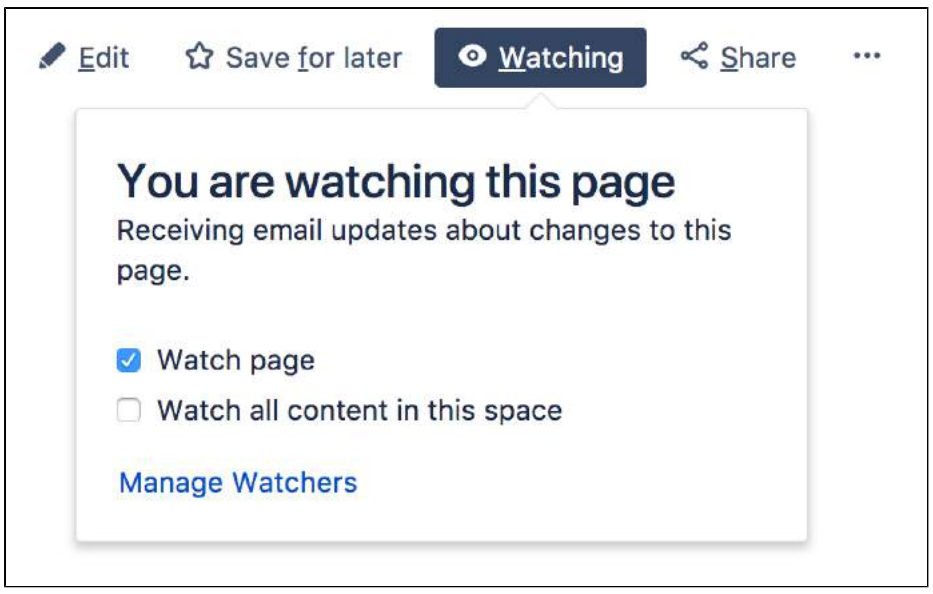

@mentions

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Use @mentions for any work where you need someone else's input or want to assign someone a task. Mentioning someone works like a tag; they'll immediately receive a notification that they've been mentioned, and can click through to that page or blog. If you mention someone when creating a task, it'll assign that task to them and they'll also be able to find it under their profile.

You can use this in place of emails if you want someone to look something over, add in additional information, or approve anything, simply put that work on Confluence and assign it to them as a task. They'll be able to make any changes or comments within Confluence and let you know when they're done by mentioning you back.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Set up a Space Home Page

When you create a space, Confluence automatically creates a home page for that space. Spaces created with a blueprint come with a ready-made home page populated with useful macros and sample content specific to the blueprint's use case.

However, even if you've started with a blank space, you can still customize your home page to make it easier for everyone using that space to navigate their way around and find useful content.

On this page:

- Create a kickass home page
- Top macros for common types of spaces
- Set up your personal space home page
- Set another page as your home page

Related pages:

- Organize your Space

Create a kickass home page

- **What is this space about?**
  Your home page is the first page visitors will see when they visit your space, so it helps to include some information about what the space is about, and what you're working on. You can use the Info, Tip, Note, and Warning Macrosto create a colored box in which to put this information.

- **What is in this space?**
  You can use the Content Report Table Macroto create a table of contents for your space. You can also set this macro to only display content with a particular label, so if you would like only some content to display, you can do this by adding that label to only those pages and blogs you want listed on your home page.

- **Organize your space with labels**
  You can organize content in your space with labels, so that for example, if you have a Learning and Development space, you can create different labels for online learning resources, upcoming workshops, and training strategy. You can then use the Labels List Macroto create a list of those labels to make it easy for visitors to your space to find the content relating to each of those topics.

- **Add a search box so that it's easy to find content within your space**
  The Livesearch Macroallows you to add a search box to a Confluence page, and you can set it to only find content within your space.

- **Keep everyone updated about the latest changes within your space**
  If it's important for your visitors to know about the latest changes to your space, you can use the Recently Updated Macroto display a list of the most recently updated content. You can set the space parameter to show this for just your space, or, if you have related spaces, to show the most recently updated content across all of those spaces as well.

- **Using Jira? Create and display your Jira issues on Confluence**
  If your Confluence site is connected to a Jira application, you can both create and display your Jira issues without having to leave Confluence using the Jira Issues Macro. You can connect Confluence to any Jira application, including Jira Software and Jira Service Management.

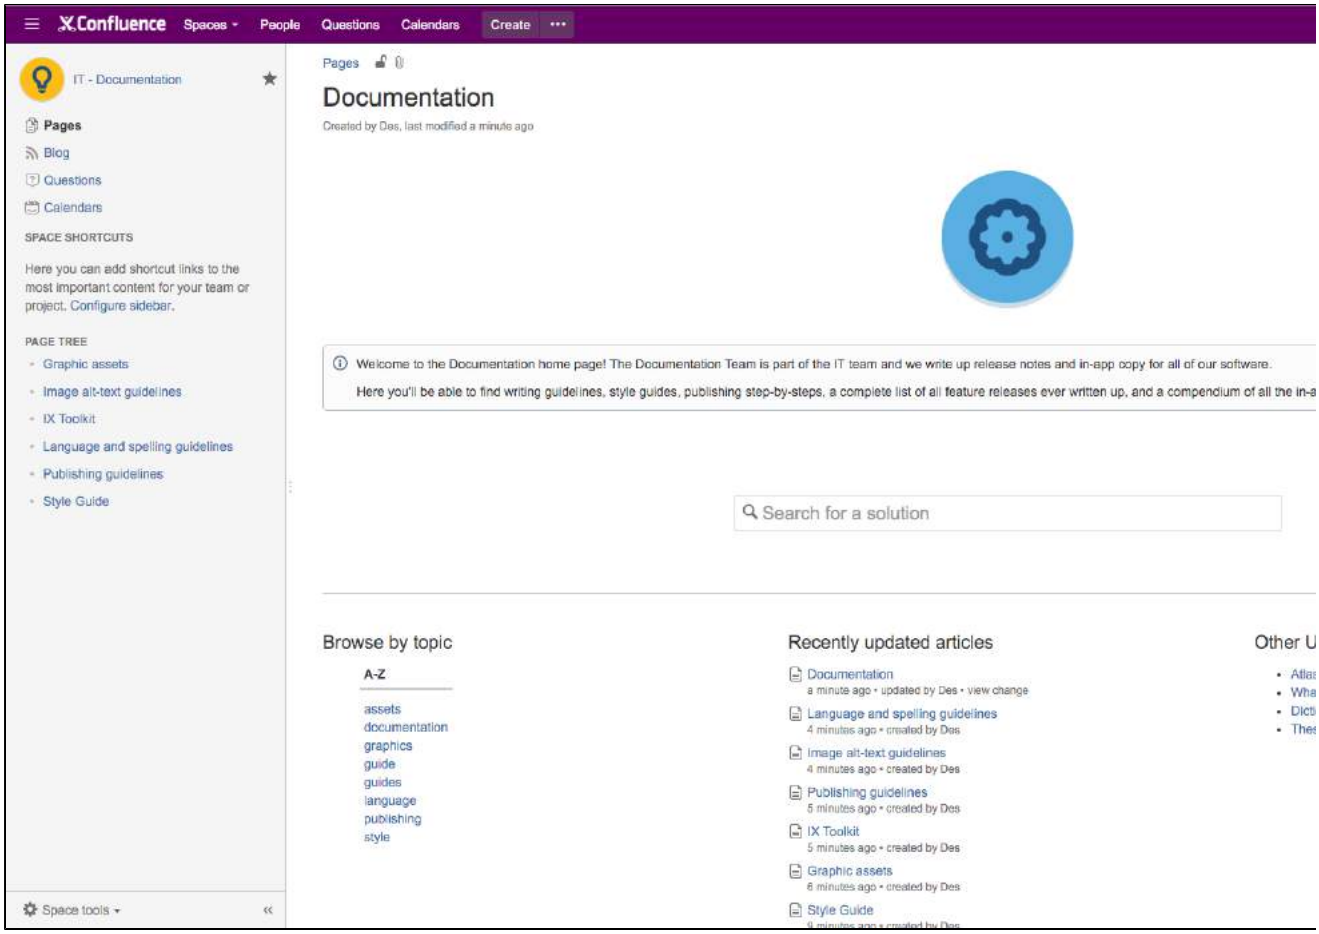

Top macros for common types of spaces

Team Spaces:

- **Introduce the team:** The User Profile Macrodisplays a short summary of a given Confluence user's profile with their role, profile photo and contact details.
- **Share news and announcements with your team:** The Blog Posts Macrodisplays a stream of your latest blog posts so your team can easily see what's been going on.

Knowledge Base:

- **Have external content that you need on your page?** Embed online videos, slideshows, photo streams, and more, directly into your page with the Widget Connector Macro.
- **Put your own multimedia content onto the page:** The Multimedia Macroembeds attached video, animation, and other multimedia files on a Confluence page.
- **Create an index of all your content:** The Page Index Macrocreates a hyperlinked alphabetical index of all page titles within the current space.

Planning/Project:

- **Keep track of everyone's tasks:** Use the Task Report Macroto display a list of tasks on a page. Filter the tasks by space, page, user, label, created date and more.
- **Is your project on track?** The Status Macrodisplays a colored lozenge (a rounded box) that is useful for reporting project status. You can choose the color of the lozenge and the text that appears inside the lozenge.
- **Let everyone see where you're going:** The Roadmap Planner Macrocreates simple, visual timelines that are useful for planning projects, software releases and much more.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Using Macros

1. From the editor toolbar, choose **Insert** + > **Other Macros**
2. Find and select the required macro

Speed it up with autocomplete: Type **{** and the beginning of the macro name, to see a list of suggested macros. In this example we're inserting the cheese macro.

To edit an existing macro: Click the macro placeholder and choose **Edit**. This will open the macro details, so you can edit the macro parameters.

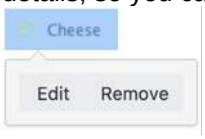

Set up your personal space home page

Whether you're using your personal space as a sandbox to draft and test content, a portfolio to show off what you're working on, a home base to navigate to your content in other spaces, or something completely different, these are some macros that should help you use your space more effectively.

- Use the Favorite Pages Macroto create a list on your home page of all the pages you've saved for later, so you can easily navigate back to any of them.
- Use the Content by User Macroto keep track of all the current pages, comments and spaces you've created so you can find everything you've been working on in one place.
- Use the Task Report Macroto keep track of all tasks assigned to you, and tick them off as you finish them.
- Use the Recently Updated Dashboard Macroto keep track of all the content across your Confluence site that you're interested in the Dashboard lets you choose which spaces, users, blogs, pages or files you would like to keep updated about.

Set another page as your home page

If at any point you decide that you would like another page within your space to be your home page, you can easily change this from the Edit Space Details tab.

To edit a space's details:

1. Go to the space and choose**Space tools**>**Overview**from the bottom of the sidebar.
2. Choose **Edit Space Details**.
3. Enter the page you want use in the **Home page** field then choose **Save**.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

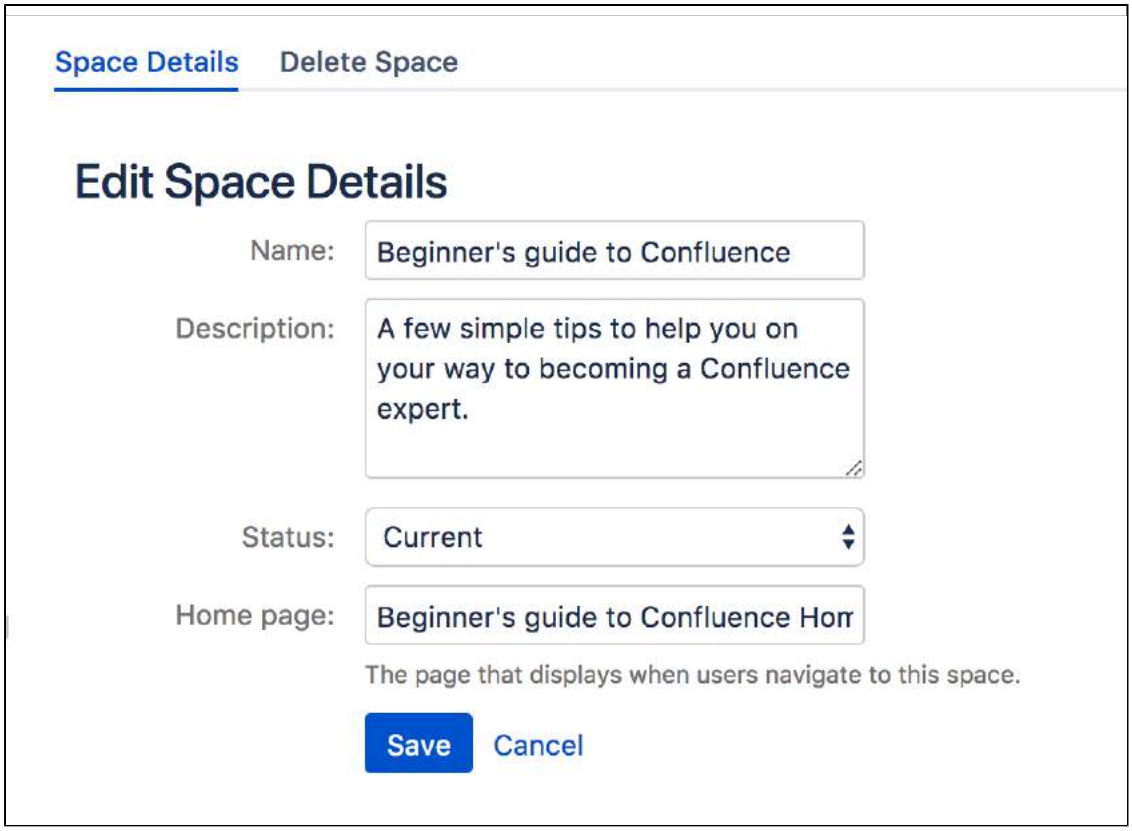

⚠️You can change the home page, name and description of your space, but you are not able to change the space key.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Use Labels to Categorize Spaces

If you've got lots of related spaces, you can use labels to group them together into categories in the**Space Directory**.

For example, if you're in an IT team who work on a number of projects, each in a different space, you could label each space 'IT-projects-open'. Then in the Space Directory you could select IT-projects-open to see all your current IT project spaces.

You can add as many space categories to a space as you need, so that if, for example, two different teams are working on a project together, you can add labels for both teams and space will appear under both categories.

Labels are easy to add or remove, so if your needs change, you can always recategorize your spaces.

On this page:

- Categorize a space
- View spaces in a category
- Remove a space from a category
- Search within a space category

Related pages:

- Add, Remove and Search for Labels

Categorize a space

You need space administrator permissions to add categories to 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**.

Add a space description

Help make it easier to find the right space within a category by adding a description to your space:

1. Go to the space and choose**Space tools**>**Overview**from the bottom of the sidebar.
2. Choose **Edit Space Details**.
3. Under the **Description** field, type a short description to tell visitors what your space is about, then choose **Save**.

View spaces in a category

To see what spaces are in a category, choose**Spaces** > **Space directory**in the Confluence header, then select one of the categories from the list.

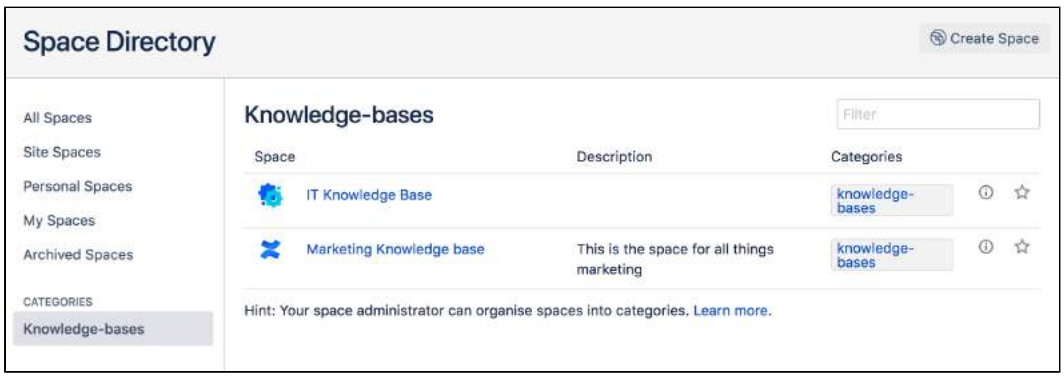

You can also view spaces by category byusingthe Spaces List Macro and filtering by category.

Head to **My Spaces** in the **Space Directory** to see all your favorite spaces. When viewing a space, you can choose the star icon next to the space title in the sidebar to add it to **My Spaces** so that it's easier to find later.

Remove a space from a category

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**, choose the x icon next to each category that you want to remove.

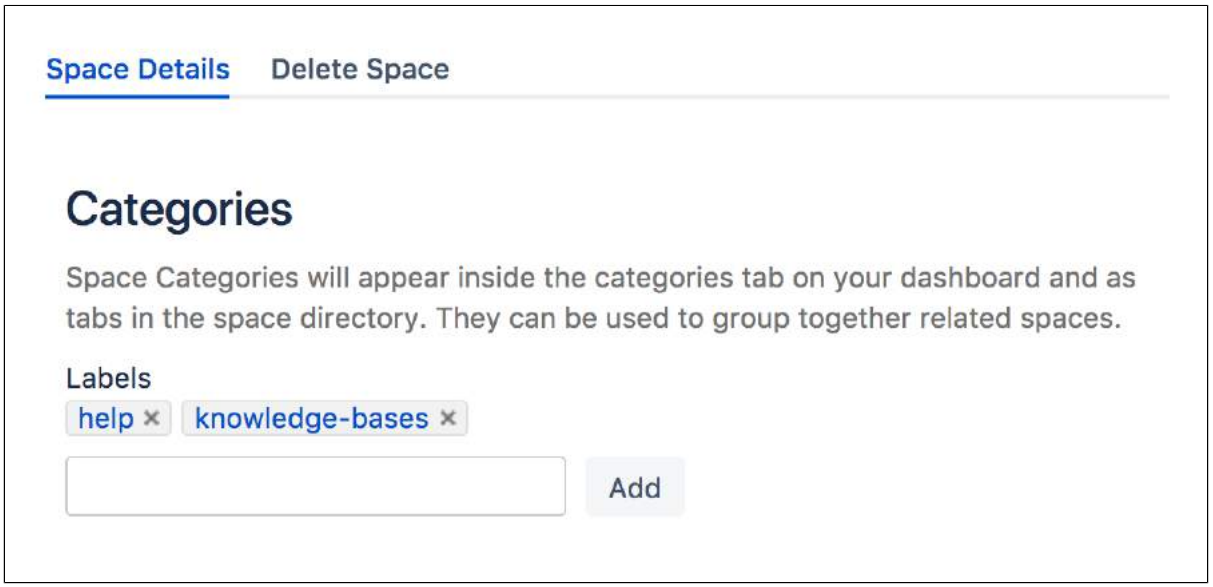

⚠If you remove all spaces from a category, the category will no longer appear in the Space Directory. There's no way to bulk remove a category, but you can choose the category in the Space Directory to find all the spaces it appears on, and then remove it from each space.

Search within a space category

You can search for content within a specific space category using the search filter.

To search within a Confluence space category:

1. Click the search field in the top-right of Confluence to open the search panel.
2. Click the **Space category** filter on the left.
3. Start typing the category name and choose from the list of possible matches.

Learn more about searching Confluence.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Customize your Space

Make your space stand out from the crowd by customizing its appearance.

If you have space adminprivileges, you can change the color scheme for your space, add your own space logo,choose what shows up in your space's sidebar,oruse Atlassian Marketplace themes to change the whole look of your space.

- Configure the Sidebar
- Edit a Space's Color Scheme
- Apply a Theme to a Space
- Customize Space Layouts
- Change the Space Logo

Related pages:

- Changing the Look and Feel of Confluence
- Styling Confluence with CSS

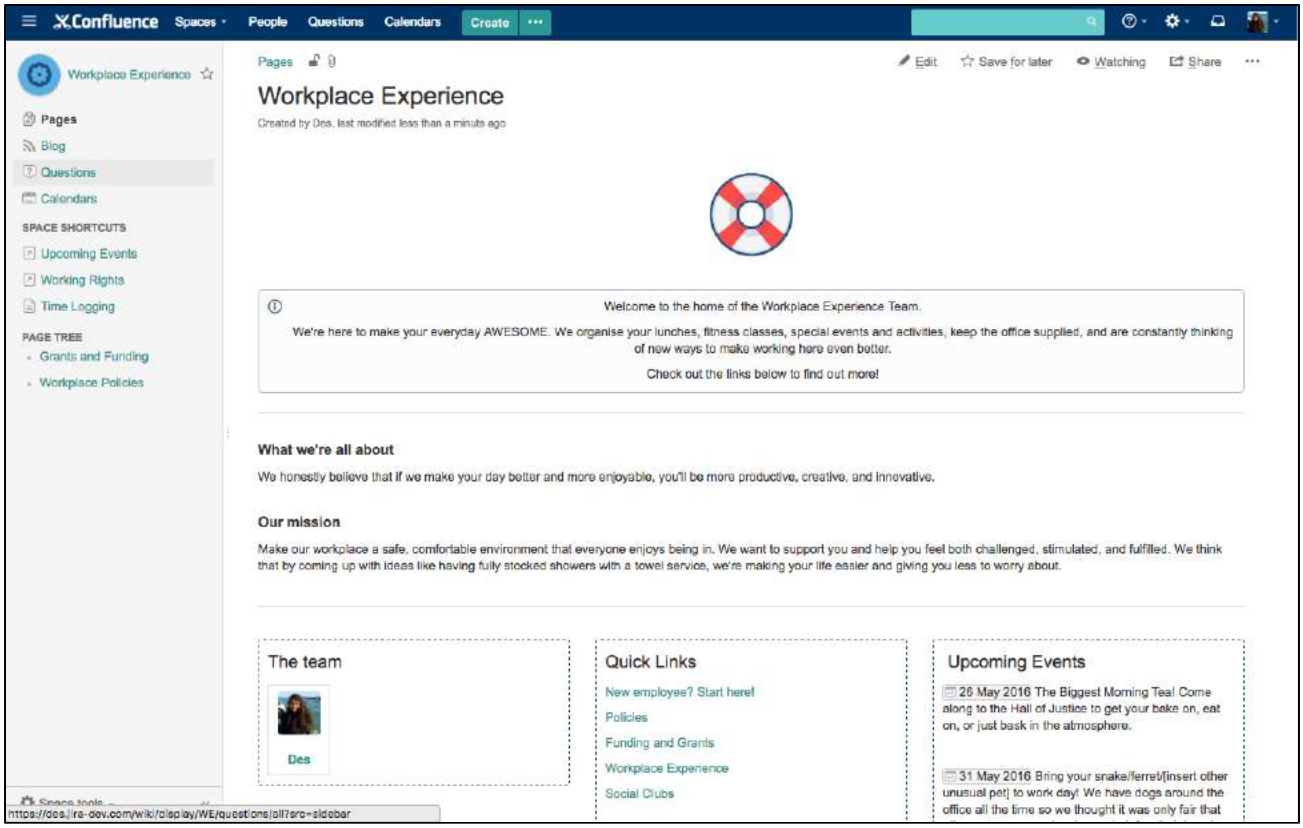

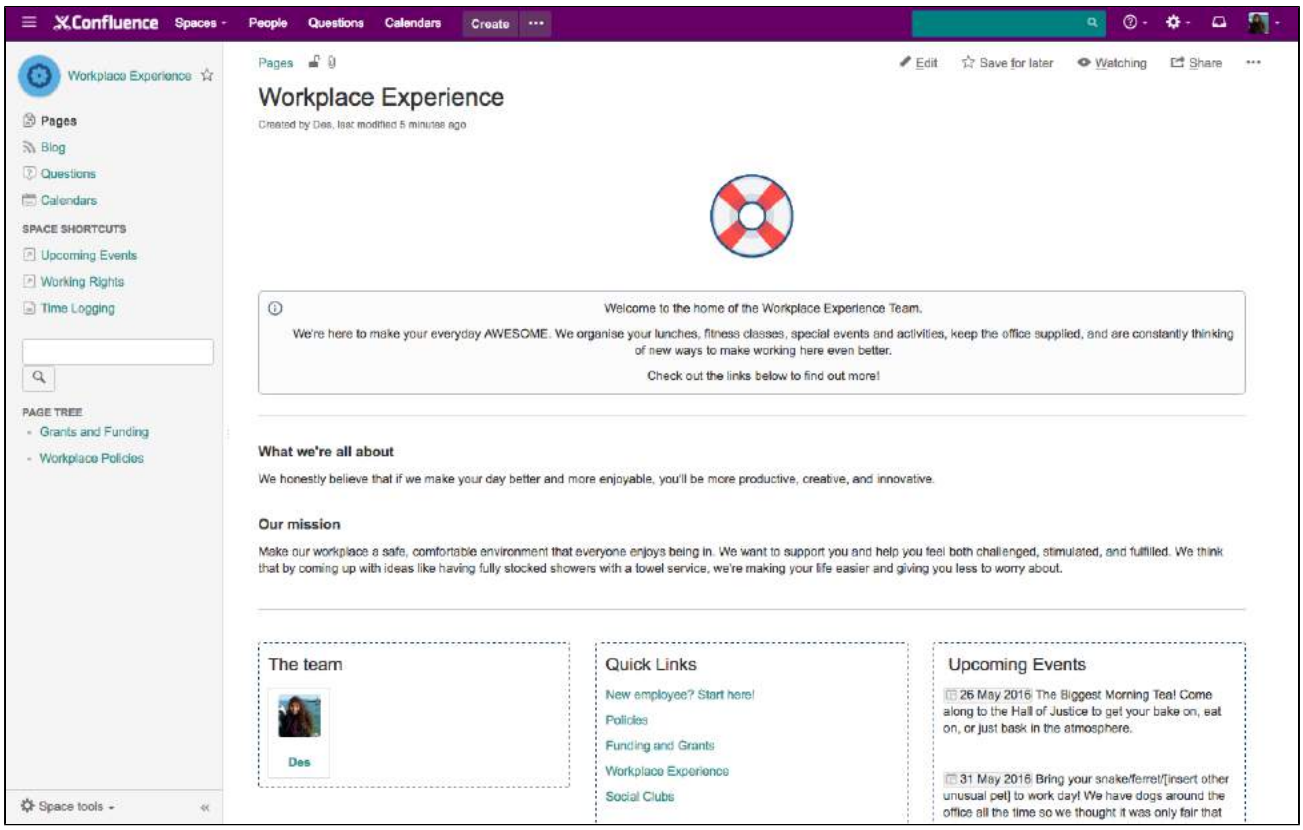

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Configure the Sidebar

If you have administration permissions for a space, you can customize the space's sidebar to have its own logo, change the way the hierarchy is displayed, and add shortcut links to help navigate to important content quickly.

To start configuring the sidebar, choose **Space tools** > **Configure sidebar**.

On this page:

- Change the space name and logo
- Configure the sidebar links
- Change the navigation display options
- Adding custom content to a sidebar

Related pages:

- Edit a Space's Color Scheme
- Organize your Space

Change the space name and logo

To change the space name:

1. Choose next to the space name.
2. Type in a new space name and**Save**.

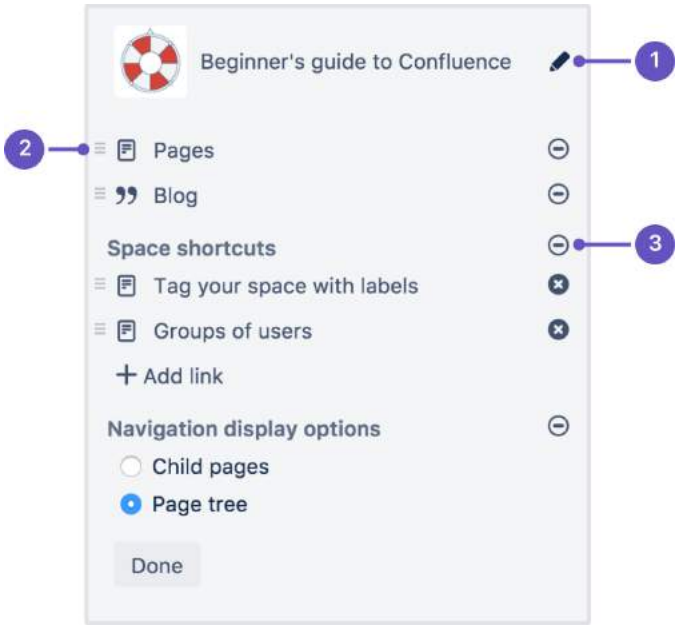

1. **Logo** - change the space name and logo
2. **Reorder** - drag to reorder shortcuts
3. **Hide** - hide or show pages, blogs, shortcuts or navigation display options.

To change the space logo:

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

Things you should know:

- Space logos are 48px x 48px. Logos smaller than these dimensions will be centred with whitespace around them.
- You can only change the space logo for a site space. For your personal space, your profile picture is used as the space icon.

Configure the sidebar links

- Choose the icons to show or hide **pages**, **blogs**, **shortcuts** or **navigation options**.
  For example, if you want your space to be used primarily as a blog you can hide the 'Pages' link.
- Add-ons such as Team Calendars for Confluence Server may add other links in this section of the sidebar and you can also show or hide these.
- Choose **Add link** to add a shortcut link to the sidebar. This could be a link to an important page for your team, or to an external site.
- Drag links to reorder them within each section (you can't move a link from one section to another).

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.
- You can also choose to completely hide the navigation display options and instead add the pages you want to be visible as shortcut links. Alternatively, you can remove the sidebar navigation altogether and use a Page Tree Macro or Children Display Macro on your homepage for navigation instead.

Adding custom content to a sidebar

You can further customize a sidebar using wiki markup to add custom content.

To add custom content to your sidebar:

1. Go to the space and choose **Space tools** > **Look and Feel** from the bottom of the sidebar.
2. Choose **Sidebar, Header and Footer**.
3. Enter your custom content in the **Sidebar** field.

The sidebar, header and footer fields all use wiki markup, check our guide to wiki markup for help, or check out some common customizations below.

To add a search field to the sidebar add the following wiki markup for a search macro in the **Sidebar** field.

For the Livesearch Macro macro:

```
{livesearch}
```

For the Page Tree Search Macro macro:

```
{pagetreesearch}
```

To add a panel with some custom content to the sidebar add the following wiki markup for the Panel Macro in the **Sidebar** field:

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

To hide the default page tree and add your own, with additional parameters:

1. Add the wiki markup for the Page Tree Macro in the **Sidebar** field.
   The following example includes parameters to expand the top three levels of the page tree by default and include an **Expand All** and **Collapse All** link above the tree.

   ```
   {pagetree:root=Page Name|startDepth=3|expandCollapseAll=true}
   ```

2. Go to **Space Tools** > **Configure sidebar**.
3. Use the **Show** and **Hide** icons to hide the default page tree from the sidebar.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Edit a Space's Color Scheme

Spaces inherit the global color scheme by default, but if you have admin permissions for a space, you can jazz it up with your very own customizable color scheme.

To change the color scheme for a space:

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

Related pages:

- Configure the Sidebar
- Apply a Theme to a Space

Customizable Elements

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

Some UI elements below are for specific themes, and color 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 color of top navigation bar menu items when selected (e.g. spaces)
- **Top Bar Menu Selected Text**- text color 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 color on top navigation bar drop down menu items
- **Menu Item Selected Text**- text color on highlighted top navigation bar drop down menu items
- **Search Field Background** - the background color of the search field on the header
- **Search Field Text** - the color of the text in the search field on the header
- **Page Menu Selected Background**-the background color 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
- **Links**- all links throughout the space
- **Borders and Dividers**- table borders and dividing lines

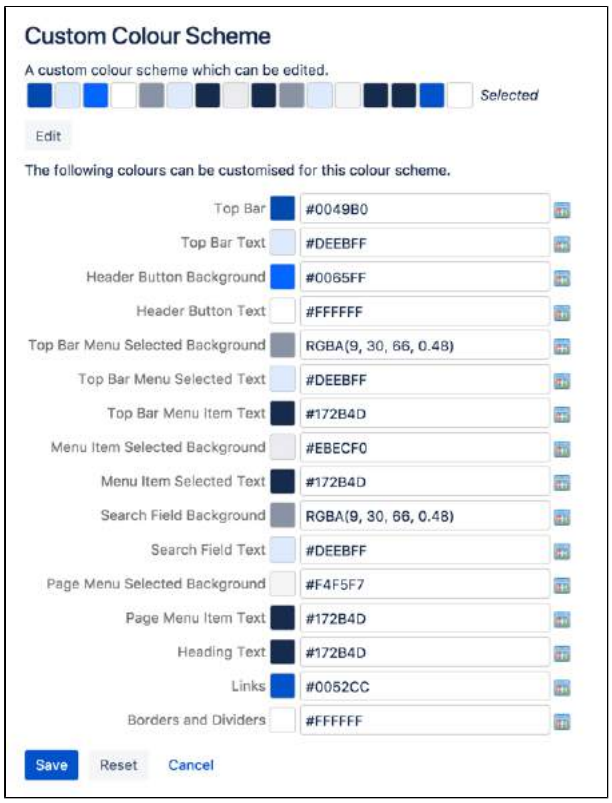

Handy Hint

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

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Apply a Theme to a Space

Themes are used to change the appearance of your space. Your Administrator can download and install themes fromThe Atlassian Marketplace.

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

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 needSpace Admin permissionsto do this.
2. Choose**Themes**and select a theme option.
3. Choose**Confirm**.

Related pages:

- Applying a Theme to a Site
- Customize your Space
- Edit a Space's Color Scheme
- Configure the Sidebar

What is the global look and feel?

When a new space is created, whichever theme is applied to the whole site will be applied by default to the new space. This is the global look and feel, and any changes made globally will flow through to all spaces that inherit the global look and feel.

If a space has its own theme applied, or if changes have been made to customize the look and feel of the space, it will no longer inherit changes from the global look and feel.

If you want to go back to inheriting the global look and feel choose**Global look and feel**from the**Themes**page.

Documentation theme migration FAQ

Aspreviously announced, the documentation theme is not available in Confluence 6.0. We know you'll have a lot of specific questions, so we've created this FAQ to help you prepare for upgrading to Confluence 6.0.

If you have further questions, you can ask them at the bottom of the page and we'll do our best to provide an answer.

- What does the default theme look like?
- How can I check if I am using the documentation theme in my space?
- How can I check if the documentation theme is being used anywhere in my site?
- What will happen to my theme customizations during the upgrade?
- Will my space break after the upgrade?
- Can I add custom content to the sidebar, header and footer globally?
- Can I still use macros in my sidebar, header or footer?
- Where should I add custom content to the sidebar, header or footer?
- How do I turn off the Pages and Blogs shortcuts at the top of the sidebar?
- Can I edit the default theme's sidebar globally?
- I want to see the page tree, not child pages. how do I do this?
- Where is Space Administration and Space Operations?
- Do I have to use wiki markup in the sidebar, header, and footer?
- Can I hide or change the space logo appearance?
- Can I hide the built in page tree and insert my own?
- How can I make my page titles wrap in the sidebar?
- Can I change the order in which things appear in the sidebar?
- Can I still use the space jump macro?
- Why don't child pages appear below the page?
- Can I restrict search to just this space?
- How do I view the pages in my space alphabetically?
- What will happen if I import a space that uses the documentation theme?

What does the default theme look like?

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

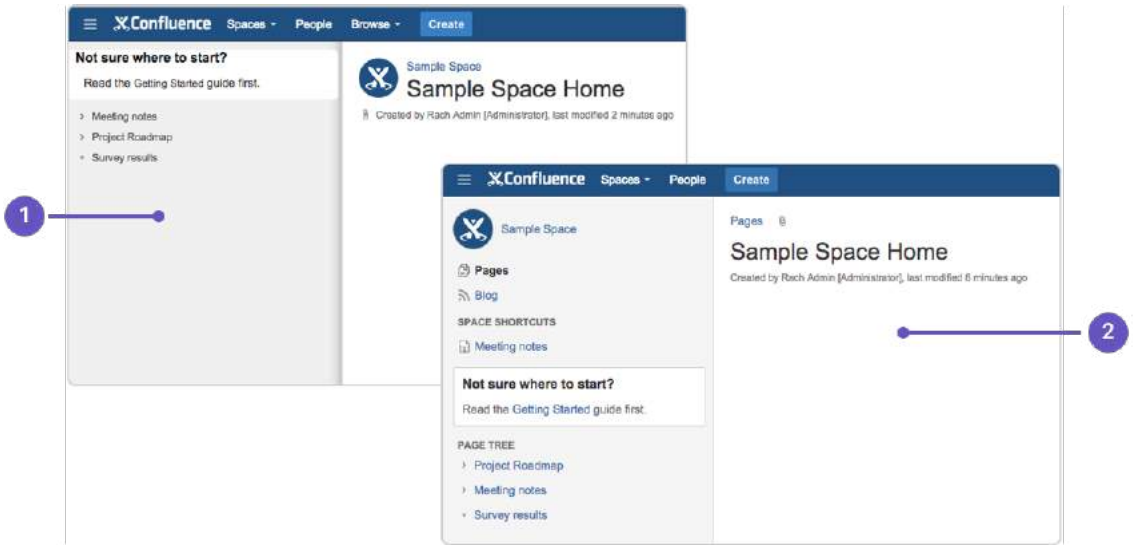

1. **Documentation theme**: with custom sidebar content.
2. **Default theme**: with the same custom content in the sidebar.

How can I check if I am using the documentation theme in my space?

The easiest way to check whether your space is using the documentation theme is to look for a**Browse** menu in the header, near the **Create** button. (If you're using the default theme already, you'll see a Space Tools menu at the bottom of the sidebar instead.)

How can I check if the documentation theme is being used anywhere in my site?

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).

What will happen to my theme customizations during the upgrade?

During the upgrade we'll automatically turn on the default theme for any spaces that currently use the documentation theme. If you've customized the documentation theme (by adding wiki markup to the sidebar, header or footer) we'll take this wiki markup and drop it into the sidebar, header and footer in the default theme.

The default theme adds some new sections to the sidebar, such as links to pages, blogs and space shortcuts. You can choose to hide these - head to **Space Tools** > **Configure Sidebar** and use the button to hide any items you don't want to see.

Will my space break after the upgrade?

This depends on the amount of customization you have. In most cases your space sidebar may look a little different but the changes shouldn't be dramatic.

If you've used CSS to change the appearance of your space (either in the space stylesheet or by using the Adaptavist Content Formatting macros like{style} and {div} in the sidebar, header, or footer of the documentation theme), you may need to make a few changes to some class names in your CSS to get your space looking right. For example, if you specified`#splitter-sidebar`in the doc theme, you'll need to use`.acs-side-bar`for the default theme.

If you have customized default theme layouts through the Confluence UI, you may find that your space looks strange or broken when the default theme is re-applied to spaces previously using the documentation theme.

If you experience problems, you'll need to reset the broken layouts.

> This method will only work if you have more than one theme available in your site. You'll need System Administrator global permission to do this.

1. Switch to another 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. In the space administration options go to **Layouts** (if available) or use the following link, replacing `YOURSPACEKEY` with the space key for the space.

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


Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

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

Layouts can also be customized for the entire site - head to > **General Configuration** > **Layouts** if you need to reset the layout for the entire site.

If you're unable to reset the layouts via the Confluence UI, you can remove the affected layouts directly in the database. Be sure to take a full database backup before you try this.

First, use the following query to identify customized layouts:

```
Select *
FROM DECORATOR
ORDER BY SPACEKEY
```

Then, you can selectively remove records for spaces that are affected.

Can I add custom content to the sidebar, header and footer globally?

Yes. Head to > **General Configuration**>**Sidebar, Header and Footer**. All spaces that use the global look and feel will inherit your custom content. Any custom content added to the sidebar, header and footer in a particular space will override any custom content added globally.

Can I still use macros in my sidebar, header or footer?

Yes! If a macro worked correctly in the documentation theme it'll work in the default theme too.

Where should I add custom content to the sidebar, header or footer?

You can add custom content to the sidebar, header and footer in each space individually (**Space Tools** > **Look and Feel** > **Sidebar, Header and Footer**) or globally (> **General Configuration**>**Sidebar, Header and Footer**).

Confluence displays global custom content in all spaces, except where a space has its own custom content defined (space custom content overrides global custom content). This behavior applies field by field, so a space can display a combination of custom content. For example you could define the content of a footer globally, and content of a header in each space individually, or only in some spaces.

How do I turn off the Pages and Blogs shortcuts at the top of the sidebar?

Go to**Space Tools**>**Configure Sidebar**and use the icons to hide any items you don't want to see.

Can I edit the default theme's sidebar globally?

No. You can add custom content to the sidebar globally, but showing and hiding sections of the sidebar, setting space logos, and adding shortcut links are done on a space by space basis.

I want to see the page tree, not child pages. how do I do this?

Head to**Space Tools**>**Configure Sidebar**and select**Page Tree**in the navigation options. The default for all new spaces is Page Tree.

Where is Space Administration and Space Operations?
Instead of choosing between Space Operations and Space Administration, the documentation theme has a single **Space Tools** menu that lets you jump right to Permissions, Content Tools, Look and Feel, or Integrations.

Old Space Operations location	New Location in the sidebar
Browse > Pages	Pages on the sidebar or Space Tools > Reorder pages

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Browse > Blogs	Blogs on the sidebar
Browse > Labels	Click a label on any page, then use the menu to see all labels, popular labels or pages from other spaces
Browse > Space Operations > Space Details	Space Tools > Overview
Browse > Space Operations > Orphaned Pages	Space Tools > Content Tools > Orphaned Pages
Browse > Space Operations > Undefined Pages	Space Tools > Content Tools > Undefined Pages
Browse > Space Operations > Attachments	Space Tools > Content Tools > Attachments
Browse > Space Operations > PDF, HTML, XML Export	Space Tools > Content Tools > Export
Browse > Space Operations > RSS Feeds	Space Tools > Content Tools > RSS Feeds
Browse > Space Operations > Watch this space	Pages > Watch this space (or use the Watch button on any page)
Browse > Space Operations > Watch this blog	Blogs > Watch this blog (or use the Watch button on any blog post)
Browse > Space Operations > Remove from My Spaces	Use the icon in the sidebar (or in the space directory)
Browse > Space Admin > Space Details	Space Tools > Overview
Browse > Space Admin > Space Categories	Space Tools > Overview
Browse > Space Admin > Templates	Space Tools > Content Tools > Templates
Browse > Space Admin > Delete Space	Space Tools > Overview > Delete Space
Browse > Space Admin > Trash	Space Tools > Content Tools > Trash
Browse > Space Admin > Permissions	Space Tools > Permissions
Browse > Space Admin > Restricted Pages	Space Tools > Permissions > Restricted Pages
Browse > Space Admin > Application links	Space Tools > Integrations > Application Links
Browse > Space Admin > Themes	Space Tools > Look and Feel > Themes
Browse > Space Admin > Color Scheme	Space Tools > Look and Feel > Color Scheme
Browse > Space Admin > PDF Layout	Space Tools > Look and Feel > PDF Layout

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Browse > Space Admin > PDF Stylesheet	Space Tools > Look and Feel > PDF Stylesheet
Browse > Space Admin > Change Space Logo	Space Tools > Configure Sidebar then choose the edit icon on the space logo
Browse > Space Admin > Hipchat Integration	Space Tools > Integrations > Hipchat

Do I have to use wiki markup in the sidebar, header, and footer?

Yes. Our main focus when adding this feature was to help people move from the documentation theme to the default theme with a minimum of effort. Keeping these fields as wiki markup means that your existing customizations can be pasted straight in.

Can I hide or change the space logo appearance?

You can upload any image to use as your space logo, but you can't change how it appears in the sidebar (it's always round and always at the top).

Can I hide the built in page tree and insert my own?

Yes! If you want to have complete control over how the page tree appears in your sidebar you can hide the built in page tree, and then add aPage Tree macro`{pagetree}` in the sidebar custom content.

How can I make my page titles wrap in the sidebar?

Page titles do not wrap in the sidebar of the default theme (regardless of whether you're using the built in page tree or have added a `{{pagetree}}`macro as custom content). There's no way to change this.

Can I change the order in which things appear in the sidebar?

You can change the order of some items in the sidebar, such as the shortcuts, but the order of the sections themselves can't be changed. Anything that has a ☰ icon can be moved.

Custom content appears above the page tree. You have the option to hide the built in page tree, and then add it back in the custom content area using wiki markup. This can be useful if you want more control over the order of the page tree and your custom content.

Can I still use the space jump macro?

No, the space jump macro was provided by the documentation theme and will not be available once the documentation theme is removed. If you've used this macro on a page or in the header or footer of a space, it will show the following error after the upgrade `unknown macro: {spacejump}.`

To find out whether the Space Jump macro is used on any pages in your site, enter`macroName:spacejump` into the search bar. All pages containing the macro will be returned (it won't search the sidebar, header or footer unfortunately).

Why don't child pages appear below the page?

The default theme does not list child pages below the sidebar, but you can achieve a similar result by adding a C hildren Display macro to the footer.

Can I restrict search to just this space?

No, that is one of the features we removed with the documentation theme.

One workaround is to add a Livesearch macro to the sidebar or space homepage. Use`@selfin` the spaces parameter to restrict the search to the current space.

How do I view the pages in my space alphabetically?

The default theme does not have an option to view all pages in your space alphabetically.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

What will happen if I import a space that uses the documentation theme?

You should have no problems importing the space, but it will have the default theme applied and any wiki markup customization in the theme will not be automatically migrated to the default theme. Before exporting the space you should copy the wiki markup contents of the sidebar, header, and footer fields and keep it so that you can add it back in manually after you've successfully imported your space.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Customize Space Layouts

You can modify Confluence's look and feel by editing the layout files. This page tells you how to customize the layout files for a space. You'll need the system administratorglobal permissionplusspace administratorpermission for that space.

People with system administrator permissions can also customize the layout of the entire Confluence site. For more information, seeCustomizing Site and Space Layouts. Site layout customizations modify the default layout of all spaces in the Confluence site.

Any space layout customizations will override the equivalent site customizations.

Related pages:

- Customize your Space
- Apply a Theme to a Space
- Styling Confluence with CSS

If you modify the look and feel of Confluence by following these instructions, you'll need to update your customizations when you upgrade Confluence. The more dramatic the customizations are, the harder it'll be to reapply your changes when upgrading. Please take this into account before proceeding with any customizations.

For more information on updating your customizations, please refer to Upgrading Customized Site and Space Layouts .

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 personalize the appearance of Confluence.

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

- **Site layouts**: These are used to define the controls that surround each page in the site. For example, 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**

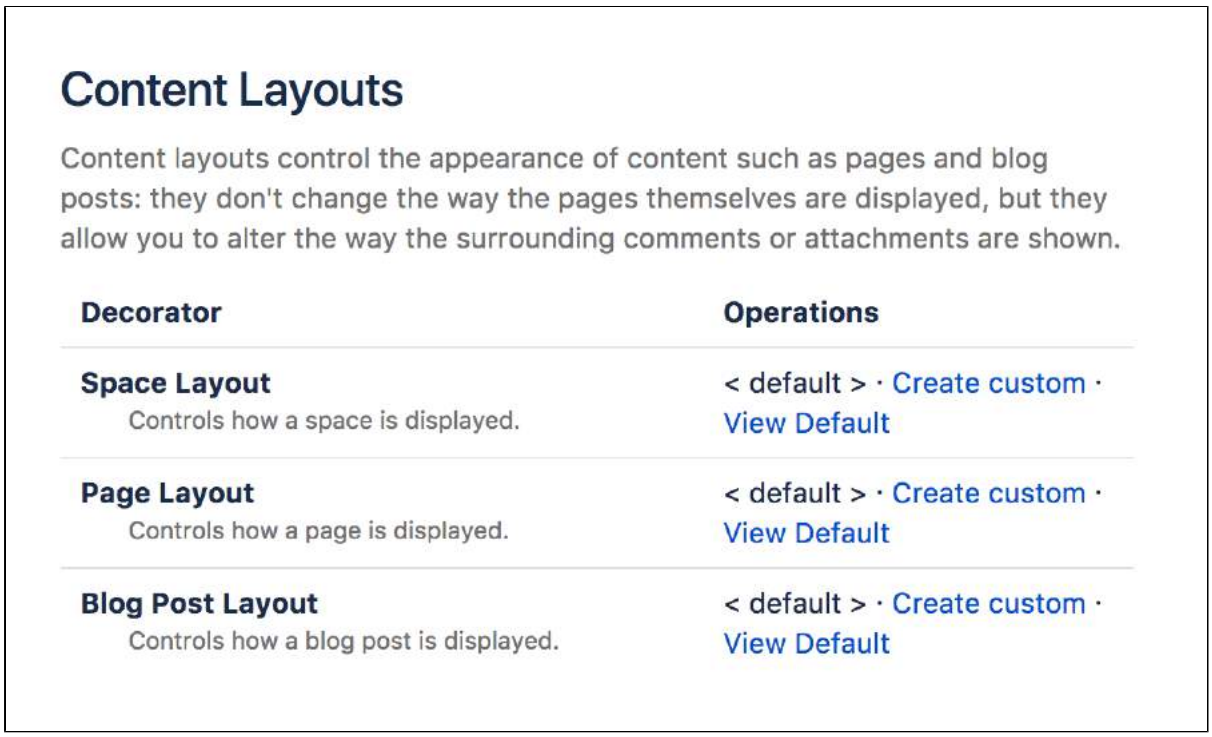

Content Layouts

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
Blog Post Layout Controls how a blog post is displayed.	< default > · Create custom · View Default

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

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 yourpersonal space, yourprofile pictureis used as the space icon.

You need to be aspace adminto replace a space's logo.

To change a space's logo:

1. Go to the space and select **Space tools** > **Configure sidebar** from the bottom of the sidebar
2. Select the edit icon next to the space name to edit space details
3. Select **Upload an image**
4. Add an image from your computer
5. Adjust the size of the image to fit within the highlighted circle
6. Select **Save**

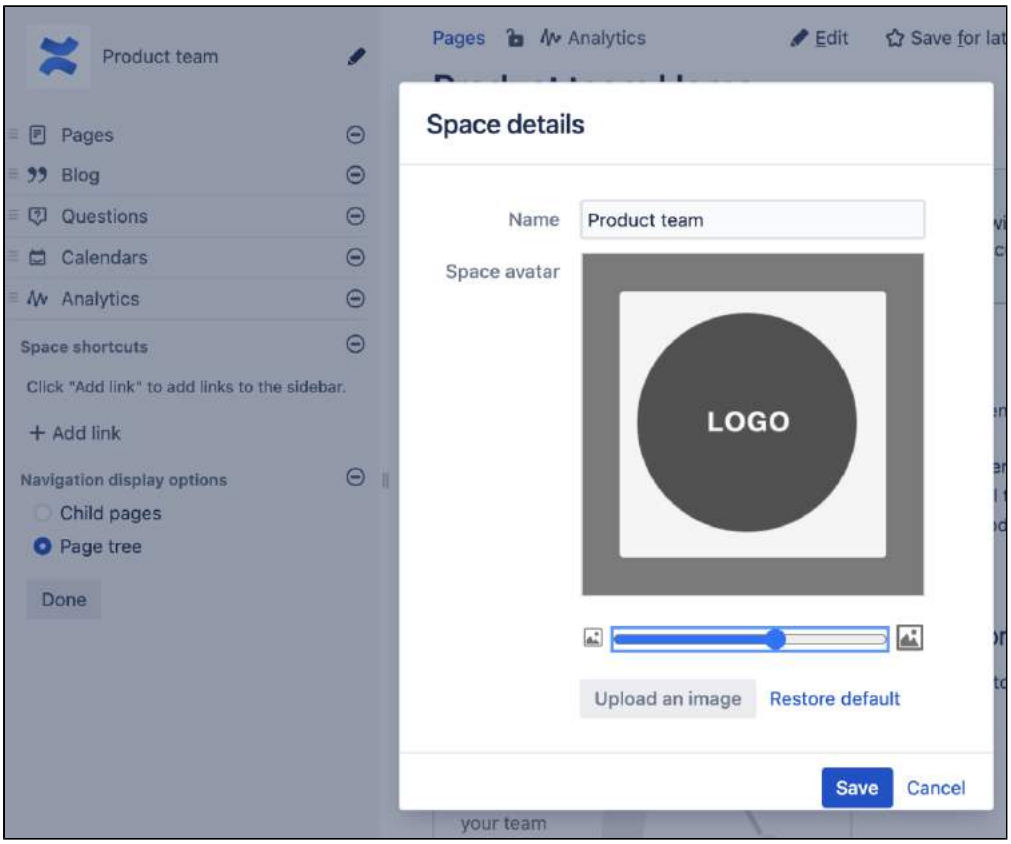

Notes

- **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 square. Personal space logos 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.

Copy a Space

Weve introduced the ability to copy spaces without installing any additional apps. If you were using our MarketplaceCopy Space app before, the built-in ability to copy spaces will automatically replace the app. No action is required from you.

To copy a space, go to**Space tools**>**Overview**>**Copy Space**.

The Marketplace version of the Copy Space app will be supported until28 July 2024.

If you're using Confluence 8.2 or any older version and you haven't installed the Copy Space app, you can copy the content of an entire space by creating a new space and copying the pages from the existing space to it.

To learn more about this workaround, go to the documentCopy a Page > Copy the contents of an entire space.

On this page:

- What will be copied from a space?
  - Known limitations
- Copying a space
- Performance tests
- Troubleshooting

Related pages:

- Create a Space
- Copy a Page

You can copy the content and settings of public and personal spaces across your Confluence instance. The ability to copy spaces provides you with full control, growth, and maintenance of the content that youre storing in Confluence. Plus, scaling up your instance will become easier and smoother.

You must be a **Confluence administrator** or **space administrator** to copy spaces. Also, you must have the **Create space** permission. Learn more about the Confluence permissions in Permissions and restrictions.

If you want to copy a personal space, you must either be its owner or have administrative permissions for it.

What will be copied from a space?

An existing space is always copied into a new space that Confluence creates automatically. The new space should have a name, a unique space key, and optionally, a description.

Not all the content and settings of your space will be duplicated in the new space. From the following table, youll find out what will be copied by default, what wont be copied, and whats optional to copy.

Content or setting	Will it be copied by default?	Additional notes
Space description	Yes	You can change the description when configuring the copy.
Space permissions	Yes	N/A
Space templates	Yes	N/A
Space look and feel	Yes	All the look and feel configurations will be copied by default, including the space header, sidebar, footer, and PDF stylesheets.

Archived spaces	Yes	N/A
Restrictions	Yes	N/A
Sidebar links	Yes	N/A
Shortcut links	Yes	N/A
Links in pages, blog posts, and comments	Yes	The following links in the original space will be copied: • To other pages in the space that youre copying • To other pages in other spaces • To blog posts in the space that youre copying • To blog posts in other spaces
Mentions	Yes	All mentions with the @ character will be copied but the mentioned users wont be notified about that in an email. They also wont see any notifications under the notification (bell) icon.
Pages	Optionally	All current versions from the original hierarchy are copied. If the original space has a home page, itll be copied too. ⓘ Copying of pages is audited in your Confluence instance. If any other option is selected for copying (for example, comments or attachments), the additional parameters will show this to you in the audit.
Blog posts	Optionally	All current versions from the original hierarchy are copied. ⓘ Copying of blog posts is audited in your Confluence instance. If any other option is selected for copying (for example, comments or attachments), the additional parameters will show this to you in the audit.
Existing authors and dates	Optionally	

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Heres how the authors and dates of the content from the original space will be set in the new space depending on whether you select this option or not:

If you copy pages and use existing authors and dates, the original page authors and dates of page creation and last update will be preserved in the new space.

If you copy pages but dont use existing authors and dates:

- Youll be assigned as an author to all pages in the new space.
- The current date when youre copying the space will be set as a creation date for all pages in the new space.

If you copy blog posts, their original authors and dates of creation and last update will always be preserved in the new space, regardless of whether you use existing authors and dates or not.

If you copy comments along with pages or blog posts, the original authors and creation dates of all comments will always be preserved in the new space, regardless of whether you use existing authors and dates or not.

Only page attachments are affected if you choose to copy existing authors and dates:

- If you select the option, the attachments copied to the new space will keep the original users who uploaded them and the original dates of the uploads.
- If you dont select the option, youll be set as a user who attached files to pages. The date when youre copying the space will be set as a creation date for the attachments.

For blog post and comment attachments, the new space will always display the original users who attached files and the original dates when files were attached, regardless of whether you use existing authors and dates or not.

> If you copy attachments but dont copy comments, comment attachments will still be copied to the new space. Youll find them in **Att achments** of a page or blog post.

If you use existing authors and dates:

- The original creator of the space youre copying will be set as the creator of the new space.
- The original date when the space youre copying was created will be set at the creation date of the new space.

If you dont use existing authors and dates:

- Youll be set as the creator of the new space.
- The date when youre copying the space will be set as the creation date of the new space.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Comments (page comments, blog post comments, inline comments, resolved comments, pinned comments on attachments)	Optionally	If you copy comments along with pages or blog posts, all comments will be copied to the new space. The original authors and creation dates of all comments are always preserved in the new space, regardless of whether the **Existing authors and dates** option has been selected. For both pages and blog posts, if you copy comments but dont copy attachments, comment at tachments wont be copied. Youll see the **Unknow n Attachment** labels instead of files. If you copy comments and attachments, pinned comments on attachments will also be copied.
Attachments	Optionally	If you copy attachments along with pages or blog posts, all files attached to them will be copied to the new space. For page attachments, the copy depends on whether you use existing authors and dates along with attachments or not: • If you use existing authors and dates, the original users who uploaded attachments and the original dates when the files were uploaded will be preserved in the new space. • If not, in the new space, youll be set as a user who uploaded attachments. The date when youre copying the space will be set as the creation date for these attachments. For blog post and comment attachments, the new space will always display the original users who attached files and the original dates when files were attached, regardless of whether you use existing authors and dates or not. If you copy attachments but dont copy comments, comment attachments will still be copied to the new space. Youll find them in **Att achments** of a page or blog post. If you copy attachments with comments, pinned comments on attachments will also be copied.
Space labels	Optionally	N/A
Labels on pages and blog posts	Optionally	N/A
Attachment labels	Optionally	N/A
Watchers	Optionally	Watchers include users who are watching: • Existing pages and blog posts in the space youre copying. • New blog posts in the space youre copying. • The space youre copying.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Space status `favorite`	No	A space acquires the `favorite` status when you star it and it's added to the list of your selected spaces.
Drafts	No	N/A
Page and blog post history	No	N/A
Page and blog post likes	No	N/A
Comment likes	No	N/A
Comments on attachments	No	N/A
Page and blog post watchers	No	N/A
Team Calendars	No	N/A
Confluence Questions	No	N/A

Known limitations

The following technical limitations were put in place to prevent errors and failures that may occur otherwise:

- You can copy only one space at a time. You should repeat the procedure for each space that you want to copy.
- You can copy spaces within only one Confluence instance. Copying spaces between different instances isnt possible.
- A copy is created in one long-running operation. Check the progress bar to find out about the status of the operation.
- From one branch of pages in the original space tree, only 2000 pages can be copied at a time. If the branch contains more than 2000 pages, the rest of the pages wont be copied and the copying will stop.
    - This limitation doesnt apply to blog posts as well as to the case where no branch in the original space contains more than 2000 pages.
- Once space copying has started, you cant cancel it. You should wait until the operation is completed to see the result.
- Other users shouldnt make changes to a space that is being copied. Making changes to the content while a copy is in progress may cause the operation to fail.
- Only the latest version of a page and blog post is copied.
- Third-party apps, like Scaffolding Forms & Templates, aren'tcopied.
- The copy operation can be a few times slower if Confluence uses Amazon S3 storage to store attachments. We recommend using a filesystem (FS) storage if you want to run copies more quickly.

Copying a space

⚠ Before copying a space, make sure youve informed other users that they shouldnt make any changes to this space until you finish copying it. Any changes they make may cause the copying to fail.

We recommend that you notify other users in advance, copy the space when the activity in your instance is low, or temporarily restrict access to the space you want to copy.

The space youre copying will also display a banner notifying other users about the copy operation and asking them not to make any changes to the space until the copy is completed. As soon as the space is copied, the banner will disappear.

To copy the content and settings of an existing space:

1. Go to the space that you want to copy.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

2. Select **Space tools** > **Overview** > **Copy Space**.

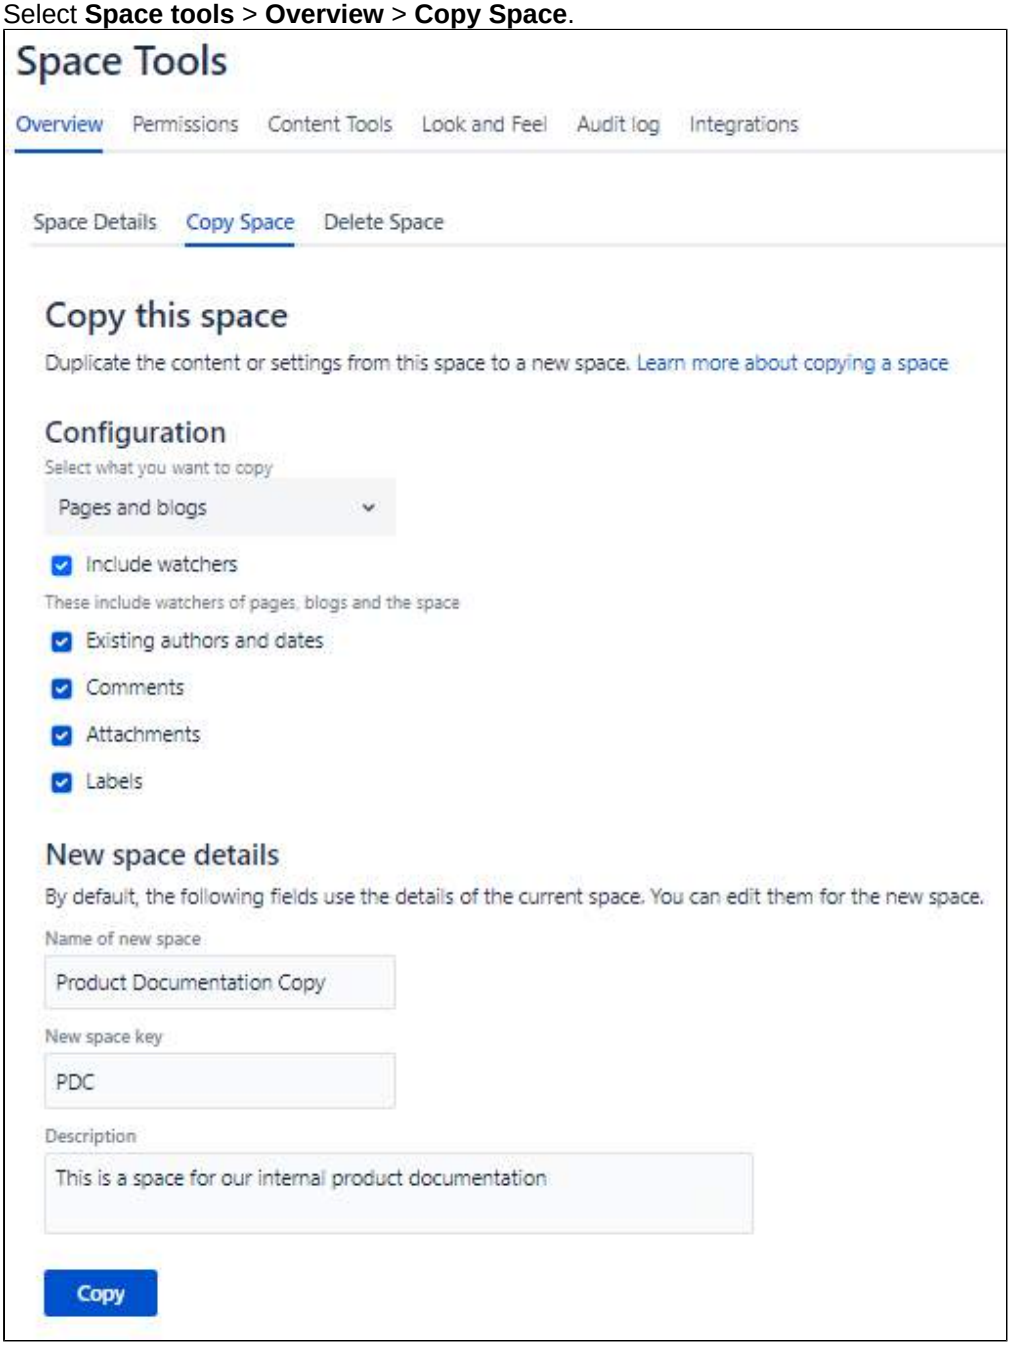

3. Select what you want to copy:
   a. **Pages and blogs** both pages and blog posts will be copied to the new space.

      > Page and blog copying is audited in your Confluence instance. If any other option is selected for copying (for example, comments or attachments), the additional parameters will show this to you in the audit.

   b. **Pages only** only pages will be copied to the new space.
   c. **Blogs only** only blog posts will be copied to the new space.
   d. **Space settings only** the new space will be created with all the settings from the original space, like permissions, templates, look and feel, etc.
      i. The new space wont have any pages, including a home page, or blog posts.
      ii. You can also select to copy the original creator of the space, the original date when it was created, and the original space labels.
      iii. If you dont copy the original space creator and creation date, youll be set as a creator of the new space, and the current date when youre copying the space will be set as the creation date of the new space.
4. If youre copying pages, blog posts, or both, choose what other content you want to copy.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Content option	What you should know
Include watchers	If you want to copy watchers, keep the checkbox selected. If not, clear the checkbox.
Existing authors and dates	ⓘ If selected, the option applies only to original authors and dates of pages and page attachments. Original authors and dates are always preserved in the new space for blog posts, comments, blog post attachments, and comment attachments. If you want to keep the original authors of pages and original dates of page creation and last update, keep the checkbox selected. If you dont want to keep the original authors and dates, clear the checkbox. In this case: • Youll be set as the creator of the new space. • In the new space, youll be assigned as the author of all pages and page attachments (if you choose to copy attachments). • The current date when youre copying the space will be set as the creation date for the new space, all pages in it, and all page attachments (if you choose to copy attachments).
Comments	If you want to copy comments, keep the checkbox selected. All comments will be copied for pages or blog posts, depending on what you want to copy. ⓘ The original authors and creation dates of all comments are always preserved on all pages and blog posts in the new space, regardless of whether the **Existing authors and dates** option has been selected. If you dont want to copy comments, clear the checkbox.

Attachments	If you want to copy attachments, keep the checkbox selected.
	• Only page attachments are affected if you select **Existing authors and dates**:
	  • If you select the option, the attachments copied to the new space will keep the original users who uploaded them and the original dates of the uploads.
	  • If you dont select the option, youll be set as a user who attached files to pages. The date when youre copying the space will be set as a creation date for the attachments.
	• For blog post and comment attachments, the new space will always display the original users who attached files and the original dates when files were attached , regardless of whether you select **Existing authors and dates** or not.
	
	If you dont want to copy attachments, clear the checkbox.
	
	ⓘ If you copy attachments but dont copy comments, comment attachments will still be copied to the new space. Youll find them in **Attachments** of a page or blog post.

Labels	If you want to copy labels, keep the checkbox selected. If not, clear the checkbox.

5. In the **Name of new space** field, edit the name of the new space.
6. In the **New space key** field, edit the key of the new space.
7. In the **Description** field, enter the description of the new space.
8. Select **Copy**.

To track the progress and status of the copy operation, check the progress bar.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Yes, you can. The copy operation will continue and the progress bar will display changes even if you leave or refresh the page.

While youre copying the space, other users can copy it too. At the same time, the space will display a banner notifying other users about the copy operation and asking them not to make any changes to the space until the copy is completed. As soon as the space is copied, the banner will disappear.

As the initiator of the copy, only you can monitor its progress.

When the copy operation is completed, the success message will display on the page. The message will contain a link to the new space.

If something goes wrong, check the reason and solution for the issue in Troubleshooting or reach out to Atlassian Support.

Performance tests

To find out how much time it takes for the copy operation to complete, check our performance test calculations.

Number of pages	Number of blog posts	Total quantity/size (Mb) of attachments	Duration (h:mm:ss)
100	100	190/944.3 Mb	0:02:05
3000	3000	6000/180 Mb	0:16:32

In both tests, copying attachments took about 50% of the total time that the entire space took to copy. This means that if your space has more attachments, copying it might take longer.

The average time for copying one page is up to 1 second. Were working on improving the performance in future versions.

⚠ The copy operation can be a few times slower if Confluence uses Amazon S3 storage to store attachments. We recommend using a filesystem (FS) storage if you want to run copies more quickly.

Troubleshooting

Is there anything you cant copy? Has the space copying failed? Look for answers and solutions here. If we havent listed your issue, reach out to Atlassian Support.

After the copy has been created, you can add anything to it without worrying about any issues. The content you add will be available only in the new space.

Yes, you can. In your new space, you can do the same thing youd do if you created a space from scratch.

In this case, creating the copy might fail. Try to contact a person whos making changes in the space and ask them to wait until you copy it successfully.

We recommend that you notify other users in advance, copy the space when the activity in your instance is low, or temporarily restrict access to the space you want to copy.

Despite the failure, part of the original space might be copied. We recommend that you delete the partial copy and retry copying the original space.

The space youre copying will also display a banner notifying other users about the copy operation and asking them not to make any changes to the space until the copy is completed. As soon as the space is copied, the banner will disappear.

To solve the issue, check the following causes:

- The space youve copied doesnt have a home page. In this case, the new space wont have a home page either.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

- Youve selected to copy only the settings of the original space. In this case, the original home page wont be copied and neither will other original pages and blog posts.
- Something went wrong during the copy operation. In this case, delete the new space and copy the original space again.

Someone might have revoked your **Create space** permission or administrative permissions. Check your permissions and contact your Confluence admin for assistance.

The following table lists the reasons why some links dont work in the new space after the copying. Youll also find the solutions for some issues that you can fix yourself.

#	Case	Solution
1	Something went wrong during the copy operation.	Delete the new space and copy the original space again.
2	In the new space, an undefined link on a blog post redirects you to a new empty page.	This is a known issue were working on. It occurs when you copy blog posts but not copy pages. In this case, youll be redirected to a new empty page in the original space from a link on a blog post in the new space because Confluence detects this link as undefined.

The following table lists the reasons and solutions for failures during the copy operation.

#	Reason	Solution
1	Someone has made changes to the space youre copying.	Despite the failure, part of the original space might be copied. Delete the partial copy and retry copying the original space.
2	Someone has deleted the space youre copying.	Search for the original space by its name. If you find the space, try copying it again. If not, a system administrator can restore the deleted space.
3	A branch of pages in the original page tree contains more than 2000 pages.	This is a known limitation in Confluence. From one branch of pages in the original space tree, only 2000 pages can be copied. If the branch contains more than 2000 pages, the rest of the pages wont be copied and the copy operation will stop. This limitation doesnt apply to blog posts as well as to the case where no branch in the original space contains more than 2000 pages.
4	Theres not enough memory or disk space to complete the copy. The space youre copying might contain a lot of attachments, but the available memory or disk space isnt enough to save their copies. As a result, the original space cant be copied.	Try clearing the memory or disk space and copy the original space again. If the issue persists, contact Confluence Support.

A failure might occur before the copy operation starts. In this case, the error message wont display any link to the partial copy of the original space.

Retry copying the space. If the issue persists, contact your Confluence admin.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Even if your internet connection has dropped, the copy operation will continue in the background. A notification about this will appear on the page instead of the progress bar.

As soon as the connection is restored, the progress bar will update to display the latest status of the copy operation.

You might have not selected the **Attachments** option when copying your space. You can delete the new space and repeat the copy operation with the option selected.

To enable logs, follow the guidelines from the section *Add logging for an additional class or package* of the document on logging configuration.

In the **Class/Package name** field, enter `com.atlassian.confluence.plugin.copyspace`.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Archive a Space

Archiving a space is useful when you have content that is no longer relevant, but you still want the option of accessing it at a later date. Archived spaces are less visible, but still available on your site. Archiving a space is easy to undo you can make a space current again at any time.

On this page:

- Archive a space
- The effect of archiving a space
  - Spaces
  - Pages
- Change a space from archived to current

Related pages:

- Delete a Space
- Export Content to Word, PDF, HTML and XML

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**.

The effect of archiving a space

Spaces

If a space is archived, that space:

- Won't appear in the **Recent spaces** section of the search panel.
- Won't appear in search results, unless you select **Search archived spaces**.
- Won't appear in advanced search results unless you select**Search archived spaces**.
- Won't appear on the **Spaces** dropdown menu.

- Won't appear in the general spaces lists in the**Space Directory**, but will instead appear under the**Archived Spaces**list. It will, however, still appear under any categories it was labeled with.
- Won't show up in activity streams when updated.
- Won't appear on your dashboard.

Pages

Pages within an archived space will appear in a few places.

If you**view**a page within an archived space, that page will appear in:

- The **Recently visited** section of the search panel.
- TheRecently viewed pagesmenu.

If you**edit**a page within an archived space, that page will appear in:

- Activity streams
- TheRecently updated macro.

Pages within an archived space won't appear in search results, unless you select**Search archived spaces**.

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 thespace permissions.
- RSS feeds, watches andnotifications remain active.

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

Change a space from archived to current

Through the space directory:

1. Go to **Spaces** > **Space directory** in the header.
2. Choose **Archived Spaces** on the left.
3. Find your space and click the on the right.
4. Choose **Edit Space Details**.
5. Change the **Status** from 'Archived' to 'Current' and hit **Save**.

Through the archived space:

1. If you know the **space key**, you can navigate straight to the archived space - `http://yoursite/display/SPACEKEY`
2. Choose **Space tools** > **Overview** from the bottom of the sidebar.
3. Choose **Edit Space Details**.
4. Change the **Status** from 'Archived' to 'Current' and hit **Save**.

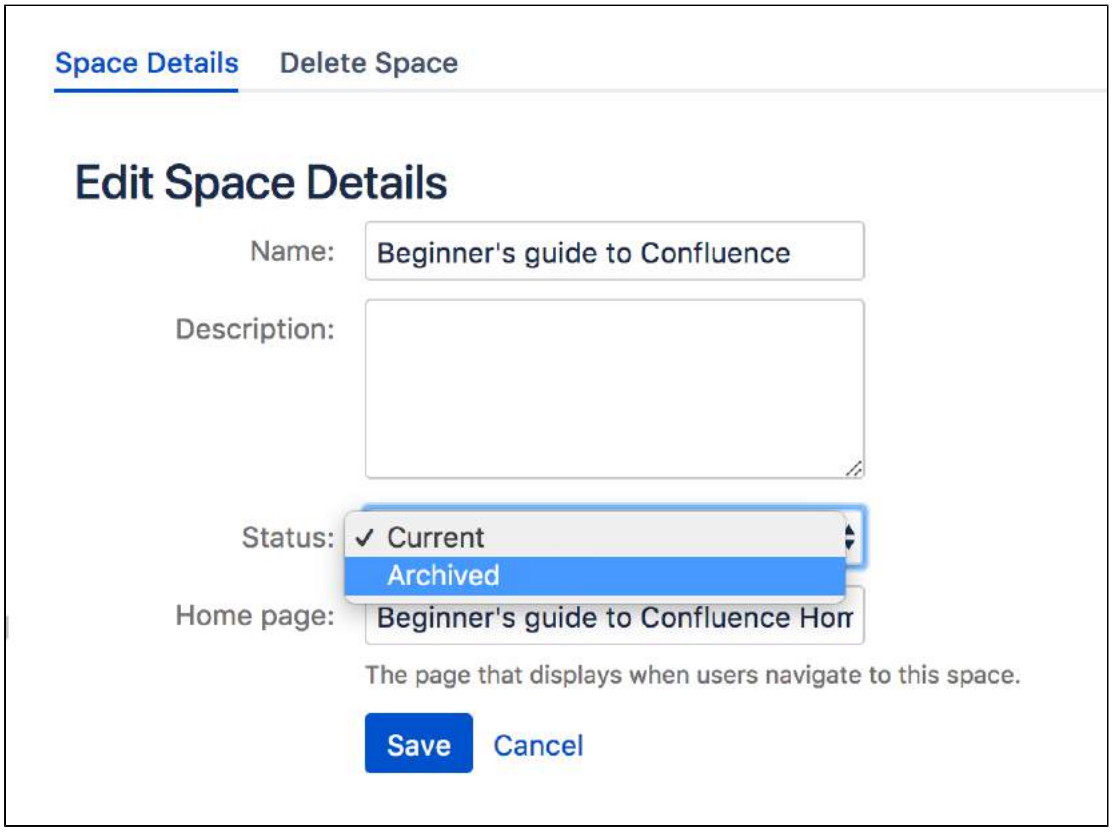

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Delete a Space

Deleting a space permanently removes the space and all of its contents, including any calendars and questions linked to that space. Only someone withspace admin permissionscan delete that space.

⚠Deleting a space is permanentit does not go to the trash and cannot be undone.

If you're unsure about deleting a space,create anXML export of the space as a backup before proceeding. You can then restore the space from the XML export file if you need to.

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**.

Related pages:

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

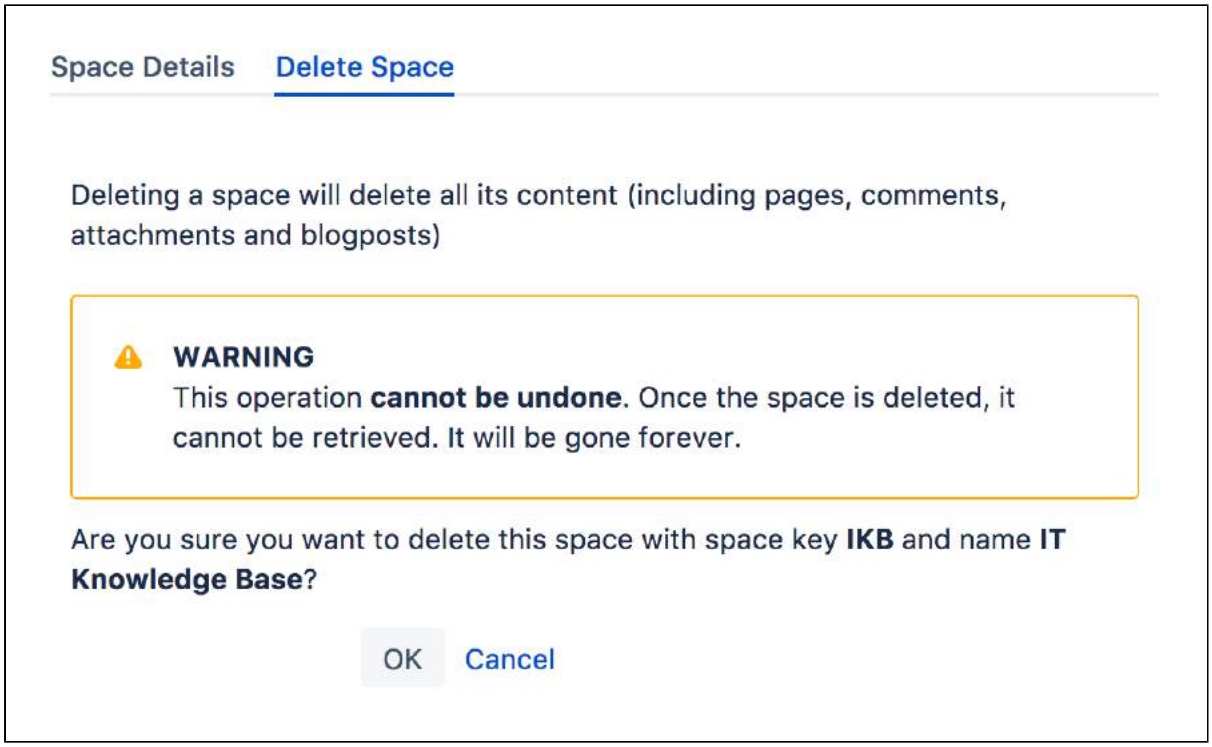

Members of the `confluence-administrators` group can also delete spaces, including personal spaces.

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 tospace permissions.

On this page:

- Export single pages to PDF
- Export single pages to Word
- Export multiple pages to HTML or PDF
- Export (back up) multiple pages to XML
- Customizing the appearance of PDF exports
- Migrating content to Confluence Cloud

Related pages:

- Customize Exports to PDF
- Restore a Space or multiple Spaces

Export single pages to PDF

If you need to send content to people who don't have access to Confluence, you can export a single page or blog post as a 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••• **(Tools)**>**Export to PDF**.

Only published content is exported. This means you can create PDF exports even while people are still working on the page.

When you export a single page to PDF, the PDF stylesheet customizations are applied, but any PDF layout customizations are not. To make your PDF layout customizations apply to a single page exported to PDF, you'll need to use the 'multiple page' method described below to export the single page. SeeCustomize Exports to PDF.

Export single pages to Word

You can also choose to export single pages to a .doc file that can be opened in Microsoft Word.

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**Tools**>**Export to Word**.

Only published content is exported. This means you can create Word exports even while people are still working on the page. Also, only the first 50 attached images will be included in the export. See the notes below for more information.

Note that due to the format of this file, it can only be opened in Microsoft Word and is not compatible with other applications such as Open Office, Libre Office or Google Docs.

Export multiple pages to HTML or PDF

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

PDF export is useful you're producing a printable user manual from your documentation space for example. The HTML export can be used to convert your site content to a static website.

To export pages to HTML or PDF:

1. Go to the space and choose**Space tools**>**Content Tools**from the bottom of the sidebar
2. Select **Export**
3. Select one of the following formats:
    a. **HTML** export an HTML file containing only the pages you have permission to view. With this format, you have two options:
        - **Normal Export** generate an HTML file for each page in the space, except for blogs, comments, and attachments.
        - **Custom Export** generate an HTML file of only the pages you've selected.

b. **PDF** export a PDF file containing only the pages you have permission to view. With this format, you have two options:
- **Normal export** generate a PDF file for each page in the space, except for blogs, comments, and non-image attachments.
- **Custom export** generate a PDF file of selected pages excluding any comments.

For normal export and custom export, you can also choose to **Include Pages Numbers** at the bottom of every page in your PDF.

4. Review your selection, and select **Export.**

You can download the zipped archive or PDF when the export process has finished.

Export (back up) multiple pages to XML

If you want to export a space as an XML, Confluence can create a zipped archive of the XML.

XML space exports can be used to import your space content into another Confluence space running the same or newer version of Confluence.

To export pages to XML:

1. Go to the space and choose**Space tools**>**Content Tools**from the bottom of the sidebar
2. Select **Back up** an XML export contains every page, blog posts, comment, and attachment in the space, but excludes blog posts.
3. Under**Save to restore directory**, you can:
   a. Give your XML export a **file name** (optional)
   b. Select**Save permanently** if you want to keep your file on the server, otherwise, all space backups are only temporarily saved to your `<home-directory>/restore/space`, where `<home-directory>` is `shared-home-directory` for Confluence Data Center, or `home-directory` for Confluence Server. You'll need access to the server to retrieve the file this way.
4. Select **Back up**

You can download the exported XML backup by selecting the file name when the export process has finished.

What's included in the export?

The following content is included in space exports.

	PDF export	HTML export	XML export (backup)
Pages	Yes	Yes	Yes
Blogs	No	No	Yes
Comments	No	Optional	Yes
Attachments	Images only	Yes	Yes
Unpublished changes	No	No	Yes
Page numbers	Optional	N/A	N/A
Restricted pages and blog posts	Only those you can view	Only those you can view	Yes, they remain restricted

Customizing the appearance of PDF exports

You can add a title page, table of contents, and customized headers and footers to the PDF output. For more advanced customizations, you can apply Cascading Style Sheet (CSS) modifications. These customizations are specific to each space, and you need the 'Space Administrator' permission to apply these customizations. For more information, seeCustomize Exports to PDF.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Notes on PDF exporting

- To export a PDF containing international text, seeCreate 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 unrecognized customization. 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.
- Tables that exceed the width of a page, particularly those with images in them, might be cut off in the PDF. SeeWide tables are cut off in PDF exportsfor some suggested workarounds.
- In Confluence Data Center, PDF exports are handled by the external process pool.

Notes on Word exporting

- Only the first 50 images attached to the page are exported to your Word document. This is to prevent out of memory errors affecting your whole Confluence site. See

  ~~CONFSERVER-34211~~ - If a page with big number of images Exported to Word, some images are invisible CLOSED

  for more information, and to find out how you can temporarily increase this limit using a system property.

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.

- To customize 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, seeHow to Edit Files in Confluence JAR Files

Notes on XML exporting (backup)

- When you create a space backup, it does't mean it will start immediately. If someone has created another backup task on your Confluence site, your job will be added to a "queue". This will be reflected in your task status. Only one backup is processed at a time; we do this to provide a stable, fast, and reliable backup experience.
- If your task was "queued" and someone restarts Confluence, your task won't lose its place in the queue.
- If your task was "in progress" and someone restarts Confluence, it will fail and you will need to run the backup again.
- The **Space backup recent activity** table lists XML exports of that space only over a two-week period.
- You can only import a space into the same or later Confluence version.You **can't import a space into an earlier**version.
- **Team Calendars** aren't included in XML exports.
- If you're doing the export for backup purposes, consider another means of backup. See Production Backup Strategy.
- See Restore a Space or multiple Spaces for notes on restrictions when importing a space into Confluence Server.

Migrating content to Confluence Cloud

If you're migrating from Confluence Server to Confluence Cloud, you can use theConfluence Cloud Migration Assistantto migrate your content and spaces.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Customize Exports to PDF

Confluence provides a basic PDF export that you can customize and style to suit your needs.

You need Space Admin permissions to customize PDF exports in a space, and Confluence Administrator global permissions to customize PDF exports for the whole site.

On this page:

- How it works
- Limitations
- Change the PDF layout
- PDF Layout examples
  - Title page
  - Header or footer
- Change the PDF stylesheet
- PDF Stylesheet examples
  - Page size
  - Page orientation
  - Page margins
  - Page breaks
  - Word wrapping
  - Title page
  - Table of contents
  - Tables
  - Page numbers
  - Headings
  - Other formatting

Related pages:

- Advanced PDF Export Customizations

How it works

When someone exports a space to PDF, Confluence will apply any layout and styling directions it finds in the current space or set globally for the whole site.

- The **PDF Layout** allows you to add a title page to your PDF, and add a header and footer to all pages.
- The **PDF Stylesheet** allows you to change the look of the PDF. You can change just about anything, including the paper size, fonts and colours, spacing, and control behaviours like page breaks.

Both the PDF Layout and PDF Stylesheet can be customized on a space by space basis, or globally for the whole site. Space customizations will always completley override any global customizations. This means you can't mix and match and set some items globally and others at the space level.

Limitations

There are a few limitations to be aware of:

- Changes to the PDF layout only apply to space exports, not to single page exports (via ••• >**Export to PDF**).
- Confluence Server and Data Center process space exports slightly differently. This means that some options, like adding page numbers via CSS, aren't available in PDFs created with Data Center. We recommend selecting **Include page numbers** on the export screen if you need to number your pages.
- We provide a number of example customizations to get you started, however Atlassian Support can't help you with styling your PDFs or problems introduced by your customizations. If you're newto CSS, you might want to get help from anAtlassian Solution Partner, or check out a Marketplace app likeScroll PDF Exporterwhich has a WYSIWYG editor to help you produce beautifully styled PDFs.

Change the PDF layout

The PDF Layout is where you add a title page, header, or footer to your PDF exports.The PDF layout fields accept HTML. You can include inline CSS in the HTML too.

To change the PDF layout for the whole site:

- Go to > **General Configuration**>**PDF Layout**.
  You need Confluence Administrator global permissions to do this.
  You need
- Choose **Edit**, then add your customizations in the **Title**, **Header**, or **Footer** fields.

To change the PDF layout for a specific space:

- Go to the space and choose**Space tools**>**Look and Feel**from the bottom of the sidebarYou'll needSpace Admin permissionsto do this.
- Choose **PDF Layout**.
- Choose **Edit**, thenadd your customizations in the Title, Header or Footer fields

PDF Layout examples

Here are some examples of things you can do in the PDF Layout. The PDF layout can accept HTML and inline CSS.

Title page

In this example we've added the title "Documentation for Confluence", a logo, and an additional title "Contents" above the table of contents.

```
<div class="fsTitlePage">
 <img src="/download/attachments/169118009/atlassian_logo.png" />
<div class="fsDocTitle">Documentation for Confluence</div>
</div>
<div class="tocTitle">Contents</div>
```

The logo image we've used is attached to a Confluence page in the same site. You can find out the attachment ID by right clicking the image on the page, and copying its location.

Header or footer

In this example we've added plain text to the footer with some copyright information, and included a link.

```
Created in 2018 by Atlassian. Licensed under a <a href="http://creativecommons.org/licenses/by/2.5/au/"
>Creative Commons Attribution 2.5 Australia License</a>.
```

Change the PDF stylesheet

The PDF Stylesheet allows you to change the appearance of your PDF. This includes things likepaper size, fonts, colours, spacing, and other styling.

To change the PDF stylesheet for the whole site:

- Go to > **General Configuration**>**PDF Stylesheet**.
  You need Confluence Administratorglobal permissionsto do this.
  You need
- Choose**Edit**, then add your CSS.

To change the PDF stylesheet for a specific space:

- Go to the space and choose**Space tools**>**Look and Feel**from the bottom of the sidebarYou'll needSpace Admin permissionsto do this.
- Choose**PDF Stylesheet**.
- Choose**Edit**, thenadd your CSS.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

PDF Stylesheet examples

Here are some examples of common CSS overrides that you can use in your PDF Stylesheet.

The default CSS rules will apply unless you have specified an override in the PDF Stylesheet.

Page size

The default page size is US Letter(8.5 inches wide by 11 inches long). To override this behaviour and specify a particular size, add `asize`property to the`CSS @page`rule.

For example to export your spacein A4 size:

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

Page orientation

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

For example, to generate an A4 PDF in landscape, your`@page`rule might look like this:

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

Page margins

The default margins are 0.5". To set all margins to 15 mm, with a paper size of A4, edit the`CSS @page`rule in the PDF Stylesheet, like this:

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

To set the margins independently, edit the`@page`rule like this:

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

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

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

```
@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 */
}
```

In the example above we've also used the `:first` pseudo-class to define different margins for the title page.

Page breaks

By default, each Confluence page will start on a new page in the PDF. If you don't want each Confluence page to start on a new page, you can override the default page breaks using the following CSS:

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

This behaviour changed in Confluence 6.13. If you're using Confluence 6.12 or earlier, page breaks are **not** added before each page title.

If you're using Confluence Data Center, you won't be able to change this behavior, as PDFs are generated page by page in the external process pool, and then combined together once all pages are complete.

Word wrapping

Long, unbreakable words or strings (such as a URL) will automatically wrap to fit the page width, or cell width if in a table.

If you don't want words or long strings to break you can use the following CSS:

```
div {
    word-wrap: normal !important;
}
```

This may mean that the table formatting in your PDF is problematic, and very long content may overflow, and be cut off the page.

Title page

If you have added a title page in the PDF layout, you can use the following rules to change the appearance of the title page and title text.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

```
.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;
}
```

Table of contents

A table of contents is included by default when you export a space to PDF. It will appear at the start of the document, or after the title page, if you've configured a title page in thePDF layout.

To omit the table of contents, use the following override:

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

The table of contents uses a leader character to visually connect the page title with it's page number. By default this is a dot. Allowed values are`dotted`,`solid`and`space`. You can also use a string, for example`leader(". . . ")`.

The example belowuses solid line, instead of dots.

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

Tables

When you export a page that contains a table, we'll reduce the width of the table columns as much as possible, so that the whole table fits comfortably on the page. Individual columns are resized to fit the contents of each column.

If you prefer table columns to always be of equal width, you can use the following CSS:

```
table.fixedTableLayout {
    table-layout: fixed !important;
    width: 98% !important;
}
```

Any images in a table will be exported using the size set in the editor. If your table contains large images, part of the table may be cut off when exported to PDF. To ensure that nothing is cut off, we recommend resizing images in the editor, so that the total width does not exceed about 600px (for an A4 page in portrait orientation).

Alternatively you can use the following CSS to fit images to the available width:

```
table img.confluence-embedded-image {
    -fs-fit-images-to-width: 100%  !important;
}
```


Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Page numbers

The best way to add page numbers to your document is to select **Include page numbers** on the export screen. If you're using Confluence Server, and want more control over where page numbers appear, you can use CSS to add numbers instead.

> ⚠ If you're using Confluence Data Center, you can't add page numbers using these methods, as PDFs are generated page by page in the external process pool, and then combined together once all pages are complete. Use the **Include page numbers** option on the export screen instead.

To add page numbers in the format "Page x of y" to the bottom of your page, add the following CSS to the PDF stylesheet:

```
@page
{
    @bottom-center
    {
        content: "Page " counter(page) " of " counter(pages);
        font-family: ConfluenceInstalledFont, Helvetica, Arial, sans-serif;
        font-size: 8pt;
    }
}
```

Alternatively you can add page numbers into the footer. This requires making a change in the PDF layout and the stylesheet.

First, add an elementin the PDF layout. In this example we'll call it`pageNum`:

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

Then, in the PDF stylesheet, style the `pageNum`element as follows:

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

Headings

Heading sizes in the PDF export roughly match the sizes used on Confluence pages. You can easily override them as follows:

```
h1 {
    /* Custom styling */
}


h2 {
    /* Custom styling */
}
```

This behaviour changed in Confluence 6.13. In Confluence 6.12 and earlier, headings were demoted based on the position of the page in the page tree. Now headings are a consistent size on every page.

Other formatting

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

You can use the stylesheet to customize the output of just about anything on the page,including fonts, tables, line spacing, macros, etc. The export engine works directly from the HTML output produced by Confluence. Therefore, the first step in customizing 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.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Advanced PDF Export Customizations

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

⚠️ The information below is for advanced users. Be aware that the advanced customizations described below require knowledge of certain parts of Confluence, and of CSS and HTML. Customizations are not supported by Atlassian, so oursupport 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:

- Customize 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 ruleswithin your`@page`rule:`@top-left`,`@top-center`, and`@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 `@page` 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` attribute 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; }
```


Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

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 thetitle 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 recognize the NoPrint macro:

CSS - PDF Stylesheet

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

Indexing

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

Notes

If styling is not working as expected, it is useful to look at theintermediary HTML source to which the CSS is applied. This intermediary HTML is created whenever you create anHTML exportthat 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
```


Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

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://corefonts.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**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.
- For symbols, if the other fonts do not work, try Seguisym
- If the font file size is bigger thanyourcurrent attachment size limit, you will not be able to upload it. Please increase theattachment size limittemporarilyand re-upload again. An improvement of the error messaging is tracked at ~~CONFSERVER-24706~~ CLOSED
- To make use of an installed font in yourPDF Export style sheet (CSS)refer to it by the font-family ConfluenceInstalledFont.

Pages and blogs

Pages and blog posts allow you to capture and share information 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 can create it as a Confluence page or blog post.

Pagesare great for when you want the information to last and evolve over time. If it's a point-in-time update or one-time communication then a blog postis the way to go. These aren't hard-and-fast rules; they're just pointers to give you a place to start.

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 onSpaces.

Related pages:

- Create and Edit Pages
- Blog Posts
- The Editor
- Page Templates
- Delete or Restore a Page
- 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, mentions and emojis
- Page Layouts, Columns and Sections
- Create Beautiful and Dynamic Pages
- Page Templates
- Blueprints
- Import Content Into Confluence
- Undefined Page Links
- View Page Information
- Page History and Page Comparison Views
- Confluence Markup

Create and Edit Pages

Create a page

On this page:

- Create a page
- Edit together
- Collaborate or restrict
- Organize and move
- Other page actions
  - Notes

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 add rich text,tasks,images,macrosandlinks, or use one of the usefulblueprintsto capturemeeting notes,decisions, and more.

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.

1. **Create blank page**
2. **Create from template**

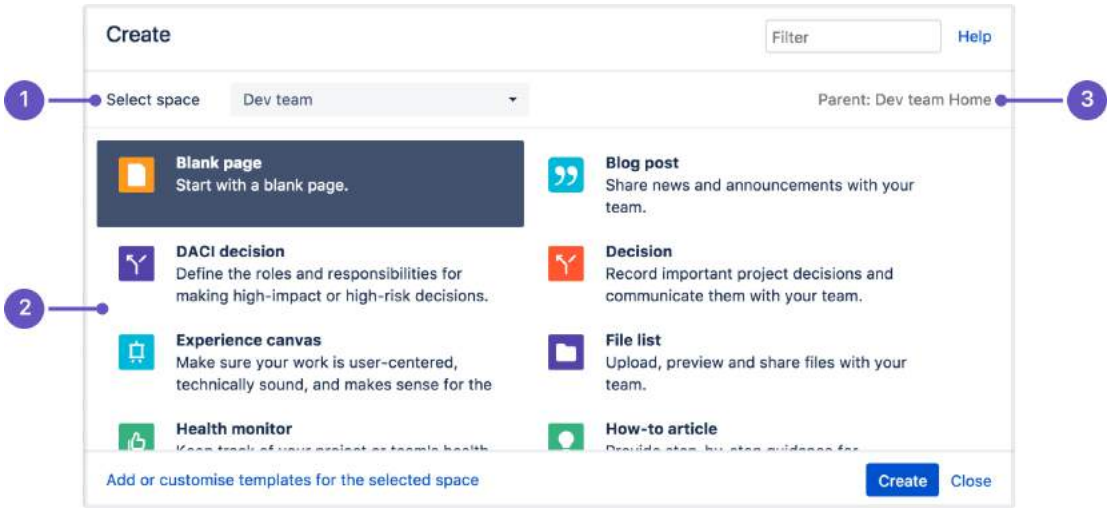

1. **Select space**: choose the space where you'll create the content.
2. **Page templates**: create a page from a template or create other types of content.
3. **Parent page**: your page will be a child of this page.

Once youdecide 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**Publish** when you're ready to make it appear in the space.

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.

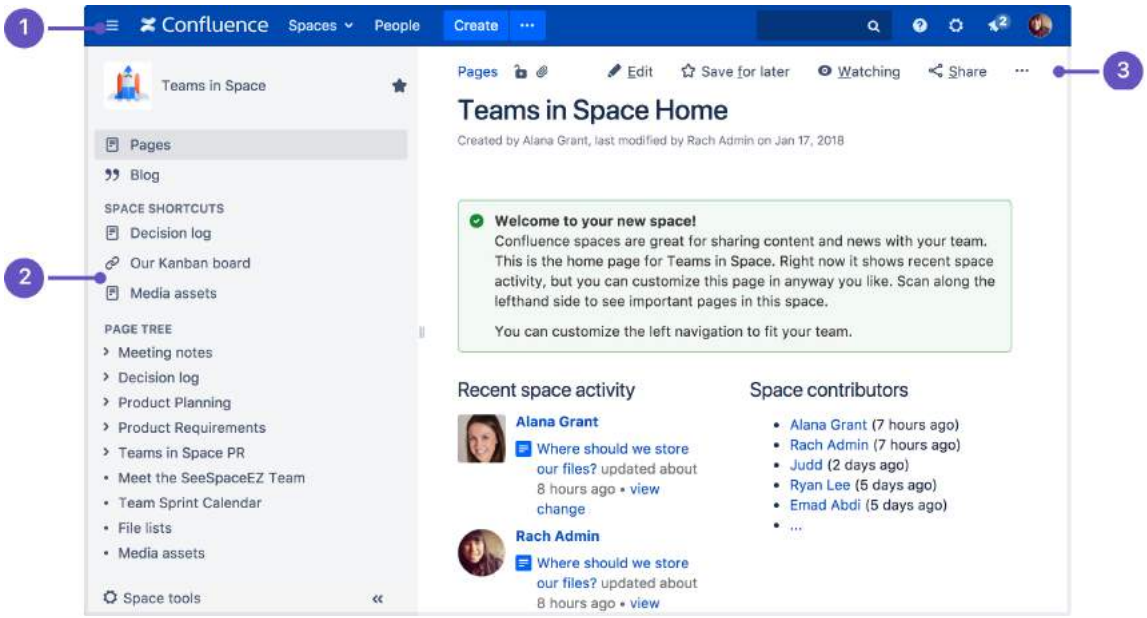

1. **Confluence header**: create blank pages, pages from templates and visit spaces or your profile.
2. **Space sidebar**: access pages, blogs and administer the space.
3. **Page tools**: edit or share the page, watch it to get updates and perform more actions.

Another useful way to create a page is to use theCreate 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.

Edit together

Need input from your team members? Multiple people can edit your page at the same time.

Hit the **Invite** button in the editor and either grab the link, or enter some people or groups to invite by email (t hey need the appropriate Confluence and space permissions of course).

SeeCollaborative editing for all the ins and outs.

Collaborate or restrict

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

Organize and move

You can also organize 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, andmove pagesif you want to change their location.

If you want to view all pages in a Confluence space, choose**Pages** in the sidebar.

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 organization, for projects, a combination of both, or for any reason you want to group pages together. SeeSpacesfor more information.

Other page actions

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

- Copy a page and its child pages
- 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 toWord, 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 Linksfor more information.

Notes

You may experience problems saving extremely large pages. Confluence can accept approximately 5mb of content (not including attached files) which is roughly equivalent to 800,000 words. If you do experience errors that indicate the page is too large to save, you should break up the page into several smaller pages.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

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 likingyour 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.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.

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

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.

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. SeeSpace Permissions.

Move a blog post

If you create a blog post in the wrong space, or want to reorganize 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 aspage 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 organization.

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

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

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.

On this page:

- Edit a page or blog post
- The editor
- Editor toolbar
- Emoji menu
- Insert menu
- Restrictions, labels, and notifications
- Things to help you work faster
- Find and replace text
- Invite people to edit with you
- Record change comments and notify watchers

Related pages:

- Tables
- Page Layouts, Columns and Sections
- Display Files and Images
- Links
- Symbols, emojis and special characters

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 changes as you type. Changes are only visible when viewing the page after you publish or update. SeeCollaborative editing for more information on how this works.

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 aresome things you should take into account.

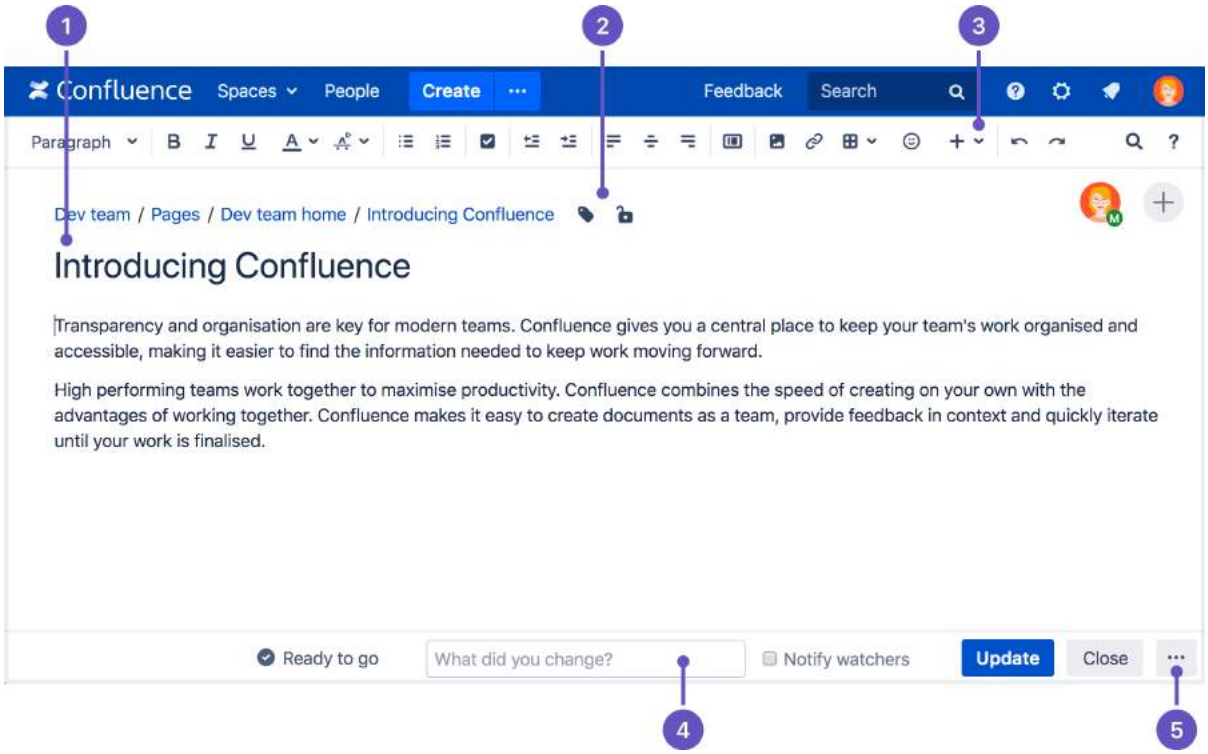

1. **Page content**: add your page title and body text.
2. **Page tools**: add labels, or restrict the page.
3. **Insert**: add files, links, images, tables and macros
4. **Notify**: notify others and leave a comment when you change a page.
5. **Preview or revert**: preview your page, view changes since last published, or revert back to last published version (or delete the draft page, if it has never been published).

Editor toolbar

The editor toolbar is where you format your page layoutand 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 thehelp icon ?.

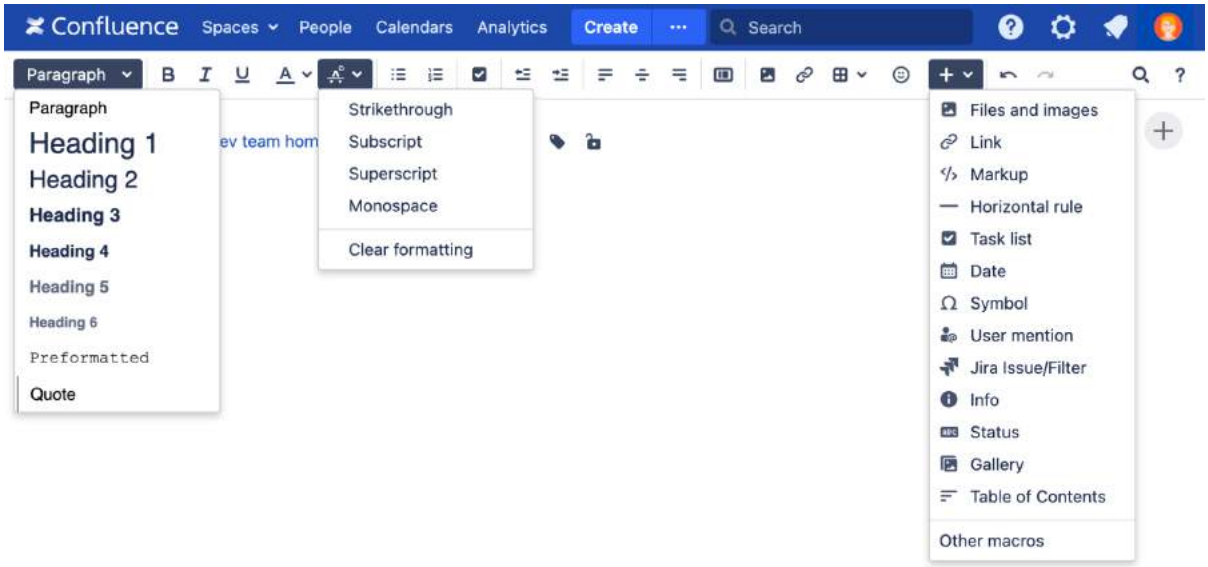

Emoji menu

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Use the emoji menu to preview, search for, and insert anemoji or symbol into your page. You can see the emoji's keyboard shortcut in the preview bar at the bottom of the emoji menu.

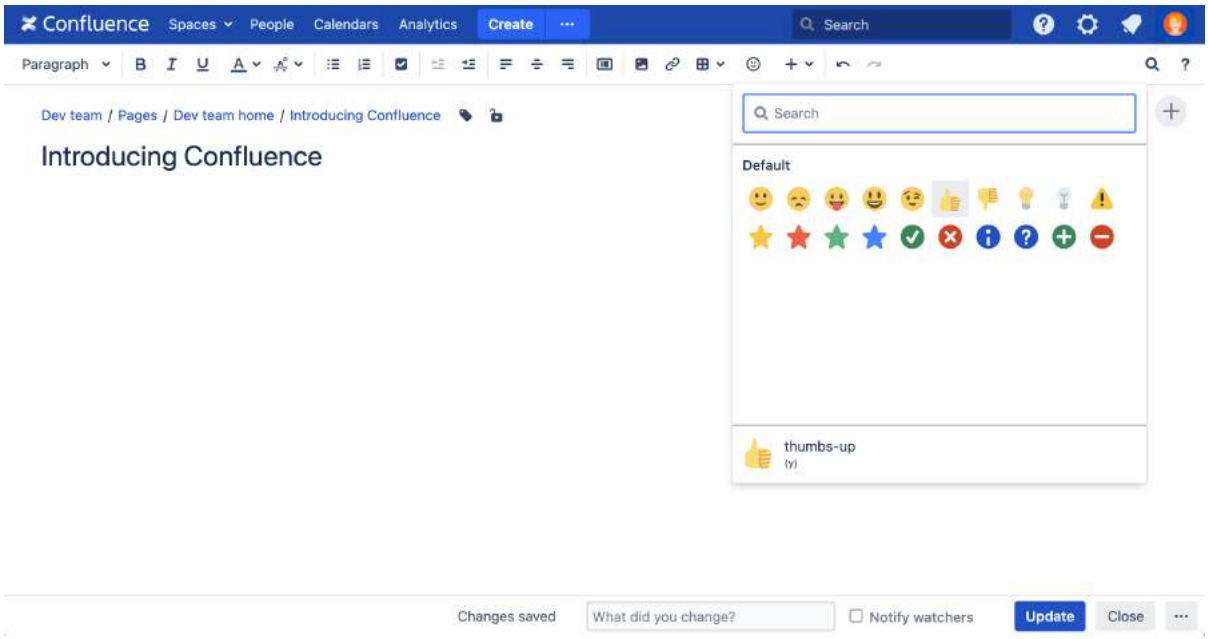

Insert menu

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
- 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**.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Autoconvert for pasted links

When you paste certain URLs into Confluence, the editor will analyze 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.

Keyboard shortcuts

To view the availablekeyboard shortcuts, choose the help icon ? in the editor toolbar.

Find and replace text

Click the iconon 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.

Invite people to edit with you

Speed up your draft and review cycles and get input from the right people by inviting them to edit the page with you. The page does not need to be published.

Hit the button in the editor and either grab the link, or add people, groups or email addresses to invite, just likesharing a page.

Sharing or inviting someone to edit a page or blog post does not automatically grant any permissions - they will still need the appropriate Confluence permissions to access Confluence and view or edit the page.

Up to 12 people can edit the same page at the same time (your administrator can change this limit).

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 thepage 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 for each page, so if you choose not to notify people, the checkbox will be deselected for you next time you edit that page.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

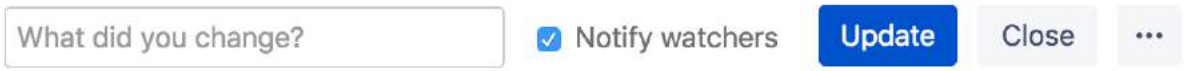

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Symbols, emojis and special characters

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

> **Using Confluence Cloud?**Check out our info on how toUse symbols, emojis and special charactersin the new editor.

On this page:

- Insert symbols and special characters
- Insert emojis
  - Select it from the emoji menu
  - Type the emoji shortcut
  - Trigger the emoji autocomplete
- Add your own emojis
  - View and delete custom emojis
- Prevent emojis from appearing

Related pages:

- The Editor

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

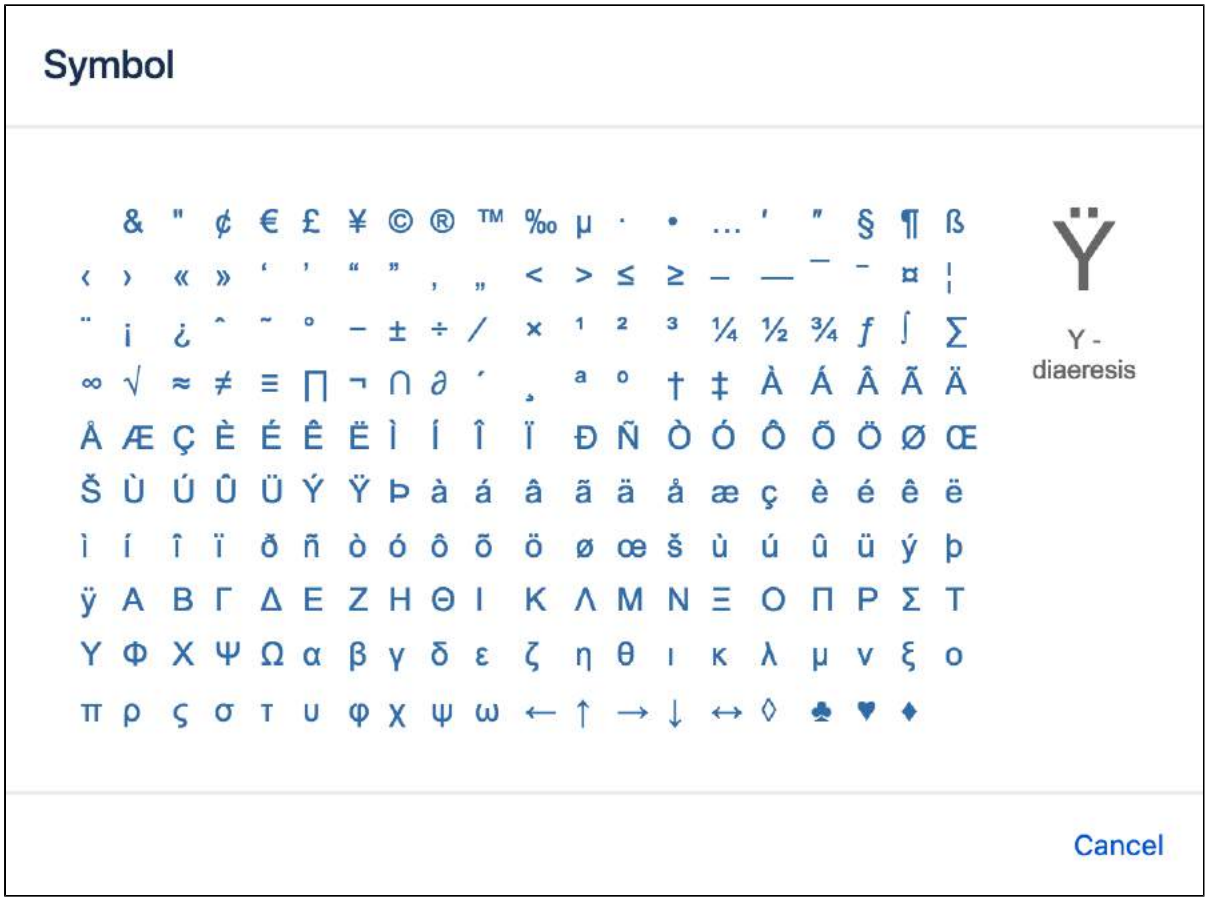

Insert emojis

There are a few ways you can add an emoji to your page.

Select it from the emoji menu

1. Edit the page
2. Select where you want to add the emoji
3. Open the emoji menu from the editor toolbar
4. Select an emoji to insert it.

Tip: Open the emoji menu using **Ctrl+Shift+M** (Windows) or **Cmd+Shift+M** (Mac).

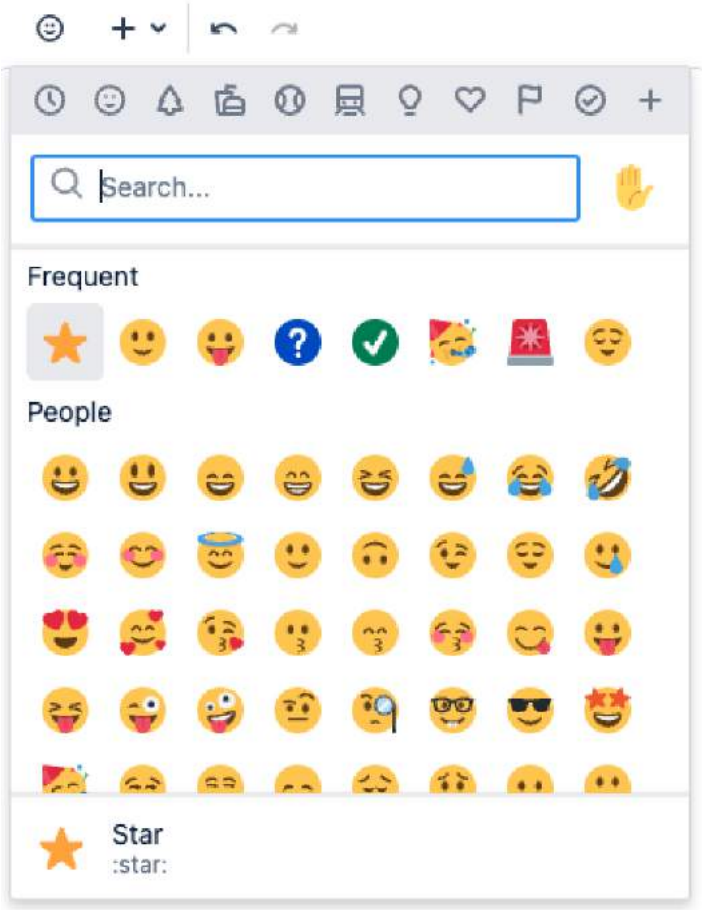

You can also personalize skin tone for certain emojis. Select the hand emoji (next to the emoji search bar) and it will expand to show the skin tone variations.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

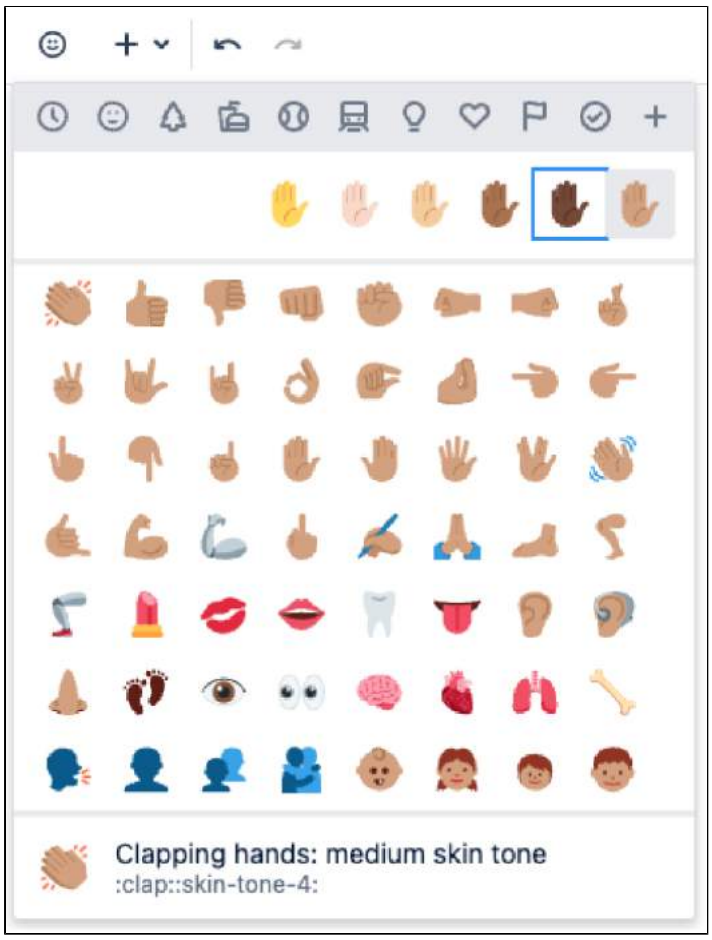

Type the emoji shortcut

1. Edit the page
2. Select where you want to insert the emoji
3. Type its keyboard shortcut

Emoji shortcuts start and end with a colon ":", for example :wave:. To learn an emoji's shortcut, hover over an emoji in the emoji menu. The emoji shortcut is visible in the preview panel.

Knowing an emoji shortcut is helpful when the emoji menu isn't available, for example in an inline comment.

You must have autoformatting enabled in the editor for this to work. See Edit Your User Settings

Screenshot: The keyboard shortcuts for our original emoticon set

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

:)	:(	:P	:D	;)
(y)	(n)	(on)	(off)	(!)
(*)	(*r)	(*g)	(*b)	(*y)
(/)	(x)	(i)	(+)	(-)
(?)	<3	</3		

Trigger the emoji autocomplete

1. Edit the page
2. Select where you want to add the emoji
3. Type a colon "**:**" to activate the emoji autocomplete menu
4. Type the emoji name, shortcut, or browse the suggestions using your up and down arrows or mouse
5. Select the relevant emoji to insert it

The emoji autocomplete menu will continue to change as you type to predict and match your input.

You must have autocomplete enabled in the editor for this to work. See Edit Your User Settings

Screenshot: The emoji autocomplete feature is activated when you type a colon in the editor

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

HipChat plugin emojis

If your team has migrated HipChat plugin emojis to Confluence, you will be able to see them in the emoji menu under **All uploads**.

Learn how to migrate Hipchat emoticons to Confluence Server and Data Center

Add your own emojis

You can upload your own emojis using the emoji menu in the editor.

Files must be JPG, PNG, or GIF, and no larger than 1 MB. Regardless of the original file's dimensions, the emoji created will be 36 x 36 pixels.

To upload a custom emoji:

1. Open the emoji menu from the editor toolbar
2. Select **+Add your own emoji**
3. Select the icon under Select an image'
4. Browse for and select the image file (we accept JPG, PNG, and GIF files up to 1MB)
5. Select **Open**
6. Give your emoji a name
7. Give your emoji a shortcut (this must be unique)
8. Select **Add emoji**

System administrators can view and delete custom emojis uploaded to their site at `https://[yourdomain]/admin/emojis/globalsettings.action`

Learn about managing emojis

View and delete custom emojis

Custom emojis will appear in the emoji menu under 'All uploads' (+).

To see the emojis youve uploaded:

1. Open the emoji menu from the editor toolbar
2. Select the **+** tab
3. Any emojis you've added will appear under Your uploads

To delete an emoji youve uploaded:

1. Open the emoji menu from the editor toolbar
2. Select the **+** tab
3. Under Your uploads, select the emoji you want to delete
4. Select the cross (x) to the top right the image or use the backspace keyboard shortcut
5. You will receive a confirmation message, select **Remove**

Prevent emojis from appearing

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

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

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Collaborative editing

Collaborative editing takes teamworkto the nextlevel by lettingyou and your teamwork together in real timeon software requirements, meeting notes, retros, and any other Confluence page. See who's editing the page with you, and see changes as they happen.Hit + to invite more people to edit with you.

Changes save and sync automatically,so everyone editing sees the same thing. And, because we're saving all the time, there's no need to manually save. Publish now or keep the draft and publish it lateryou're in control.

Up to 12 people can edit the same page at the same time. Your administrator can change this limit using asystem property.

On this page:

- Drafts and unpublished changes
- Things you should know

Once you and your team are done editing you can:

- **publish**(or update if the page has previously been published) to make everyone's changes visible
- close the editor and**keep**everyone's work to finish later
- **revert**to the published version of the page, discarding everyone's unpublished changes
- **delete**the draft page entirely, if it has never been published.

We'll warn you if you're about to publish (or discard) other peoples' changes along with your own.

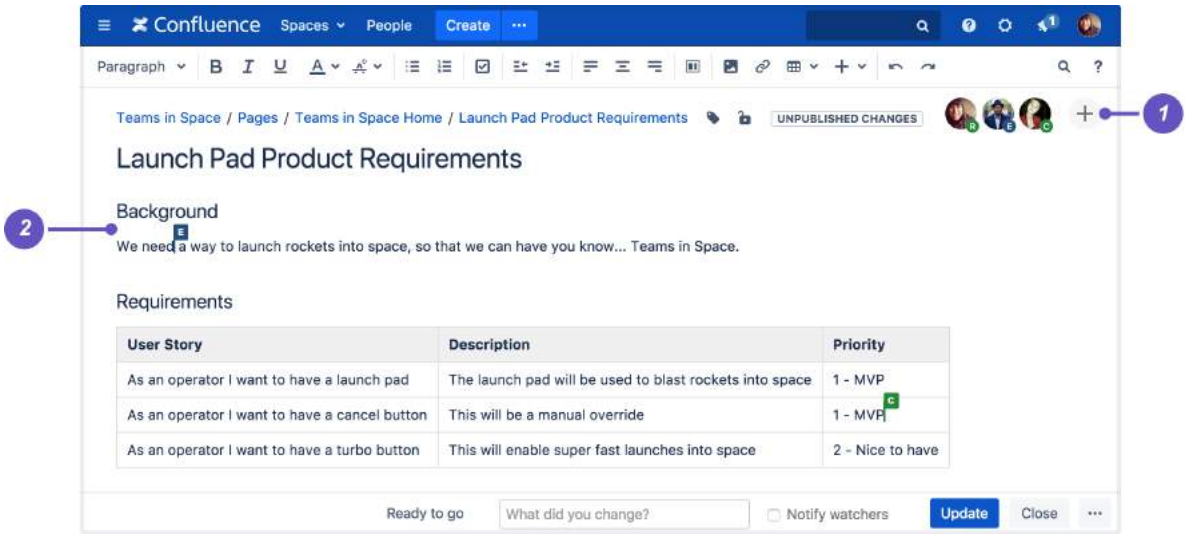

1. **Invite more people:**see who is editing the page and invite others to edit with you.
2. **See what they're doing:**watch others edit the page in real time.

Drafts and unpublished changes

What is a draft?

A draft is a page that has**never**been published before. Draft pages have a lozenge that says**draft**, and are only visible to their author, and to anyone that author shares their draft with. Nobody else will be able to see your draft, as it is only accessible from the**Recently worked on**list of each of the people who've contributed to it.

What are unpublished changes?

A page with**unpublished changes**is a page that has been published, and has then had edits made to it, but which has not yet been republished. Anyone who has unpublished edits will see the page in their**Recently worked on**list, with a lozenge saying**unpublished changes**. People who haven't contributed to the unpublished changes won't see this lozenge.

Those unpublished changes, however, are visible in the editor, and anyone can access them by editing that page. Therefore, if you have unpublished changes and do not want someone else making additional changes before they can be published, you might want to temporarily restrict editing on that page (leaving the published version of the page visible).

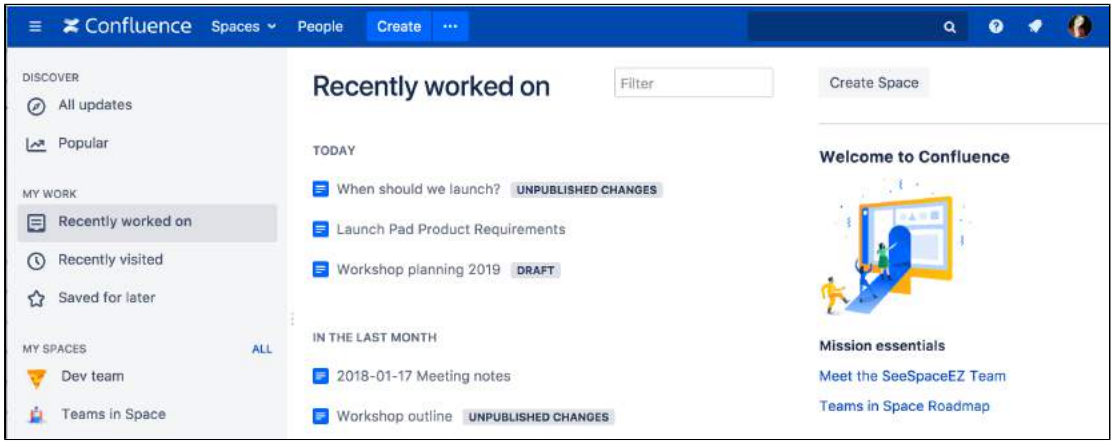

Things you should know

Limited content auditing

We dont yet have the same auditing capabilities with collaborative editing. All page changes are currently attributed to the person that publishes the page, rather than the person who made each specific change.

Changes in drafts aren't versioned

Were saving all the time in collaborative editing, but we dont save versions in a draft.When restoringan earlier pageversion,you can only roll backto published versions(the page draft is deleted when you restore a previous version)

No more personal drafts

Collaborative editing introduces a new type of draft, a*shared*draft. Previously, when you edited a page but didn't save it, Confluence would create a draft that was only visible to you (a personal draft). When collaborative editing is turned on, Confluence creates a shared draft whenever anyone edits a page.All page editors work on this same shared draft, and it exists untilsomeone publishes the page.

When you publish a shared draft, you're publishing all the changes you have made and changes made by others. Publishing creates a version in the page history.

If you discard a shared draft, you're discarding all changes, including changes made by others. Because shared drafts aren't versioned, there's no way to get a discarded draft back.

Any existing personal drafts are still available, but are no longer editable. If you edit a page, you'll see the shared draft of the page, not your personal draft (if one exists).

If you need to get content out of your previous personal drafts head to**Profile**>**Drafts**, locate your page and copy the contents.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

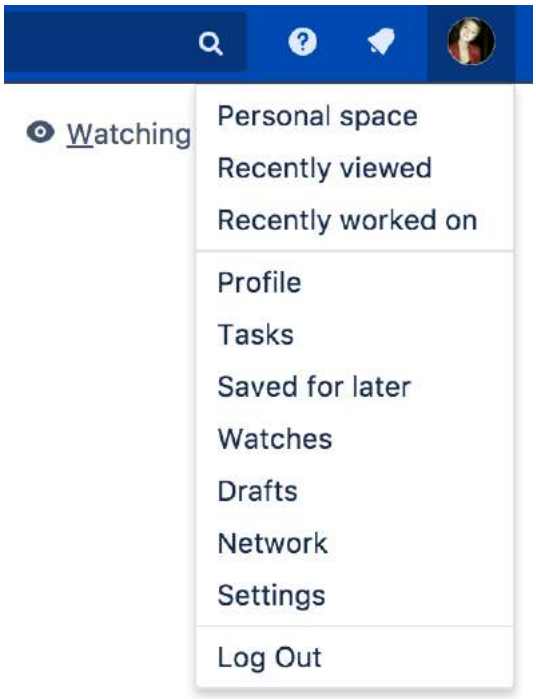

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

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).

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

Set page location or move a page

1. Do either of the following:
    - **While creating a page** choose the locationicon 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**)

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

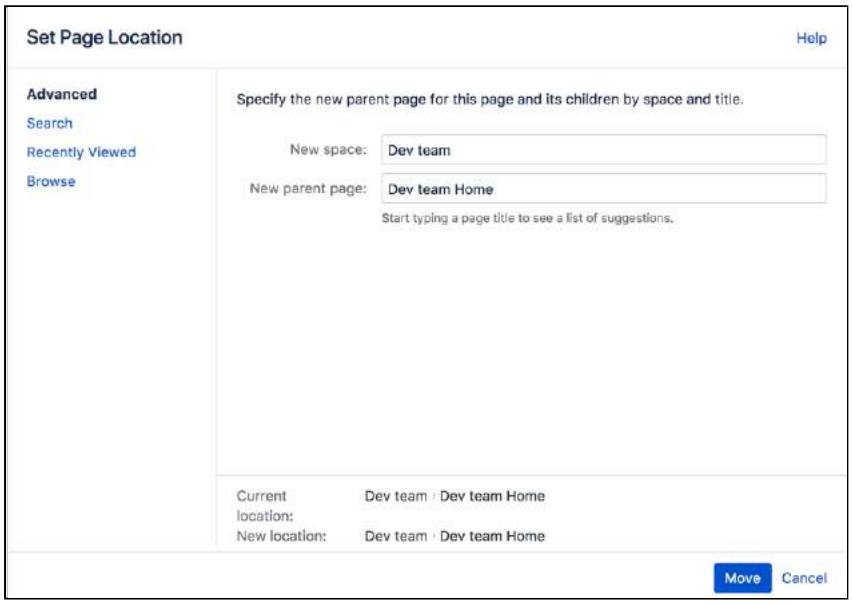

Reorder pages within a space

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.

Alternatively, you can choose to order a group of child pages alphabetically by choosing the**Sort 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 availableimmediatelyafter sorting the page, while you're still on the **Reorder Pages**t ab, and haven't performed any other action.

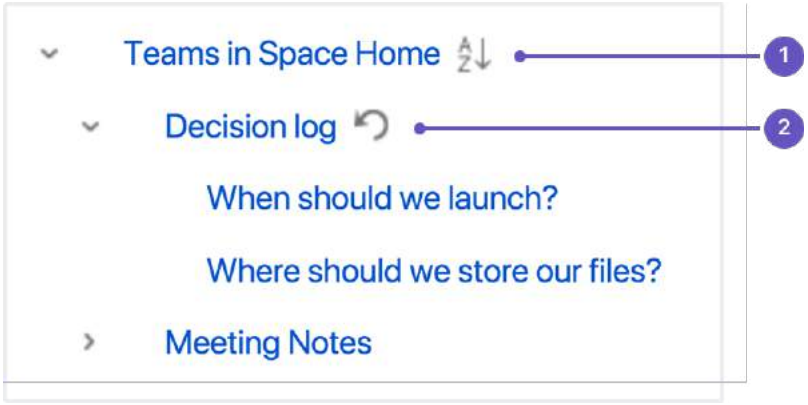

1. **Alphabetical**: sort all child pages alphabetically.
2. **Undo**: undo sorting.

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.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Copy a Page

If you need to duplicate the contents of a page, the easiest way is to copy the page.You can copy pages into the same space or to a different space.

When you copy pages into the same space, you'll need to rename them, as two pages with the same name can't live in the same space. We give you some handy tools to help rename your pages during the copy process.

You need the 'Create Page' permission to copy pages. Seespace permissions for more information.

On this page:

- Copy pages

Related pages:

- Create and Edit Pages
- Move and Reorder Pages

Copy pages

Single page

To copy a single page:

1. Go to the page and choose ••• >**Copy**.
2. Choose a**Location**for the new page.
3. Choose whether to **include attached files and images**.
4. Hit**Copy**.

Confluence will open the copy of the page in the editor and name it 'Copy of [original page title]'. You can then rename the page and work in the editor like any ordinary page.

Any restrictions are not copied over. If the page contains private information, click the padlock icon in the editor to apply restrictions *before* you publish.

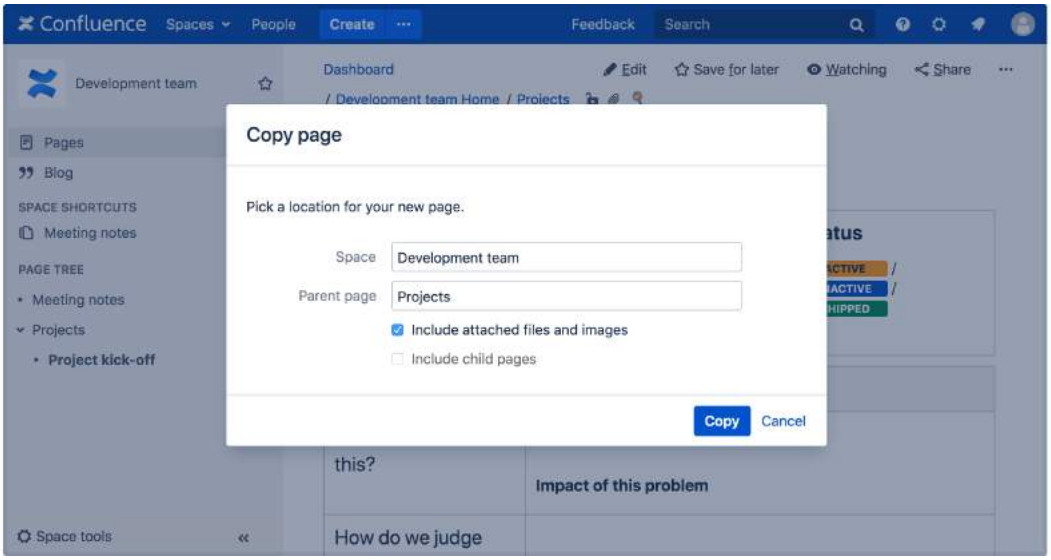

Page hierarchy

To copy a page and all its child pages:

1. Go to the page and choose ••• >**Copy**.
2. Choose a **Location** for the new pages.
3. Select**Include child pages**.
4. Use the options provided to customize your new page titles.
   You can add a prefix, replace a keyword or phrase in existing titles, or both.The preview will give you a good idea of what your new page hierarchy will look like, and warn you if there are any problems.

5. Deselect any items you do not want to copy over (attached files, labels, restrictions)
6. Hit **Copy**.

You'll find a link to your newly copied pages in the copy complete message.

⚠️Note:It's not possible to selectively copy multiple pages. You will be copying the entire hierarchy.

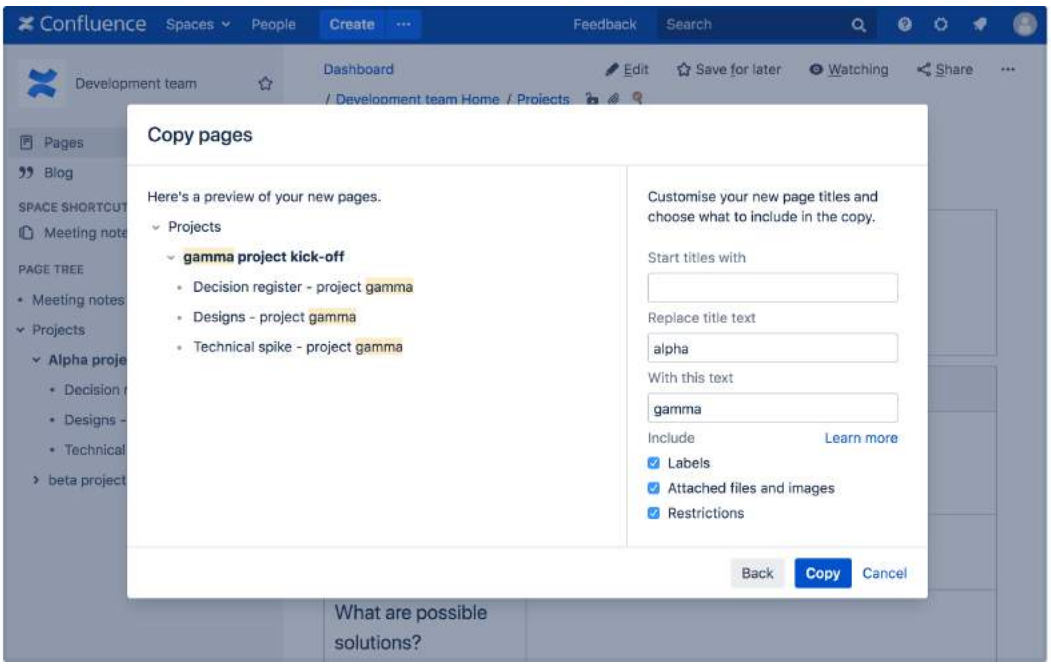

What's included when you copy a page?

Here's some more info on what's included when you copy a page.

	Single page	Page and its child pages
Page contents (text, macros)	Yes	Yes
Attached files and images	Optional - if you choose not to include attachments, you may see 'unknown-attachment' errors on the copied pages.	
Comments	No	No
Labels	Yes	Optional
Restrictions	No	Optional - you may not be able to change this option if you don't have appropriate 'Restrict' space permissions.
Watchers	No	No
Saved for later info	No	No
History	No	No
Child pages with view restrictions	-	Yes - if you have permission to see the child pages. No - if you do not have permission to see the child pages.

Space permissions and page restrictions

Your space permissions and any existing page restrictions have an impact on what you can copy:

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

- To copy pages you need the 'Create' page permission in the destination space.
- To copy pages with restrictions intact, you need the 'Restrict' page permission in the destination space.

When copying a page and its child pages, you also have a choice to copy over any existing page restrictions. This is useful if you need to maintain the current view or edit restrictions.If you don't have 'Restrict' page permission in the destination space, you won't be able to choose this option.

Limitations

You can copy hierarchies containing up to 2000 pages. Administrators can increase or decrease this limit by setting the`confluence.cph.max.entries`system property.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Delete or Restore a Page

When you delete a Confluence page, we'll move it to the space's trash. It's not permanently deleted, and can be restored by a space admin, until the page ispermanently deleted from the trash.

Don't see a delete option? Deleteonly appears if you are allowed to delete the page. Bothspace permissions andpage restrictionscan prevent you from deleting.

On this page:

- Delete a single page
- Delete a page hierarchy
- Delete an unpublished page
- Delete a page version
- Restore deleted pages
- Empty the trash or permanently delete a page
- Automatically purge deleted pages and files

Related pages:

- Cannot edit page: Editing or Deleting a Page That Won't Render
- Copy a Page

Delete a single page

When you delete a page in Confluence, you're deleting its page history too. 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:

1. Go to the page and choose••• > **Delete**.
2. We'll warn you of any issues, such as incoming links that will break. Choose **Delete** to proceed.

The page will be sent to the trash,where it can berestored by a space admin.

Any child pages (including any restricted pages that you are not allowed to see) will move up to the nearest parent page.

Delete a page hierarchy

Only users with Delete Page permission in a space can delete hierarchies.

When deleting a page that has child pages you have the option to delete the entire page hierarchy.

To delete a page hierarchy:

1. Navigate to the parent page and choose ••• >**Delete**.
2. Choose**Also delete child pages**then hit**Next**.
3. We'll warn you of any issues, such as incoming links that will break. Choose**Delete**to proceed.

The pages will be sent to the trash, where they can berestoredby a space admin.

Any pages that are restricted (that you are not allowed to see or delete) will not be deleted and will move up to the nearest parent page.

Users who only have Delete Own permission cant delete hierarchies, even if they created all of the pages in the hierarchy.

Delete an unpublished page

To delete a page that has never been published (a draft), in the editor go to ••• > **Delete unpublished page**.

Deleted drafts are not sent to the trash, so cannot be restored.If other people have contributed to the draft, you will be deleting their work as well as your own.

Delete a page version

Space adminscan delete specific versions from the page history. This is useful if you need to prevent older versions of a page being restored in future.Deleting a page version is permanent and cannot be undone.

To delete a specific version of a page:

1. Go to the page and choose ••• > **Page History**
2. Choose**Delete**next to the version you want to delete.

If you have Confluence Data Center, the remaining versions are not re-numbered, so that other users can see that a version has been removed. If you have Confluence Server, the versions will be renumbered.

Restore deleted pages

If you're aspace adminyou can restore deleted pages back into a space. This is useful if someone accidentally deletes a page and needs to get it back.

To restore a deleted page:

1. Go to the space and choose**Space tools**>**Content Tools**from the bottom of the sidebar
2. Choose**Trash**.
3. Choose ••• > **Restore**next to the page you wish to restore.

Pages are restored to the root of the space. Head to Space Tools > Reorder Pages to drag your restored page back into the page hierarchy.

Empty the trash orpermanently delete a page

If you're a space adminyou canpermanently delete a page (and all its attached files) by purging it from the trash. Once purged, the page and all its versions and attached files will be gone for good.

To purge deleted pages:

1. Go to the space and choose**Space tools**>**Content Tools**from the bottom of the sidebar
2. Choose**Trash**.
3. Choose ••• >**Purge**next to a specific page or you can**Purge all**to completely empty the trash.

Automatically purge deleted pages and files

Confluence can automatically purge deleted items from the trash when they reach a certain age, based on the date they were deleted.

Deletion dates will not be accurate for any items deleted before Confluence 7.14. This is because we did not store this data in earlier Confluence versions. The date displayed in the trash page is the date your site was first upgraded to Confluence 7.14 or later. Anything deleted after this date has an accurate deletion date.

To find out how to automatically purge items from the trash see Set retention rules to delete unwanted data.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

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 categorize, 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 organize your Confluence spaces. SeeUse Labels to Categorize Spaces.

On this page:

- Label a page or blog post
- Label an attachment
- The Labeled content page
- Search by label
- Search for labeled pages using a URL
- Remove labels

Related pages:

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

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:

1. At the bottom of the page, choose **Edit labels** or hit**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 **Edit label** icon at the top of the page.

Labels can't contain spaces, are lower case, and can contain a maximum of 255 characters. You can add multiple labels by adding a space between each label, any capitals will be automatically converted to lower case. If you want a label to includemore than oneword, usean underscore or hyphen(the only two special characters that labels accept). For example*this_is_a_label*or*this-is-a-label*.

Label an attachment

1. Do either of the following:
   - Go to the page that contains the attachment and chooseGo 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 label** icon besidethe 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.

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 Labeled 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 **Labeled 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 labeled content from all spaces in your Confluence site.

Screenshot: The Labeled content page

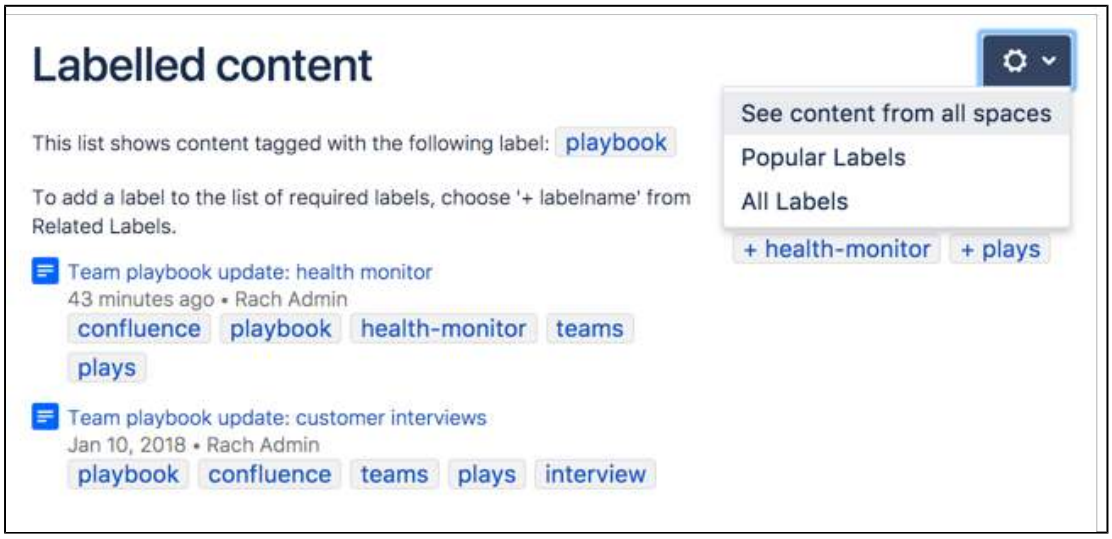

The **Popular Labels** option displays a word cloud, where the bigger a label is displayed, the more popular it is. Choose any label to view content tagged with that label.

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=SPACEKEY
```

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 labeled 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.site>/label/foo+bar`

The **Labeled 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:

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Once you're on the **Labeled 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.

You can't remove, consolidate or manage labels directly. A label is created by adding it to a page for the first time, and ceases to exist once it's been removed from all pages it was added to.

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

If you are using the Labels List Macro, the appearance of labels from deleted pages is considered a bug. Track the issue at CONFSERVER-55441 GATHERING IMPACT

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Display Pages with Label Macros

Using labels and macros, you can categorize 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 theContent by Label Macroto 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 Categorize 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 categorize 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

TheContent 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 page you've never published.**Unpublished changes**are edits that you've made to a published page, without republishing them.

Confluence autosaves your drafts and unpublished changes as you work, so if you get interrupted and close your tab or navigate away, your content lives on for you to resume editing when you're ready.

If you're creating or editing, but don't want to publish your changes yet, hit**Close**at the bottom-right of the editor. This will save those changes in the editor without publishing, and you can return to them at any point. Closing the editor will land you back on the published version of the page, or, if you're working on a draft, on your **Recently worked on** list.

On this page:

- Find drafts and unpublished changes
- Resume editing a draft
- Resume editing a page with unpublished changes
- Discarding unpublished changes
- Delete a draft
- Personal drafts

Find drafts and unpublished changes

Drafts and pages with unpublished changes appear in**Recently worked on**in the dashboard. You can easily differentiate between these as they'll have a 'draft' or 'unpublished changes' lozenge next to their titles. The 'unpublished changes' lozenge is only visible to people who have contributed to the draft or unpublished changes, so you don't have to worry about it distracting your viewers.

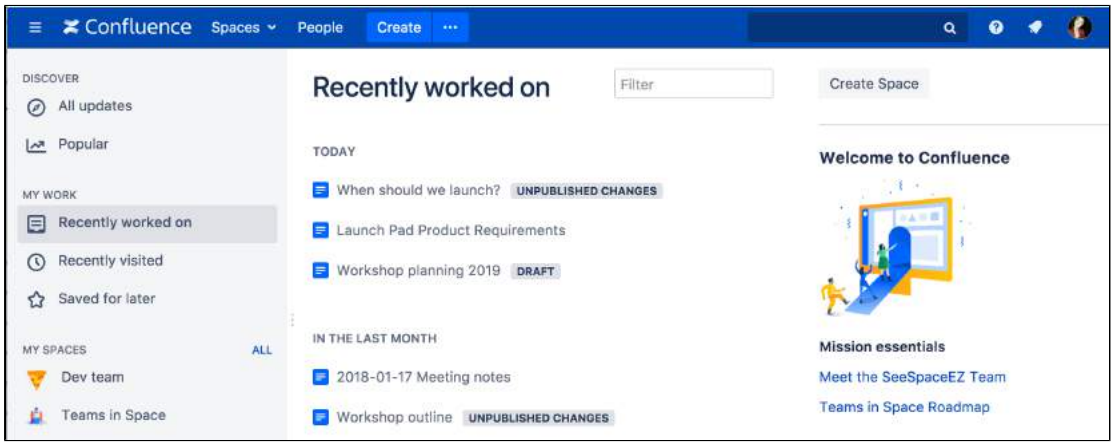

Resume editing a draft

You can find your drafts under**Recently worked on**or by heading to your profile and clicking on**Drafts** (only drafts that you created show in your profile). Clicking on a draft will drop you straight into the editor so you can keep editing and/or publish.

> If you didn't enter a page title, the draft will be called '*Untitled*'.

Resume editing a page with unpublished changes

If you've been editing a page that's already been published, you can find the page again either through the page tree or under your**Recently worked on**. The 'Unpublished changes' lozenge makes them easy to spot.

Edit the page to see the unpublished changes and keep editing, then, when you're ready, hit**Publish**.

Discarding unpublished changes

If you make changes to a published page, then change your mind, you can discard all changes by reverting to the last published version of the page.This will discard all unpublished changes made by you and any others who haveedited the page since the last time it was published.

Before you revert to the last published version you should:

- Check who else has edited the page since last publish - their avatars will be shown at the top of the editor.
- In the editor, go to ••• > **View changes** to see all changes that have been made since last publish. The changes won't be attributed to individual users.

Once you've checked to make sure you aren't going to inadvertently discard someone else's changes, go to ••• > **Revert to last published version** to discard all changes.

Delete a draft

To delete a draft go to ••• >**Delete unpublished page**.

Because drafts have never been published, you'll be deleting the entire page or blog post. Discarded drafts are not sent to the trash.

⚠ Drafts in Confluence are shared, meaning other people can work on them with you. If you delete a draft that other people have worked on, you're deleting their changes too.

Personal drafts

When collaborative editing is turned off, drafts work a little differently. Instead of a shared draft, you have a personal draft of a page. SeeConcurrent Editing and Merging Changes for more information.

You may see some old personal drafts in the Drafts page in your profile. These were created when collaborative editing was turned off.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Concurrent Editing and Merging Changes

This page covers the concurrent editing behavior in Confluence 6.0 or later when your administrator has chosen to **disable** collaborative editing.

In most cases, this won't apply to your Confluence site.

Related pages:

- Page History and Page Comparison Views
- Drafts

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: Concurrent editing notification

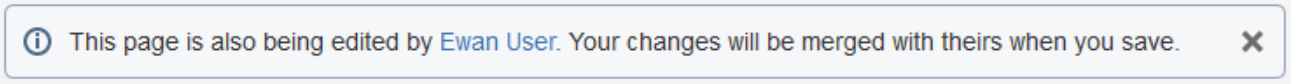

What happens if two of us are editing the same page and the other user saves before I do?

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 pagerestrictions, 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
- Copy a restricted page
- Get access to 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. Choose the Restrictions iconat 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. You can apply page restrictions to published and unpublished (draft) pages.

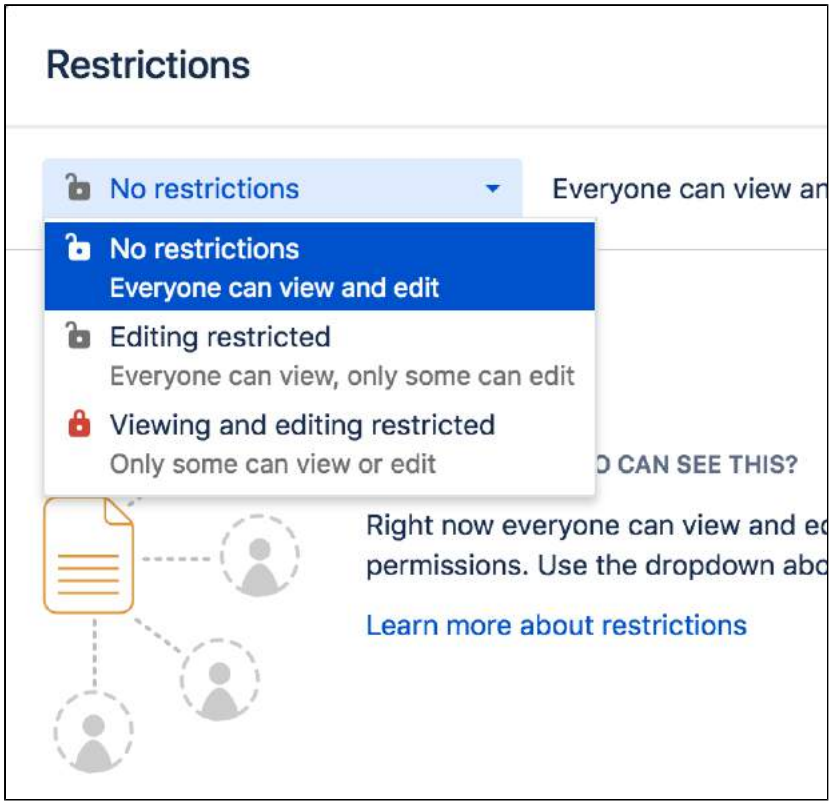

In this example, some users and groups can view only, others can also edit, plus there are inherited restrictions that might impact who can view the page.

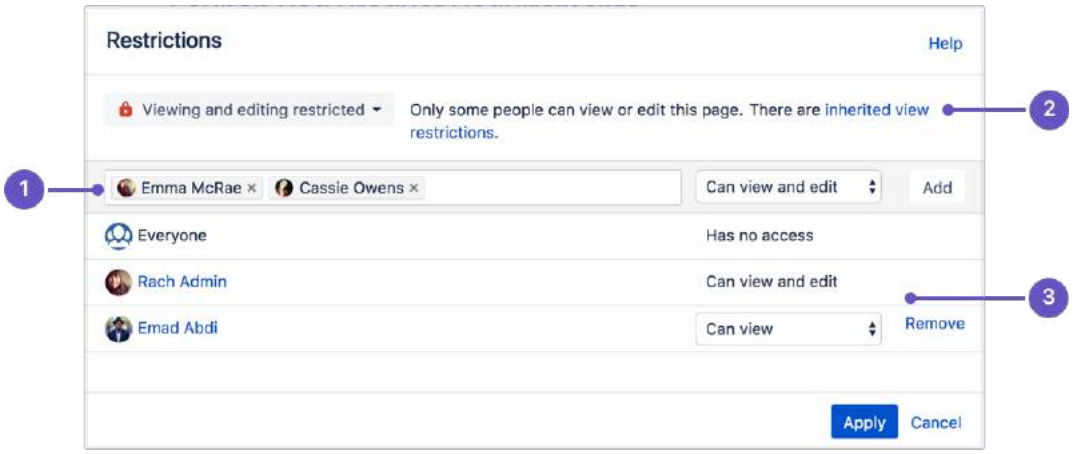

1. **Speed it up**: apply the same restriction to multiple people and groups.
2. **Watch out**: restrictions on other pages can affect this one.
3. **Be specific**: choose exactly what each group or person can do.

Who is 'everyone'?

When we say "everyone can view this page" *everyone* means all the people who can view the page by default. There are 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 allitschild 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 permissionfor 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.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

	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.

Copy a restricted page

When you copy a single page, we don't automatically copy the restrictions. If the page contains information that should be private, remember to reapply restrictions in the editor *before* you publish, to avoid notifying people who are watching the space.

When you copy a page and all its child pages, you have the option to copy all restrictions, or skip copying restrictions on all pages. SeeCopy a Page for more information.

Get access to a restricted page

If you navigate to a page that you're not able to view or edit because it has page restrictions applied (for example from a link, an invite, or page URL) you may be able to request access to the page.

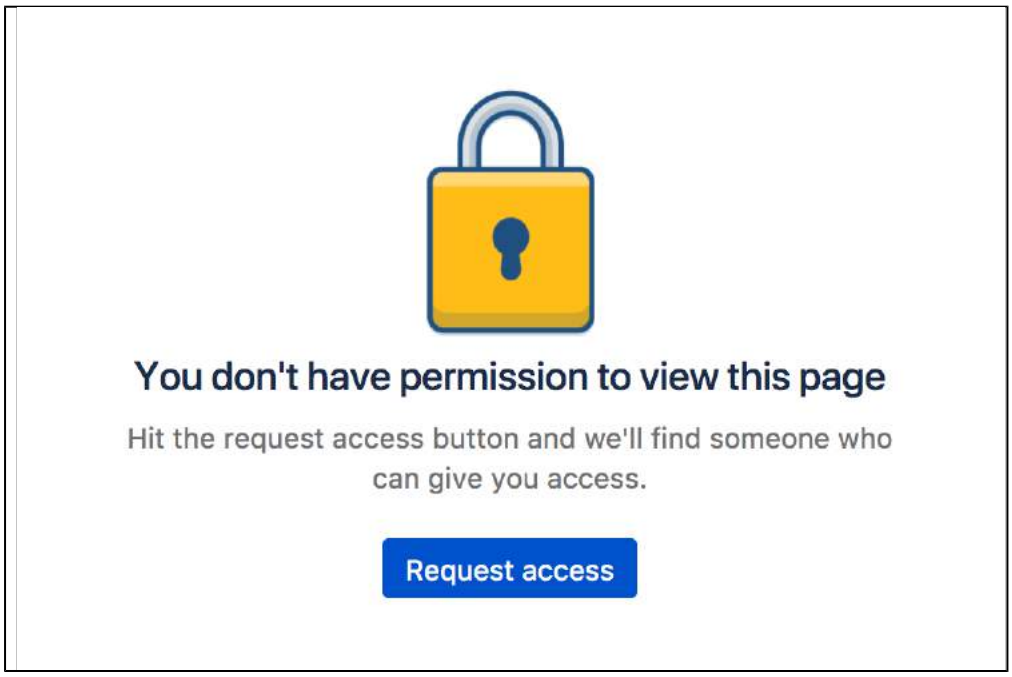

If the request access message above doesn't appear, you're not able to request access for that particular page. This usually is because the page has inherited view restrictions from a parent page, you don't have adequate space permissions, or there is no mail server set up.

Request access

To request access to a restricted page:

1. On the restricted page choose**Request access**.
2. Confluence will send an email to up to 5 people most likely to be able to grant you access.
3. Wait for an email confirming that access has been granted.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

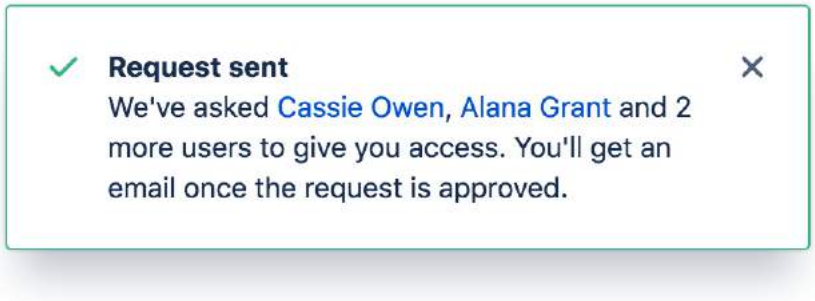

Grant access

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**.
   (We'll let you know if someone else got there before you, and has already granted access)

The user will receive an email confirming that access has been granted.

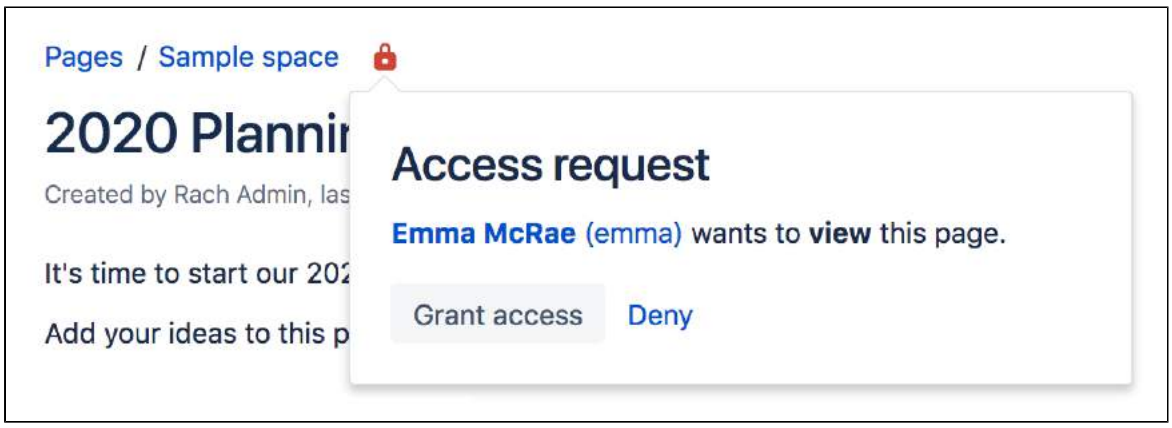

This process is the same as navigating to ••• >**Restrictions**and adding a 'View' restriction for the user.

Who can grant access?

When a user requests access to a restricted page, Confluence will send an email to up to 5 people who are most likely to be able to grant access, in the following order:

1. people who have contributed to the page in the past, can see the page and have 'Restrict' or 'Admin' space permission (sorted by last edit date)
2. space administrators who can see the page (sorted alphabetically).

This means that the request should be actioned quickly, as it prioritizes the people who have been interacting with the page most recently. There's no follow up email if none of the 5 people respond,the user will need to contact a space administrator directly to ask for access.

Disable the ability to request access

If you don't want people to be able to request access to restricted pages, for example if you're using Confluence for public documentation, you can disable the Confluence Request Access Plugin. SeeManaging System and Marketplace Apps.

View all restricted pages in a space

You need space admin permissions to view the list of restricted pages in a space.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

To view restricted pages:

1. Go to the space and choose**Space tools**>**Permissions**from the bottom of the sidebar
2. Choose **Restricted Pages**.

Screenshot: Restricted pages in a space

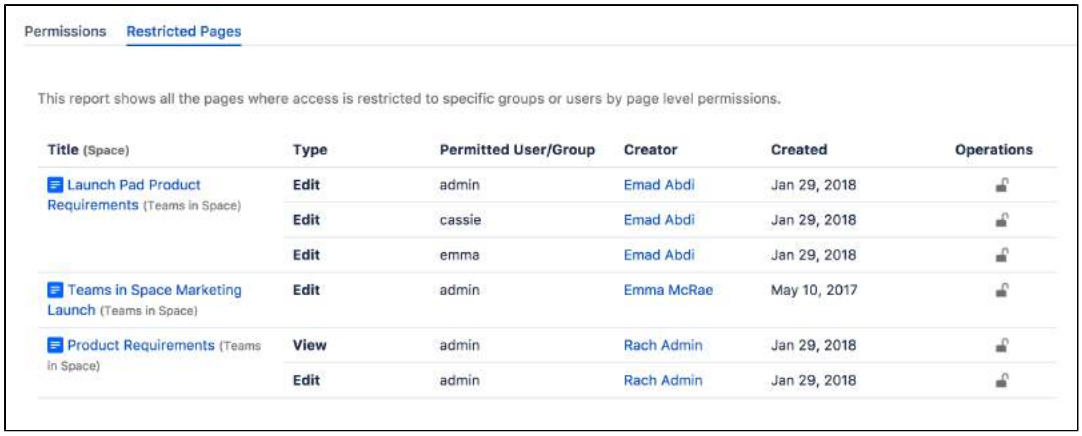

Notes

- **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**.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Links

You can create links to pages, blog posts, anchors, attachments, external websites, Jira issues and more. Links can be text or images, andcan 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.

On this page:

- Insert a text link
  - Other ways to do this
- Insert an image link
- Modify a link
- Remove a link
- Link to specific types of content
- Link to Confluence pages from other websites
- Link to a comment
- Using shortcut Links

Related pages:

- Anchors
- Inserting JIRA Issues
- The Editor

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 belowfor 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**

Other ways to do this

There are a few other ways to insert a link:

- Type**[**followed by the page or attachment name. Autocomplete will suggest matching items for you, or
- 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.

⚠️Confluence doesn't provide an option to configure a link to open in a new window or tab. Users can choose to right click / CTRL+click the link if they want to open it in a particular way.

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 (seebelowfor 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

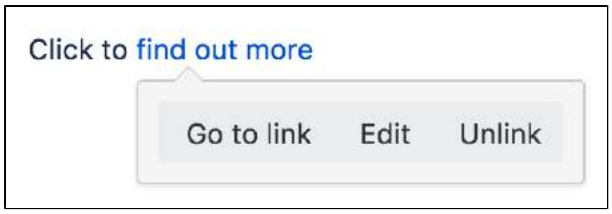

Link to specific types of content

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. or Choose **Link** > **Recently viewed** and select a page from the list. or Type **[** and enter part of the page name then select the page from the list. or 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. or Choose **Link** > **Advanced** then enter the space key followed by the page name `spacekey:mypage`. or 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. or Type **[** and enter part of the blog post name then select the blog post from the list.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Link to an attachment or image on this page	Choose **Link** > **Attachment** then upload or select an attachment from the list. or 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. or 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. or 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. or 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. or Type **[$** then enter the Comment ID ('12345' in this example): `[$12345]`

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

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). or 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. or 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 from another site or in an email, is to use the share 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**Share**.
3. Copy the**Share link**.

You do not need to use the share 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/Page+name#pagename-anchorname`
	
	Example from this page:
	
	`https://confluence.atlassian.com/display/DOC/Working+with+Links#WorkingwithLinks-InsertLinkAnchor`

Link to a heading (from an external website)	`http://myconfluence.com/display/spacekey/Page+name#pagename-headingtext`
	
	Example from this page:
	
	`https://confluence.atlassian.com/display/DOC/Working+with+Links#WorkingwithLinks-Insertalink`

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Link to a comment (from an external website)	`http://myconfluence.com/display/spacekey/pagename?focusedCommentId=commentid#comment-commentid`
	
	Example from this page:
	
	`https://confluence.atlassian.com/display/DOC/Working+with+Links?focusedCommentId=368640803#comment-368640803`

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 theAnchor.

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?focusedCommentId=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 haveconfigured shortcut linkson 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`producesthis 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.

SeeConfiguring Shortcut Linksfor more details.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Anchors

On this page:

- Step 1: Create the anchor
- Step 2: Create a link to the anchor
- Notes

Related pages:

- Links

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

Step 1: Create the anchor

Add the Anchor Macroto 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, thenselect the Anchor macro
2. Enter the **Anchor Name**(For example, 'bottom' or 'important information')
3. Choose**Insert**

Macro options (parameters)

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.

Step2: 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 anchoron 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

4. Enter or modify the**LinkText**(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**

- Anchor names are case sensitive.
- Enter page and anchor names with spaces when you link to them in the same Confluence site. A good way to do this is to copy the page title exactly as appears on the page. Don't copy the page URL, as we don't want to include the + symbol).
- If you're linking to an anchor on a different page that has special characters in its name, where the URL displays a page ID rather than a name, you should still enter the page name when linking to it.

Screenshot: The 'Advanced' option in the link dialog

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://myconfluence.com/display/DOCS/My+page#Mypage-bottom http://myconfluence.com/display/DOCS/My+page#Mypage-importantinformation

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.
- If the page name contains special characters, where the URL displays a page ID rather than a name, the link to an anchor will look more like this`http://myconfluence.com/pages/viewpage.action?pageId=54689987#Test-page1!-anchor`
  In this example the page title is `Test - Page 1!`and the anchor name is `anchor`.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

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.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Tables

Confluence tables allow you to present important information and discuss it with your team. Use familiar table formatting options resizing columns, coloring cells, rows and columns, aligning content, and sorting the table by clicking the column headers to view the information the way you like it.

Using Confluence Cloud? Check out our info on the new Confluence Cloud editor if your table looks like this one.

On this page:

- Insert a table
- Edit your table
- Shortcut keys
- Sort the table in view mode
- Sticky table headers in view mode

Related pages:

- Page Layouts, Columns and Sections
- The Editor

Insert a table

To create a table:

1. Hit the**Table**button in the toolbar
2. Click a cell in the drop-down to set the number of columns and rows in your table

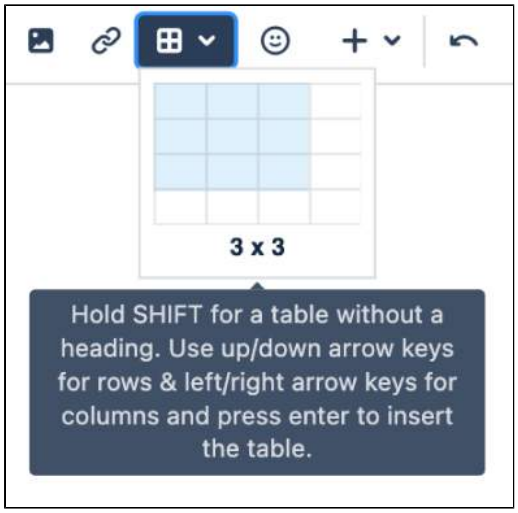

Edit your table

To resize table columns, just click and drag the column's border. To make other changes to your table, click inside it to reveal the table toolbar.

Here's a summary of the table tools:

Column width modes

- **Responsive** choose this mode if you want the table to expand as you add content. You can drag to resize the columns. It'll also resize itself to fit the page-viewer's window size (within reason).
- **Fixed width** choose this mode if you want to drag column borders to set width. Columns appear at your set size, regardless of content and window size.

	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)
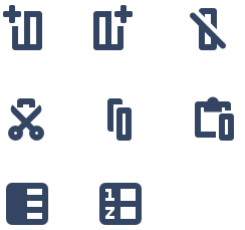	**Columns** • Insert columns before or after the current column • Delete the current column • Cut, copy and paste the current column • Mark a column as a header column (shaded with bold text) • Add a numbering column to automatically number each row
	Cells • Merge selected cells • Split selected cells • Change cell color
	Alignment • Align content to top • Align content to middle • Align content to bottom
	Table • Delete entire table

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

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

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, seeKeyboard shortcuts.

Sort the table in view mode

When viewing a table on a page, you can sort it by clicking the sort icons in the header row.

Screenshot: A colorful, sortable table

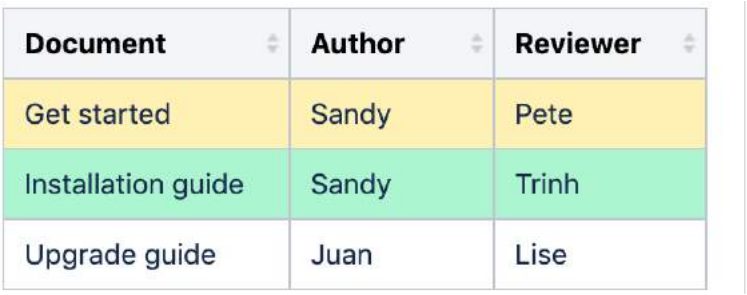

Document	Author	Reviewer
Get started	Sandy	Pete
Installation guide	Sandy	Trinh
Upgrade guide	Juan	Lise

The default sort order is the order the table rows are listed in the editor. You can use the **Cut row** and **Paste row** icons to move rows around in the editor.

Sticky table headers in view mode

In some instances the header rows of your table will stick to the top when you're scrolling down a page, making those*really*long tables easier to read.

You don't need to do anything to enable sticky table headers, however there are a lot of situations where headers won't stick. These include when your table:

- Is inside a page layout, inside another table, or inside a macro.
- Has no header row or there are cells in the top row that aren't marked as headers.
- Has a header column, instead of a header row, and scrolls horizontally.
- Contains another table, that has its own header row.

There's no way to freeze rows or columns in Confluence tables.

See

CONFSERVER-54343 - Table header with heading column not sticky when scrolling
LONG TERM BACKLOG

for issues with sticky table headers.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

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 requirements page 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.
- Tasks with no text will not be shown in the Task report macro.

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 color 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.

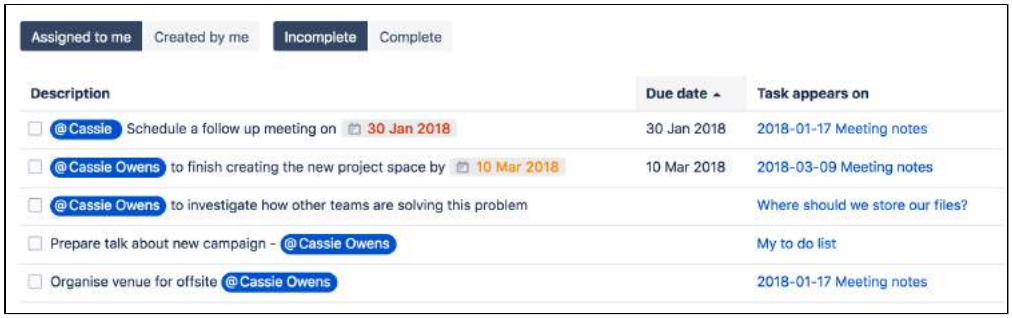

If you need to see more than just your last 6 months of tasks, use a Task Report.

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 parent 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.

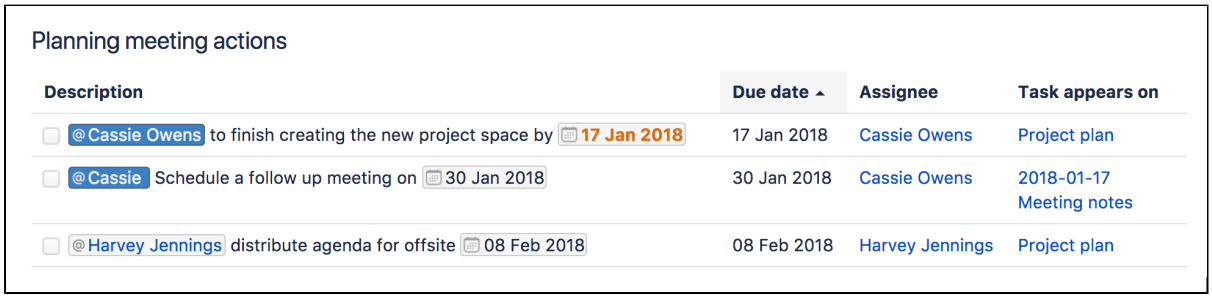

Notes

- 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 basedTasklist Macrohas 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.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Do more with Confluence

For even more ways to organize your tasks in Confluence, check out these apps from theAtlassian Marketplace:

- Comala Workflows: Improve document collaboration by assigning users to review, approve, and publish pages
- TodoMe for Confluence: Add, assign, and view tasks from every Confluence page
- Agile Retrospectives for Confluence:Interactively assign action items to team members

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Autocomplete for links, files, macros, mentions and emojis

When using the Confluence editor, you can type a trigger character or press a keyboard shortcut to see a list of suggested links, files, emojis, ormacros to add to your page, or to mention another user (and automatically notify them of this).

On this page:

- Summary of autocomplete
- Using autocomplete for links
- Using autocomplete for images, videos, audio files and documents
- Using autocomplete for macros
- Using autocomplete for mentions
- Using autocomplete for emojis
- Canceling autocomplete
- Enabling and disabling autocomplete
- Ignoring autocomplete

Related pages:

- Links
- Using Images
- Macros
- Keyboard shortcuts
- Your profile and settings

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...
Add an emoji on your page	**:**	None	See a list of suggested emojis to add to your page 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

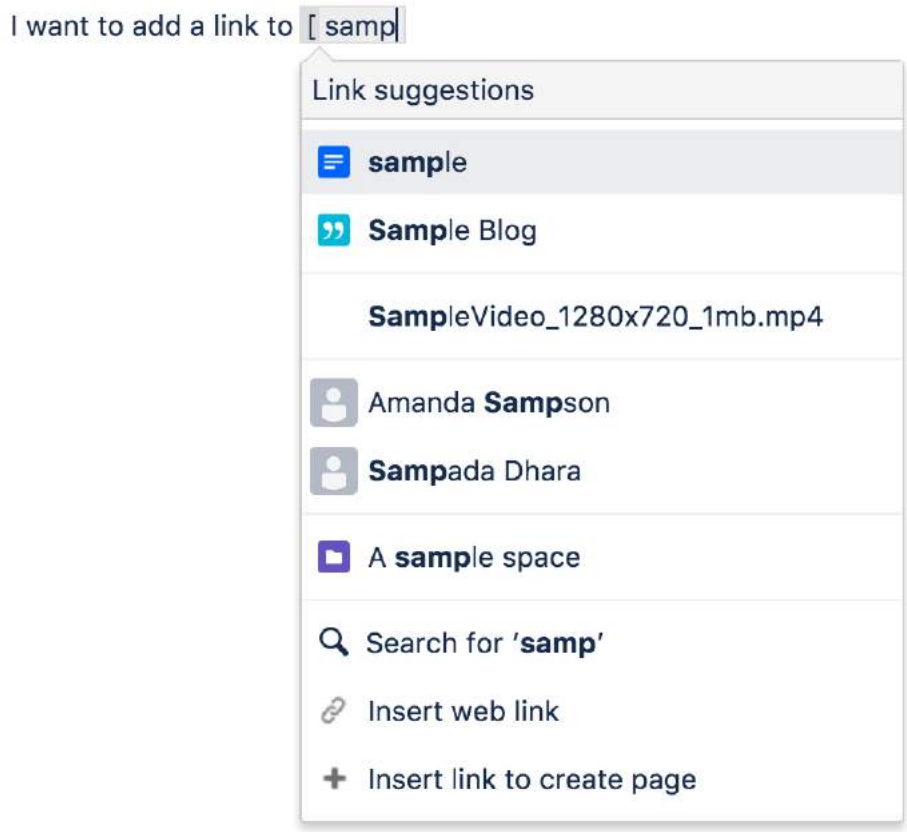

Using autocomplete for images, videos, audio files and documents

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 (or Cmd)+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.

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 (or Cmd)+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 file library** 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.

Screenshot: Autocomplete for an image or file

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

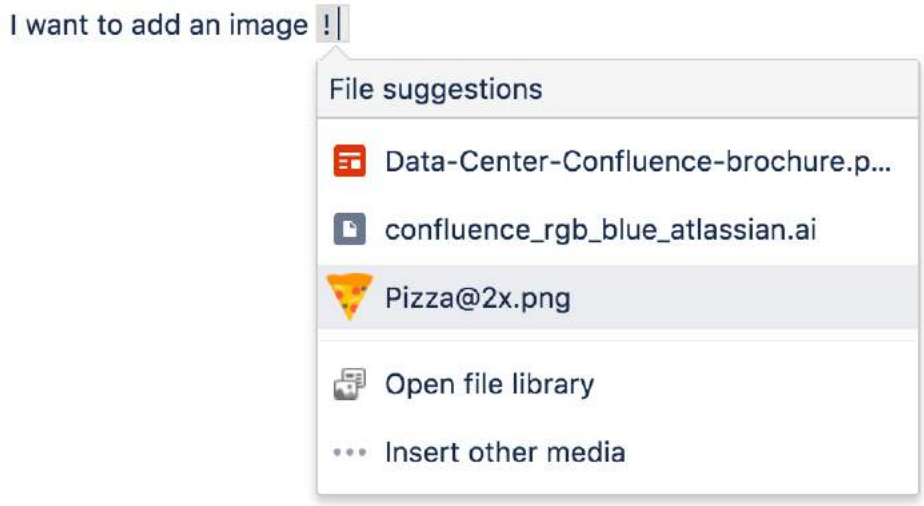

Using autocomplete for macros

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. SeeMacros.

Screenshot: Autocomplete for a macro

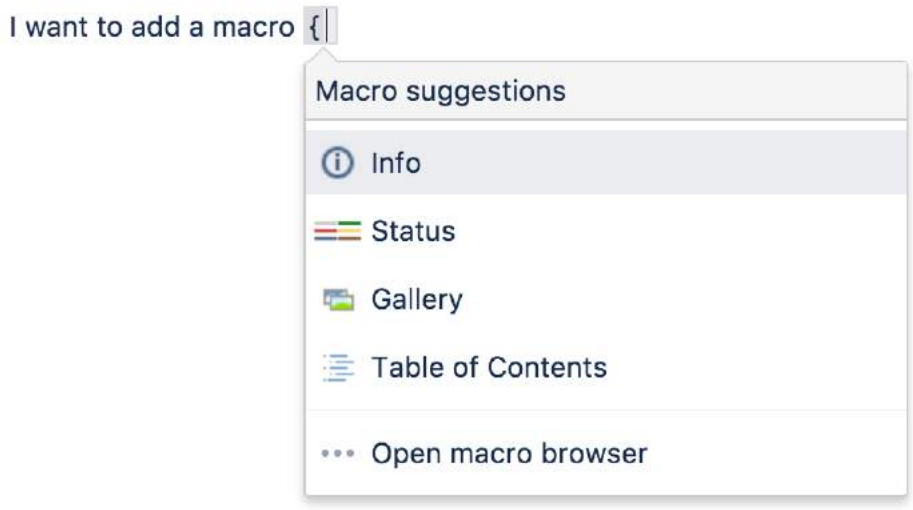

Using autocomplete for mentions

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

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 recognize users' full names only, not their usernames.

Using autocomplete for emojis

You can use the autocomplete as a fast way of searching for and inserting emojis to your page. Type ':' to see a list of emoji suggestions. If you continue typing, the list will change to predict and match your search.

To autocomplete an emoji:

1. Edit the page.
2. Click where you want to insert the emoji, and do one of the following:
    - Type ':' and then the first few characters of the emoji name or shortcut
    - Type ':' and then type the complete emoji name or shortcut
    - Type ':' and browse the emoji suggestions with your mouse or up/down arrows
3. Select the relevant emoji from the list of suggestions in the menu

Alternatively, if you type the complete emoji shortcut (which starts and ends with colons ":", for example `:wave:`), it will autoformat on the page. See Symbols, emojis and special characters.

Screenshot: Autocomplete for an emoji

Canceling 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: **]**

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

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

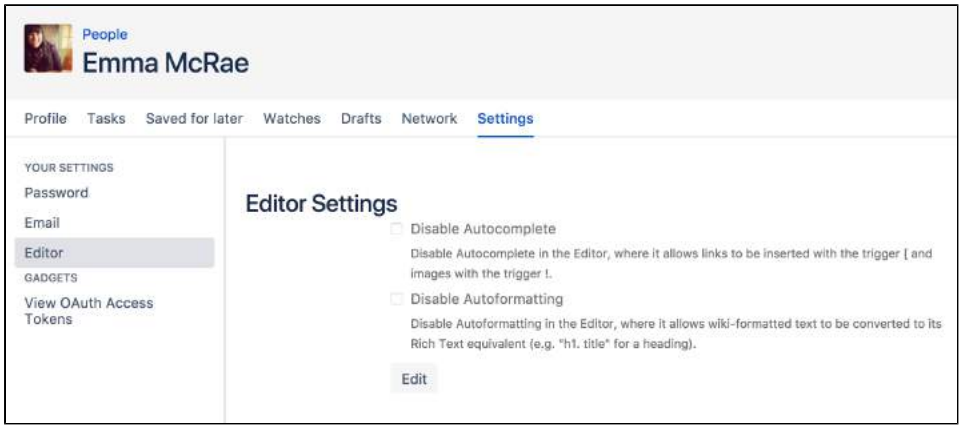

Ignoring autocomplete

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 > Keyboard Shortcuts** from the Confluence header.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

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

Using Confluence Cloud?Check out our info on how toformat your pagein the new editor.

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.

If you need to remove all layouts from a page, move all the content into one section, and set that section to be a single column. When you publish the page the layout will be removed.

Notes about Page Layouts

- **Column width** The width of the columns are a fixed percentage. 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

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

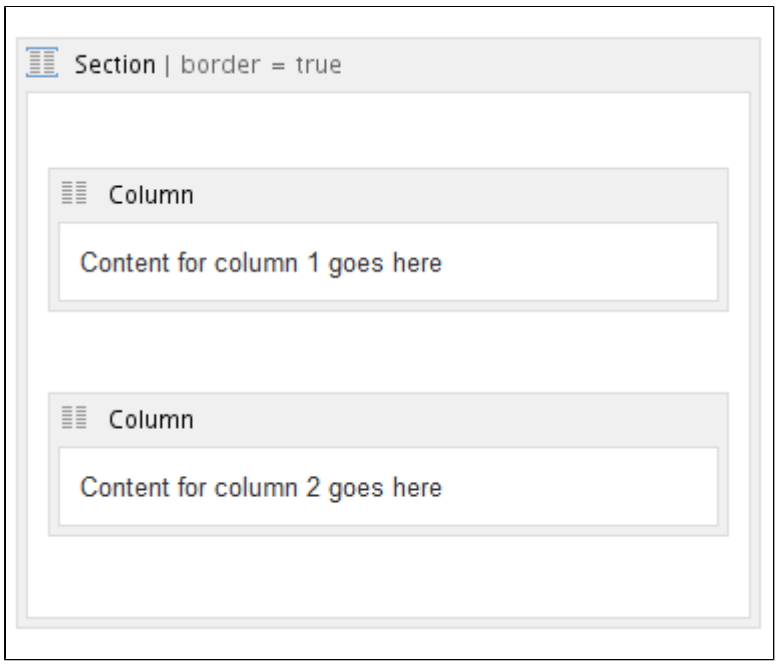

Macro parameters

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 aColumn macro, otherwise the section layout will not work as expected.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

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 summarizes those features and provides links to detailed instructions.

On this page:

- Add visual appeal
- Bring numbers to life
- Display presentations and documents
- Pull in content from Jira applications
- Tell a story in pictures
- Vary the structure of your pages
- Integrate your content with social media
- Show activity streams

Related pages:

- Macros
- The Editor
- Create and Edit Pages

Add visual appeal

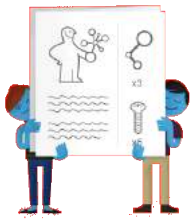

Pictures, photographs and screenshots. Confluence pages can display images from your Confluence site and from other websites. To put an image into Confluence, you can upload it and attach it to a page or blog post, then display it on any page, blog post or comment. Alternatively, display a remote image 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

Unknown macro: 'chart'

Display presentations and documents

Display your Office documents and other presentations directly in Confluence.

- 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.

Illustration: A PowerPoint slide deck

Pull in content from Jira applications

Many project teams and customers also use Jira applications such as Jira Software or Jira Service Management. 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 theJira Issues Macro.

Tell a story in pictures

A number of Marketplace apps for Confluence 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 apps.

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 onadd-on support.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

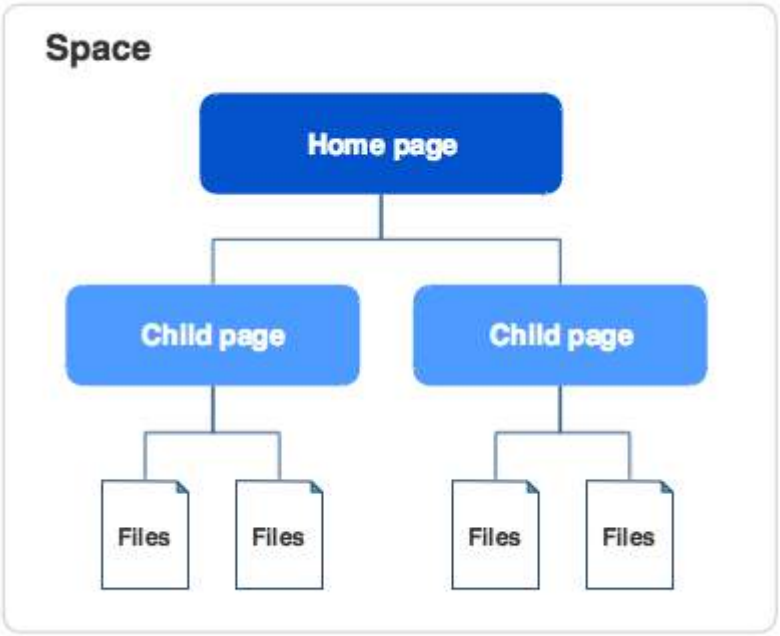

Illustration: A Gliffy diagram

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

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.

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.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Illustration: Twitter stream via the Widget macro

Show activity streams

Make your Confluence pages dynamic and interactive with:

- An activity stream showing updates and comments on Confluence and other linked applications. SeeGadgets.
- An RSS feed from within Confluence or an external site. SeeSubscribe to RSS Feeds within Confluence.
- A list of recent blog posts from within Confluence. SeeBlog Posts Macro.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

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 Marketplace apps, and you can even create your own templates.

Some examples of usefultemplates are:

- The meeting notes templatewill 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
- Create a Page from a Template

Looking for new Confluence templates? A huge range of templates are now available in Confluence Cloud. Learn more about templates in Confluence Cloud.

These templates are not available for Confluence Server or Data Center.

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 havespace 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.

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 Marketplace apps provide enhanced template functionality. You can search the Atlassian Marketplace for template apps.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 organize content in Confluence, and there's a collection of predefined ones that ship with Confluence.You can also download additional blueprints from theAtlassian Marketplace. You can customize the blueprint templates to suit your individual needs, disable particular blueprints or evendevelop 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.

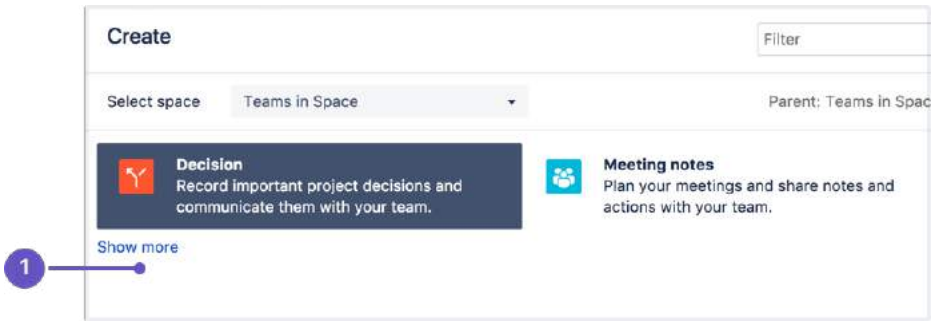

1. **Show more**: see more templates and blueprints.

To promote a template or blueprint:

1. Go to**Space Tools**>**Content Tools**
2. Choose**Promote**next to the templates or blueprints you want to appear in the Create dialog. You can only promote templates created in this space.

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.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

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 havespace 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.

On this page:

- Add a 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

Add a template

To create a new 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**.

To create a new global template:

1. Go to > **General ConfigurationGlobal Templates and Blueprints** .
2. Choose**Add New Global Template**

Looking for new Confluence templates? A huge range of templates are now available in Confluence Cloud. Learn more about templates in Confluence Cloud.

These templates are not available for Confluence Server or Data Center.

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 pageor 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

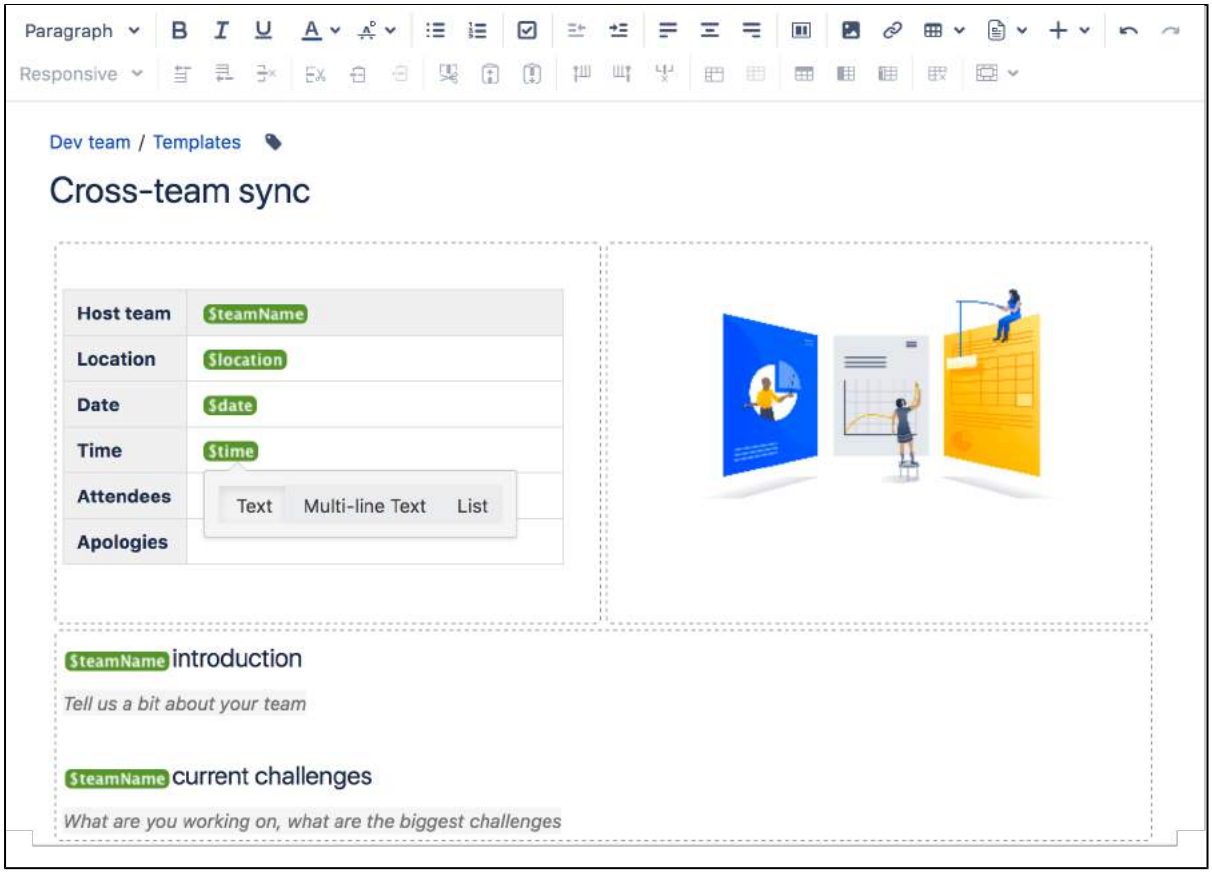

Screenshot: The form displayed when you create a page based on the template

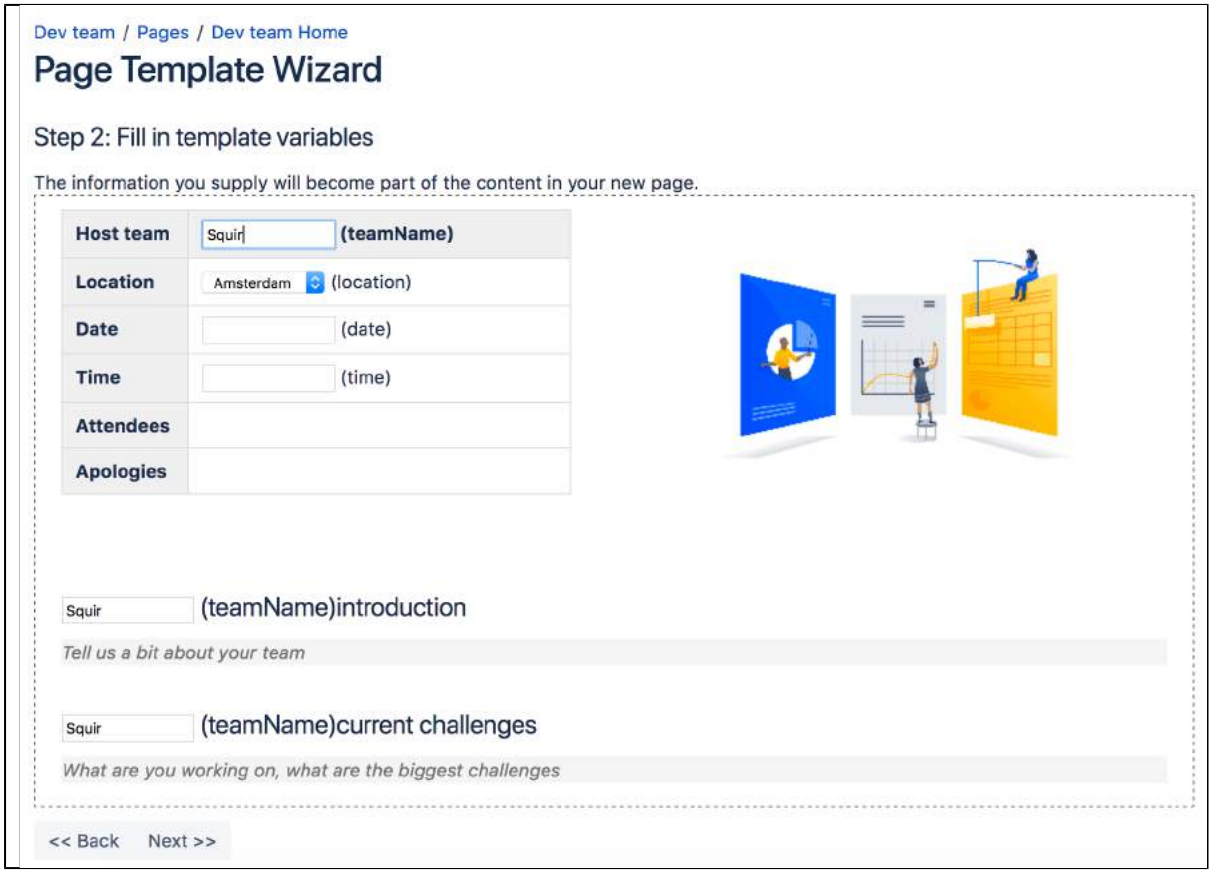

Template variables

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

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. Create a new template or **Edit** a template.
2. From the editor toolbar, select then choose**New Variable** (or choose an existing variable to add it to the page).
3. Enter a name for the variable.
4. 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.

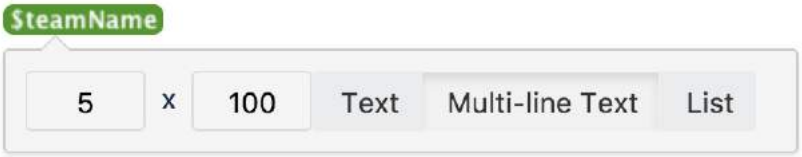

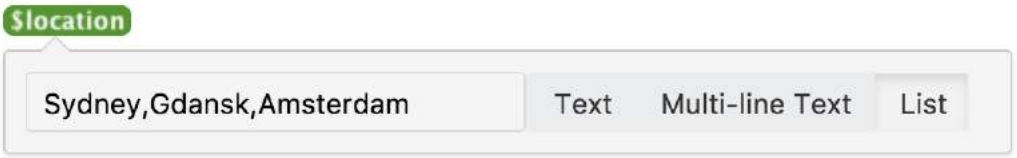

Hint: Type**$**and the variable name, then press**Enter**to add a new variable or to select an existing variable from a list of suggestions. The suggestions dialog shows variables already defined in this template.

Labels

If you'd like all pages created using this template to have one or more labels, choosethe labels icon next to the breadcrumbs at the top of the page to add them.

Note: Currently, it's not possible to edit the labels of blueprint templates.

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. Create a new template or **Edit** a template.
2. From the editor toolbar, select then choose **Instructional text.**
3. Type in your instructional text (for example, *Insert an image of the interface here.*)

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Instructional text appears in italics with a shaded background, todistinguishit from normal paragraph text.

You can also change the placeholder type from **Text**to either:

- **Usermention** Opens the user mention dialog.
- **Jira Macro** Opens a dialog 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**

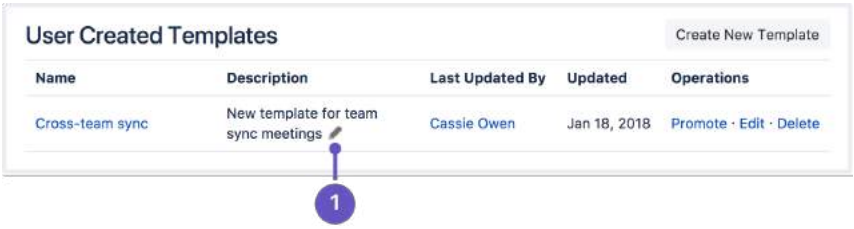

1. **Edit**: use the pencil icon to edit your template's description.

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 Marketplace apps provide enhanced template functionality. You can search the Atlassian Marketplace for template apps.
- 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.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

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).

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

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 **Cre ate from template** button.

1. **Create blank page**
2. **Create from template**

Looking for new Confluence templates? A huge range of templates are now available in Confluence Cloud. Learn more about templates in Confluence Cloud.

These templates are not available for Confluence Server or Data Center.

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

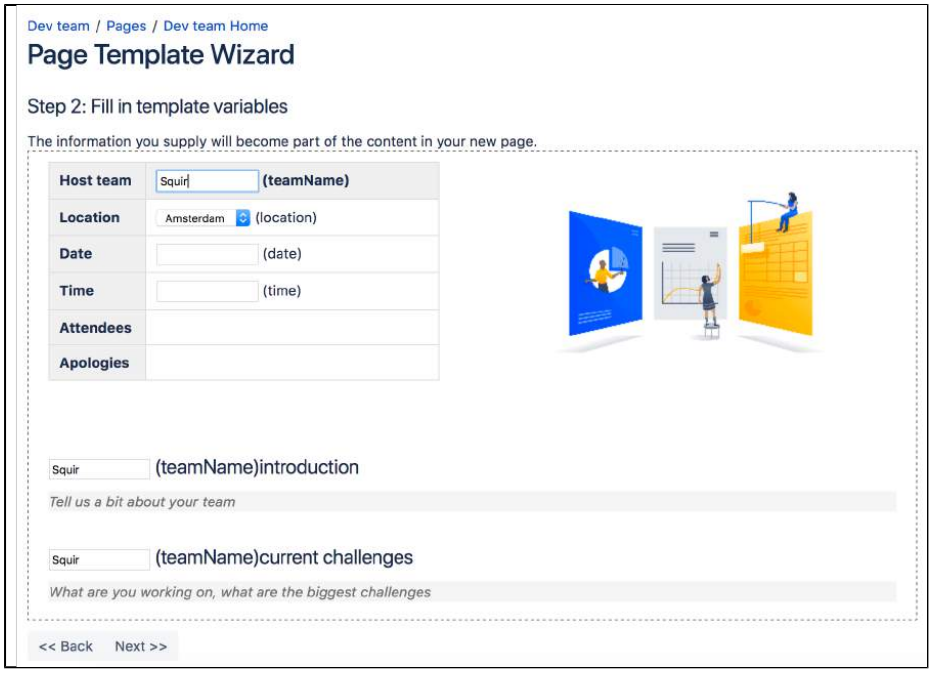

Notes

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 Marketplace apps provide enhanced template functionality. You can search the Atlassian Marketplace for template apps.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Blueprints

What's a blueprint?

A blueprint is a set of page templates with added functionality to help you create, manage and organize content in Confluence more easily.

Create meeting notes, shared file lists and requirements documentation out of the box, andCustomize the blueprint templates to suit your individual needs.You can even develop your own blueprints.

On this page:

- What's a blueprint?
- Create content using a blueprint
- Customize blueprint templates
- Promote blueprints in the Create dialog
- Add more blueprints
- Disable a blueprint
- Full list of blueprints

Related pages:

- Page Templates
- Request Marketplace Apps

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.

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 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. Here's the index page for the Meeting Notes blueprint:

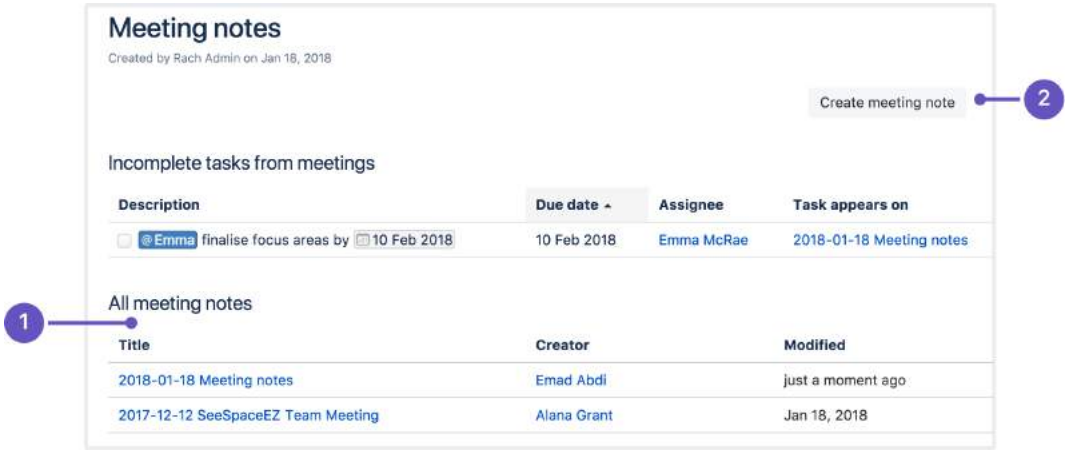

1. **Easy to find**: notes from all your meetings are listed here.
2. **Start a new meeting**: create a new meeting notes page here.

Customize blueprint templates

Blueprints are made up of templates that can often be customized 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 customize blueprint templates for the spaces you are an administrator of. You must be a Confluence Administrator to customize blueprint templates for a whole site. SeeAdministering Site Templatesfor more information.

To customize 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
- Not all blueprints are customizable. Some, including the Team Playbook blueprints (health monitor, DACI, project poster, and experience) can't be edited.
- Currently, it's not possible to edit the labels of blueprint templates.

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

SeeWorking With Templatesand Administering Site Templatesfor 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.

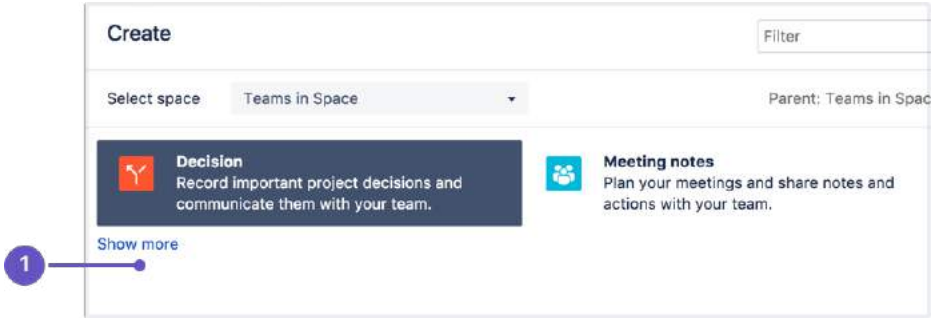

1. **Show more**: see more templates and blueprints.

To promote a template or blueprint:

1. Go to**Space Tools**>**Content Tools**
2. Choose**Promote**next to the templates or blueprints you want to appear in the Create dialog. You can only promote templates created in this space.

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

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

You can find more blueprints for Confluence in theAtlassian Marketplace. Blueprints are managed using apps (also known as add-ons, or plugins).

SeeRequest Marketplace Appsfor information on how you can search for new blueprint apps and send a request to your System Administrator.

If you are a System Administrator, seeManaging System and Marketplace Appsfor information on how to install new blueprint apps.

You can also develop your own blueprints. See our developer documentation onWriting 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**(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 • DACI decision (Atlassian Team Playbook) • Experience canvas (Atlassian Team Playbook) • Health Monitor (Atlassian Team Playbook) • Project poster (Atlassian Team Playbook)	• Documentation space • Team space • Knowledge base space

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

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 Log 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.

1. **Create blank page**
2. **Create from template**

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 log page for the space you're in, and add a shortcut to it in the space's sidebar.

Here's how the decisions page looks in the editor:

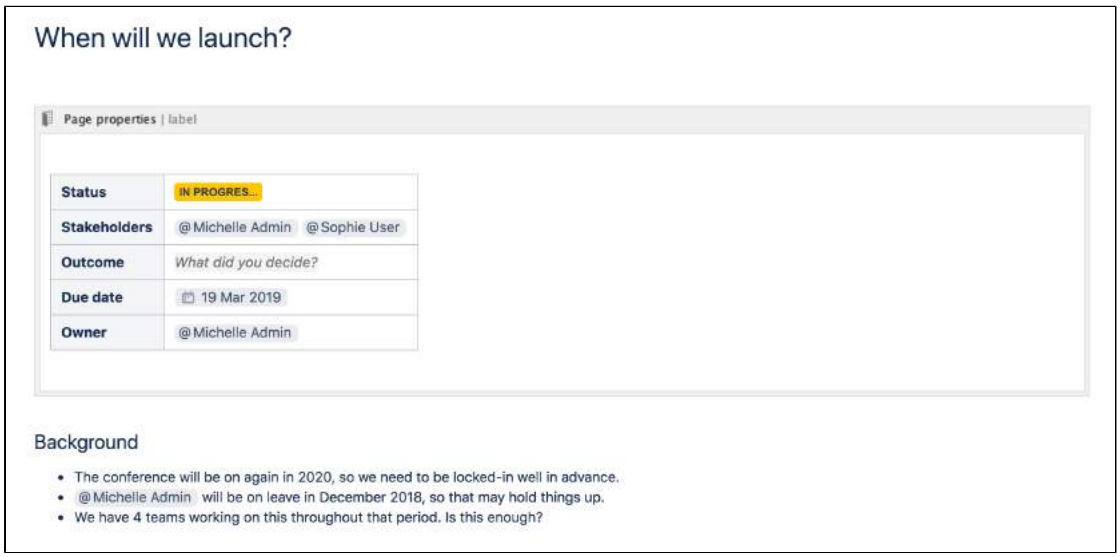

Once you save your first decision page, Confluence will create a decision log page for the space you're in, and add a shortcut to it in the space's sidebar. The decision log lists all the decisions in that space.

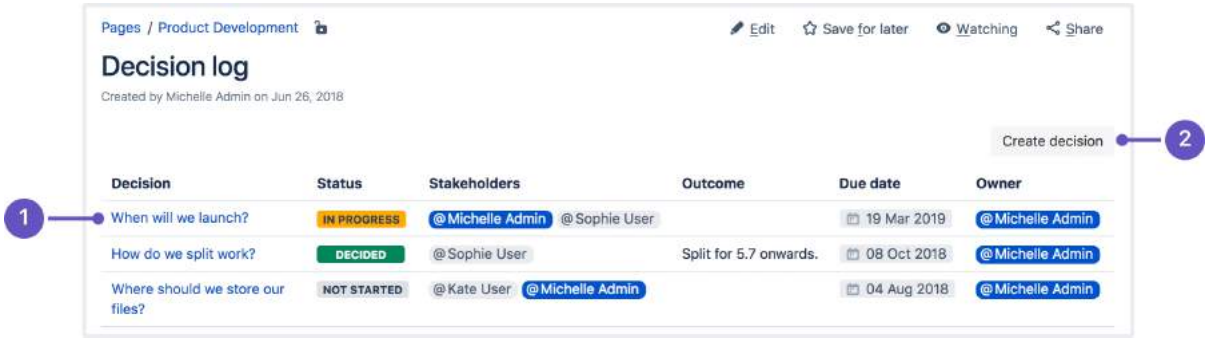

1. **Decision pages:**your existing decisions appear here.
2. **More decisions:**create new pages using the decision template.

The Decisions blueprint uses these Confluence features:

- Page PropertiesandPage Properties Reportmacro - 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.

Customizingthis blueprint

You can customize the templates that are used by the Decisions blueprint - seeCustomizing 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.

SeeInstructional textto find out more about using instructional text in templates.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

File List Blueprint

The File List blueprint helps you to create lists of files to share with your team. Great for organizing 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.

1. **Create blank page**
2. **Create from template**

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

Here's an example of a file list page:

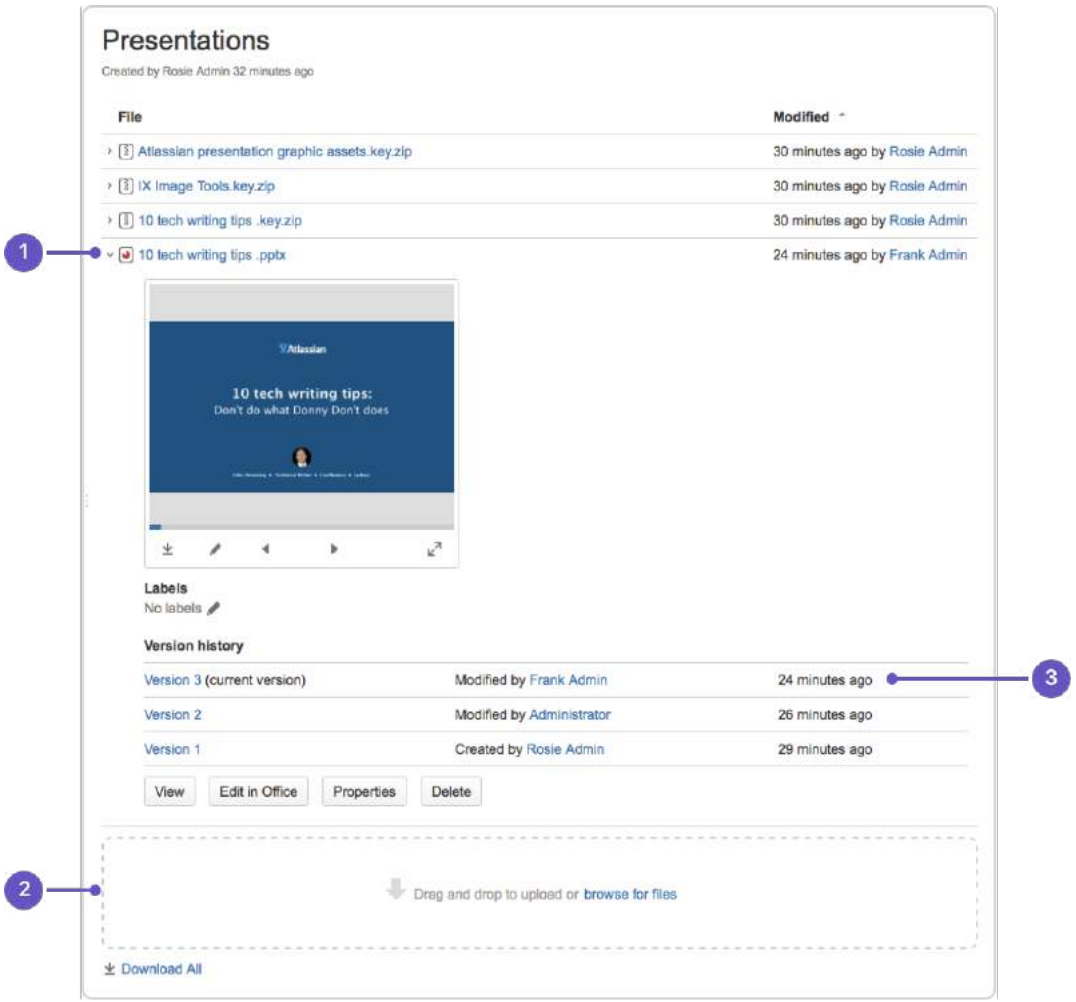

1. **Expand**: view each file's details and preview.
2. **Upload and download**: drop files to upload, and download all files.
3. **Version history**: see and manage other file versions.

Attachments appear on the page, and you can expand each attachment to preview the file and/or view its details.

In this example, we've created three file list pages 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.

Once you save your page, Confluence will create an index page and add a shortcut to your space sidebar. The index page lists the latest File List pages in the space. Create as many File List pages as you need.

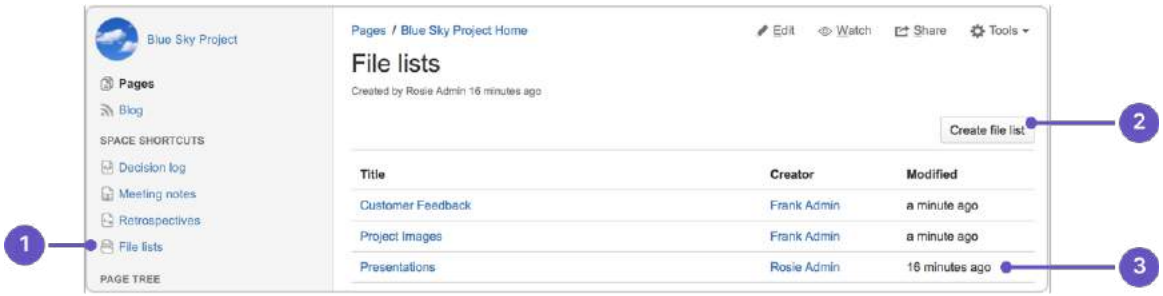

1. **Space shortcut**: a quick way to find all of your file lists in this space.
2. **Create a new list**: create more file lists in this space.
3. **Current file lists**: see all the file lists in this space.

Customizing this blueprint

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

You can customize the templates that are used by the File List blueprint - seeCustomizing blueprint templates.

The File List blueprint template uses the attachments macro. You can customize the macro to change the sort order or hide features such as version history and the upload attachment fields.

You can also edit theContent Report Tablemacro used on the Index page to specify the number of pages you want to display.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

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.

1. **Create blank page**
2. **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 @mentionusers to assign tasks to them.

Screenshot: A Meeting Notes index page

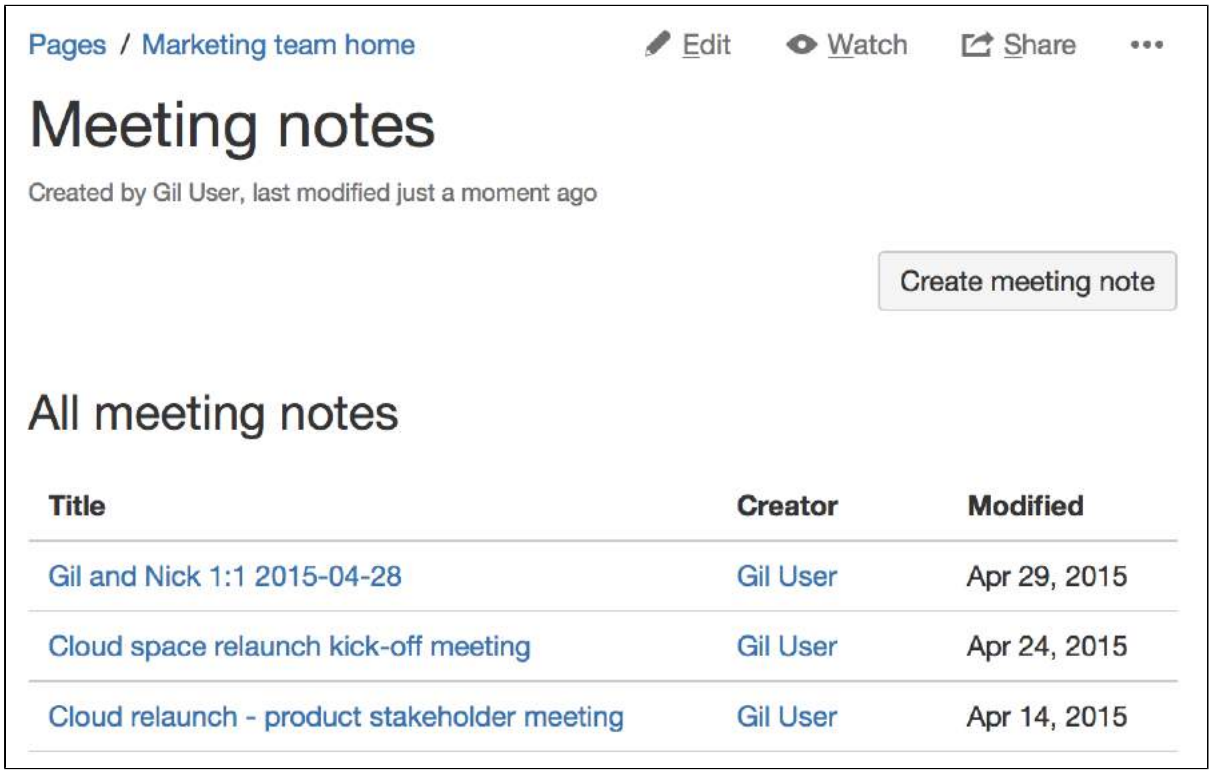

Screenshot: A blank Meeting Notes page showing instructional text

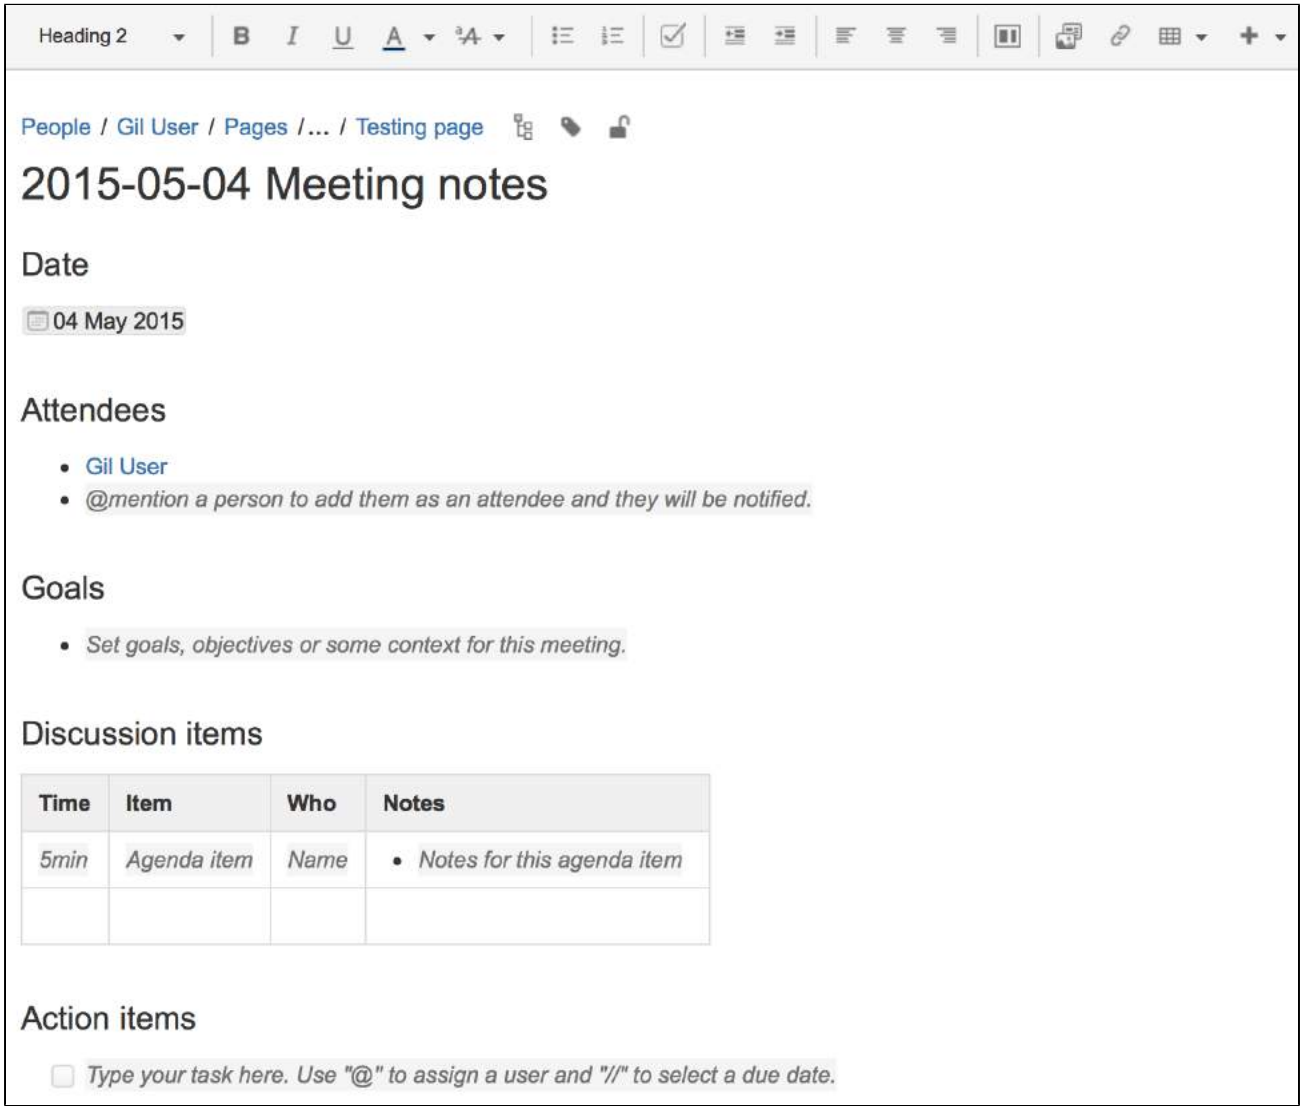

The Meeting Notes blueprint uses some cool Confluence features:

- Instructional text - this handy text prompts you to enter information anddisappearswhen 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.

Customizingthis blueprint

You can customize the templates that are used by the Meeting Notes blueprint see Customizing 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, seeInstructional text.

You can also edit theContent Report Tablemacro used on the Index page to specify the number of pages you want to display.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Product Requirements Blueprint

The Product Requirements blueprint helps you to define, scope and track requirements for your product or feature.

Learn more about writing downsized product requirements

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.

1. **Create blank page**
2. **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.

In this example we've created a series ofProduct Requirements pages. The index page shows summary information about each one.

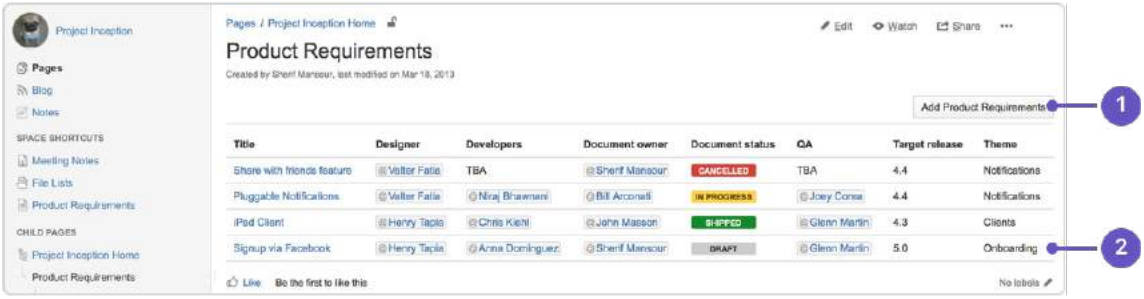

1. **New requirements**: create more requirements pages in this space.
2. **Requirements pages**: see the existing requirements pages in this space.

Here's how a requirements page looks in the editor:

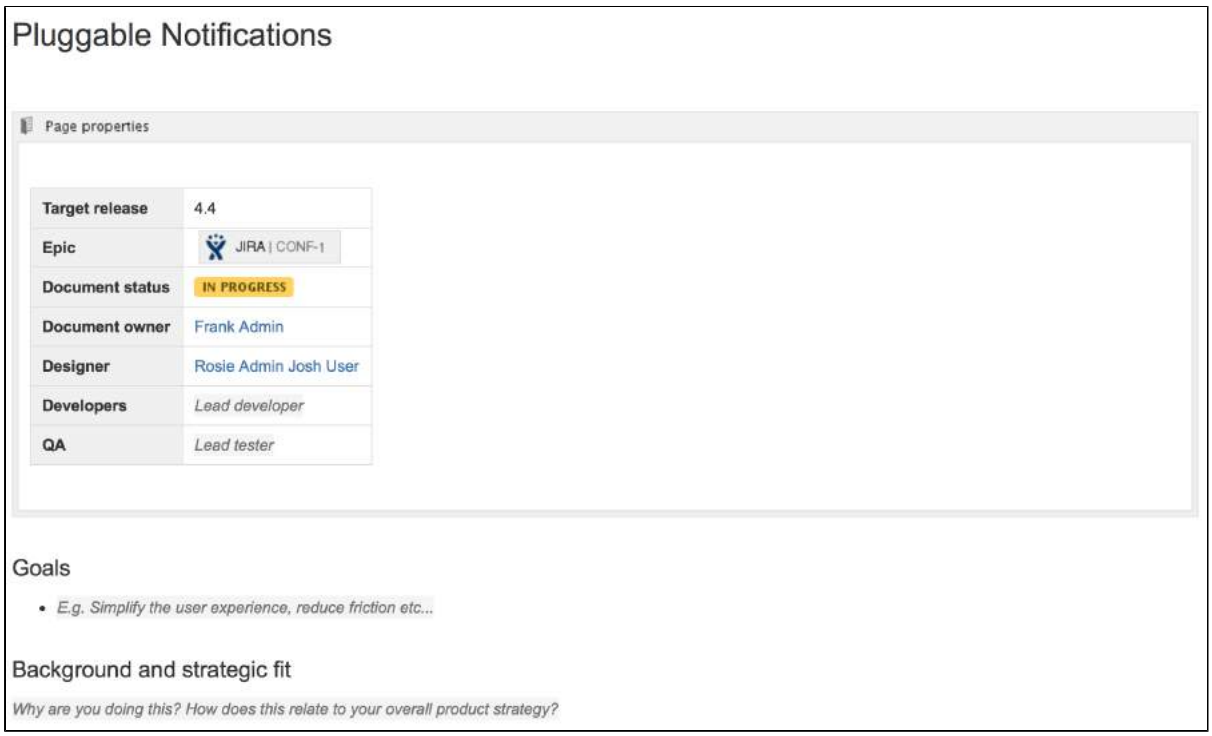

The Product Requirements blueprint uses these Confluence features:

- Page PropertiesandPage Properties Reportmacro - 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 anddisappearswhen you start typing or view the page.
- Mentions- @mention a user on the page and they will be notified in their workbox.

Customizingthis blueprint

As no two products or projects are alike, you can customize the templates that are used by the Product Requirements blueprint - seeCustomizing 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 PagePropertiesmacro.

SeeInstructional textto find out more about using instructional text in templates.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

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 centralized 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.

Related pages:

- Blueprints
- Decisions Blueprint
- File List Blueprint
- 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.

1. **Create blank page**
2. **Create from template**

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 categorize 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.

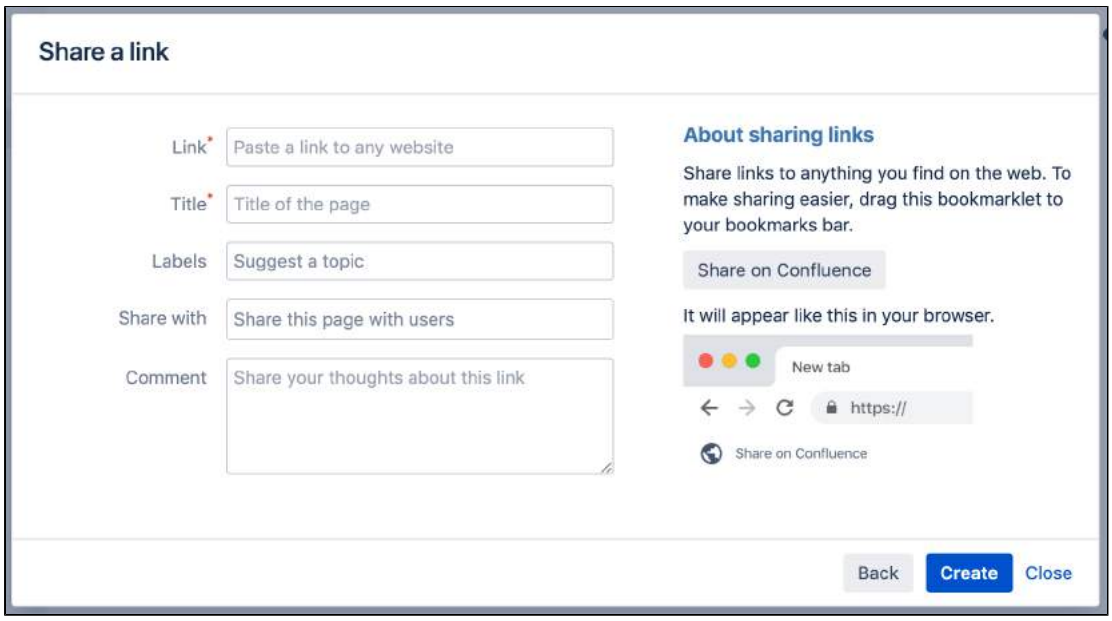

Security considerations and limitations

To prevent people from accidentally or maliciously sharing links that may pose a security risk to your site, domains must be added to the allowlist, before they can be shared using the share a link blueprint. For example if you wanted to use the blueprint to share links to this documentation site you would need to addhttps://confluence.atlassian.com/to the allowlist. The shared link blueprint is different to just inserting a link on a page because it shows a preview of the linked site. SeeConfiguring the allowlistfor more information.

If external connections are disabled for your site, you can still share a link (as long as it's in the allowlist), but we won't show a preview of its contents.Your admin can enable external connectionsat > **General Configuration** under **Connection Timeouts**.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

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).

On this page:

- Create a Change Log
- Create a Status Report
- Customizing this 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.

1. **Create blank page**
2. **Create from template**

To use the Jira Report Blueprint your Confluence and Jira Server, Data Center, or Cloud 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 **Jirareport** and hit**Next**
3. Select a **Change log**and hit **Next**
4. Enter the information required for the change logand hit**Create**

A report page will be created with sample text and a list of all issues for the project and fix versions selected, organized 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.

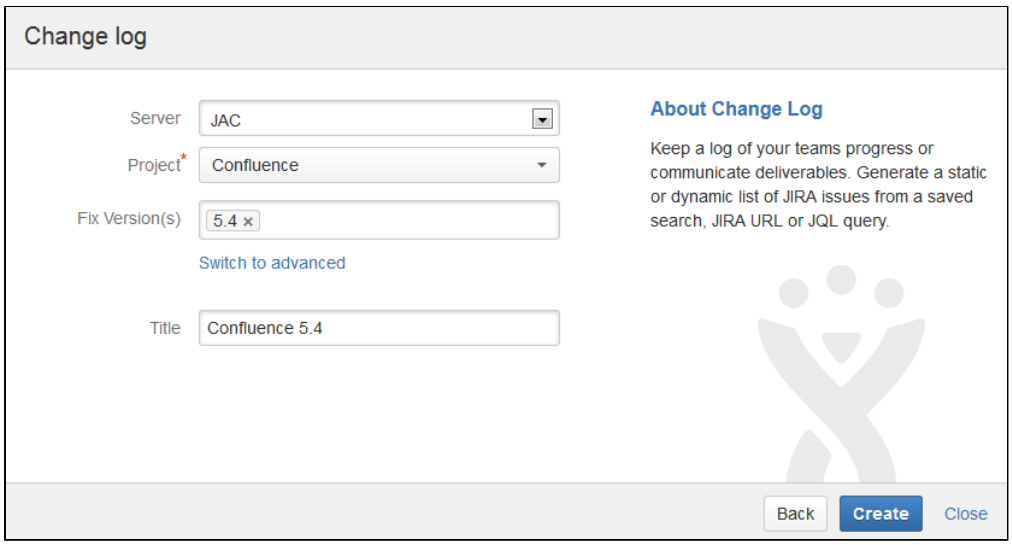

Screenshot: Static list of Jira Issues displaying in the Change Log.

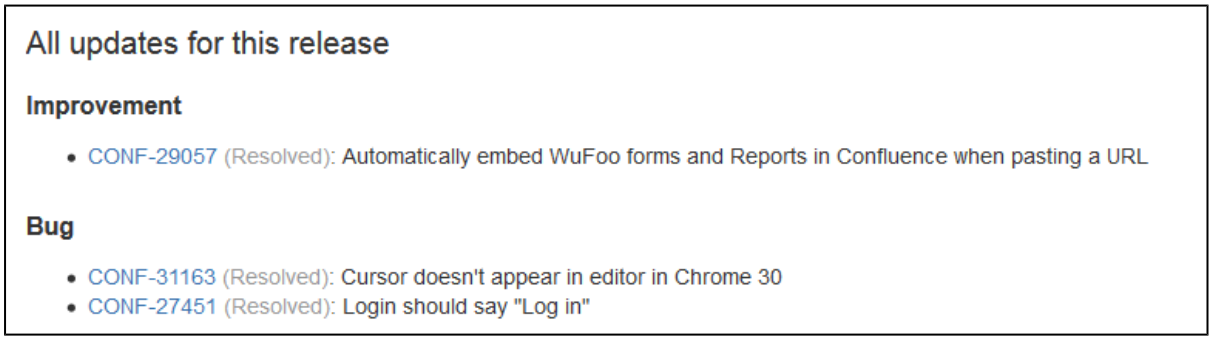

To create a dynamic change log:

1. Choose**Create from template** in the Confluence header
2. Select**Jirareport**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 theJIRA 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.

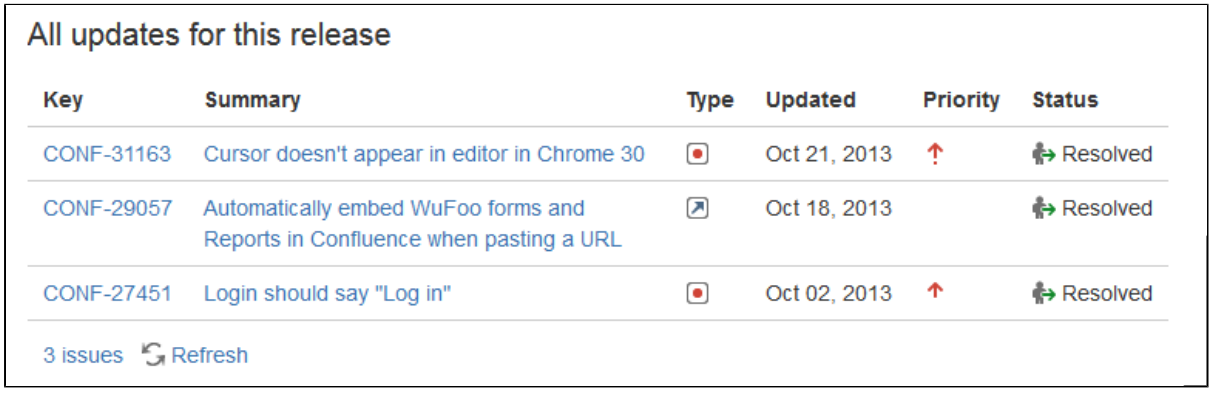

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

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**Jirareport**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.

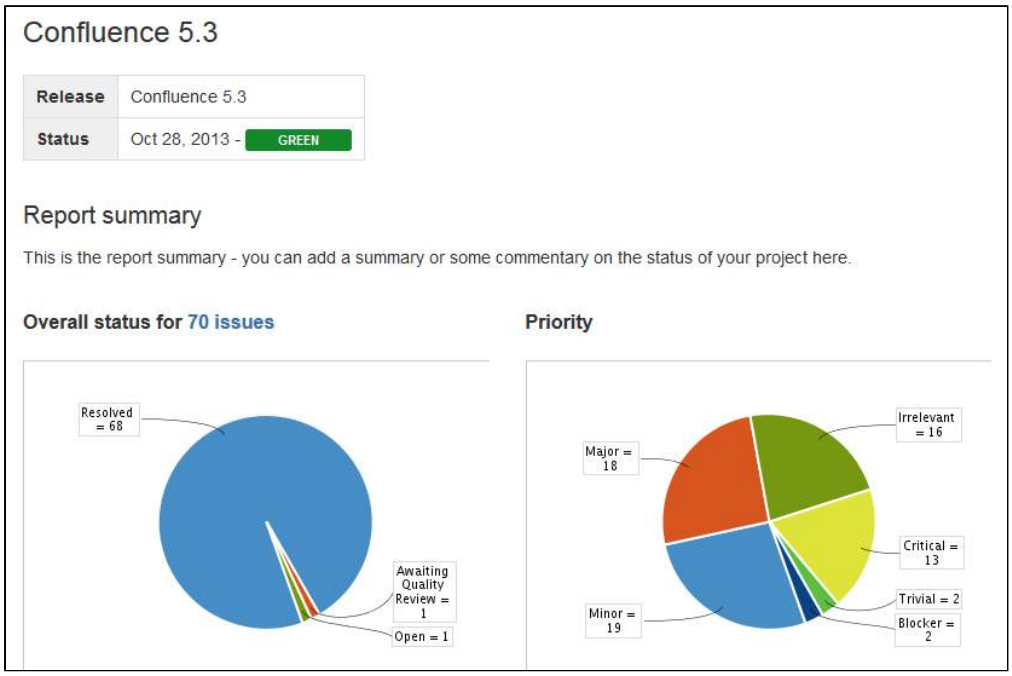

Customizingthis blueprint

You can customize 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**. SeeCustomizing 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 SeeInstructional texttofind out more about using instructional text in templates.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

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.

1. **Create blank page**
2. **Create from template**

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**

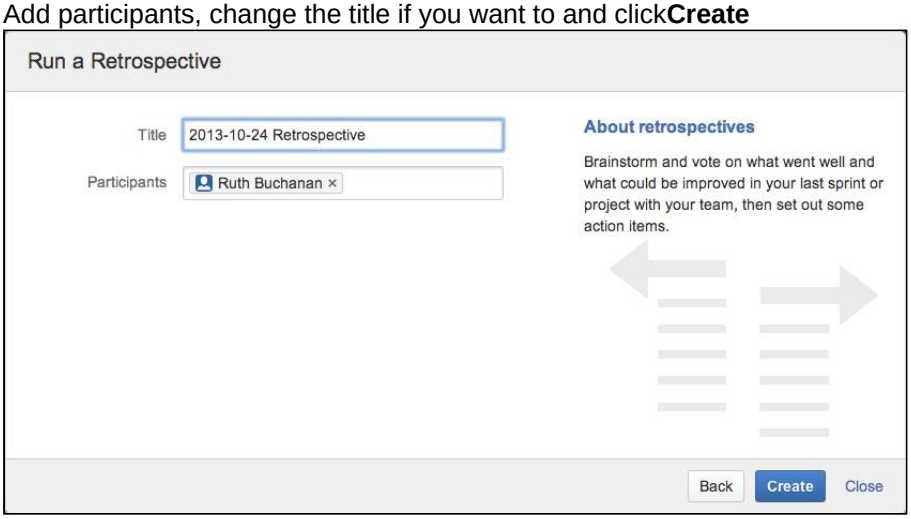

Screenshot: The 'retrospective' template page

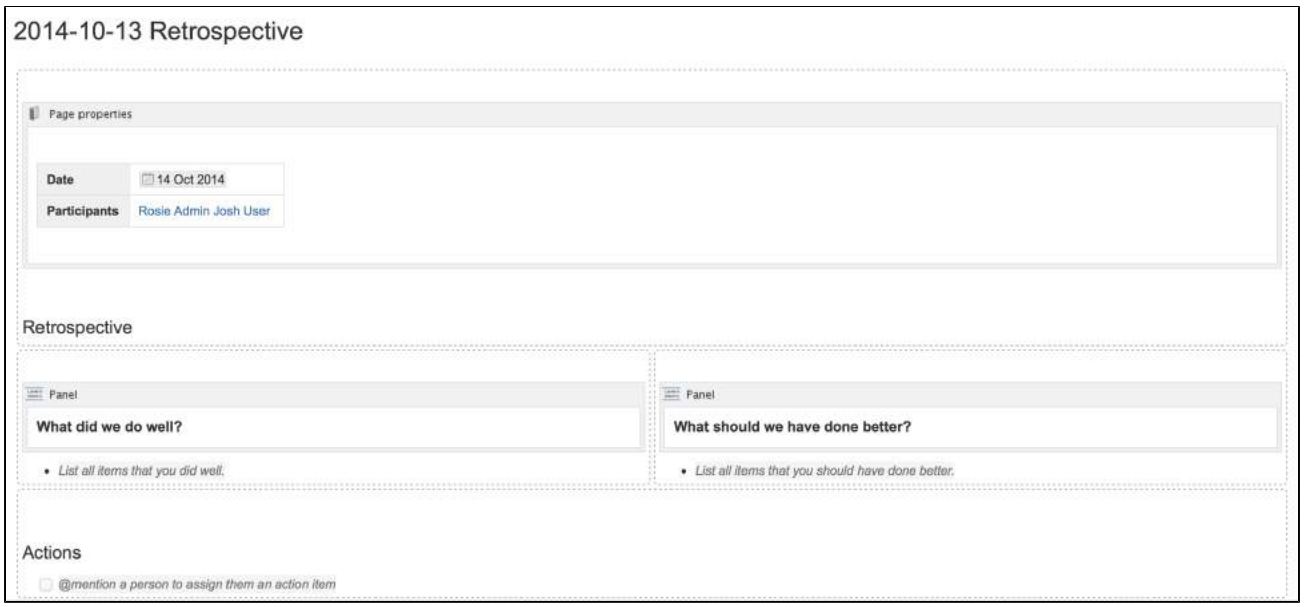

The Retrospective blueprint uses the following Confluence features:

- Page Propertiesand thePage Properties Reportmacro make content listed within the macro visible on the index page.
- Instructional textprompts you to enter information anddisappearswhen you start typing or view the page.
- Mentiona 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).

Customize this blueprint

Every team conducts retrospective meetings differently, so you can customize 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

SeeCustomize blueprint templatesfor instructions.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

How-To Article Blueprint

The How-To Article blueprint helps you to provide step-by-step guidance for completing a task.

To create a How-To Article page:

1. Choose **Create from template** ••• in the Confluence header
2. Select **How-To Article** and hit **Next**
3. Enter the article name and some labels and hit **Create**

Related pages:

- Blueprints
- Troubleshooting Article Blueprint
- Create Technical and Onboarding Documentation
- Use Confluence as a Knowledge Base

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 **Cre ate from template** button.

1. **Create blank page**
2. **Create from template**

Screenshot: A blank How-To Article page showing instructional text

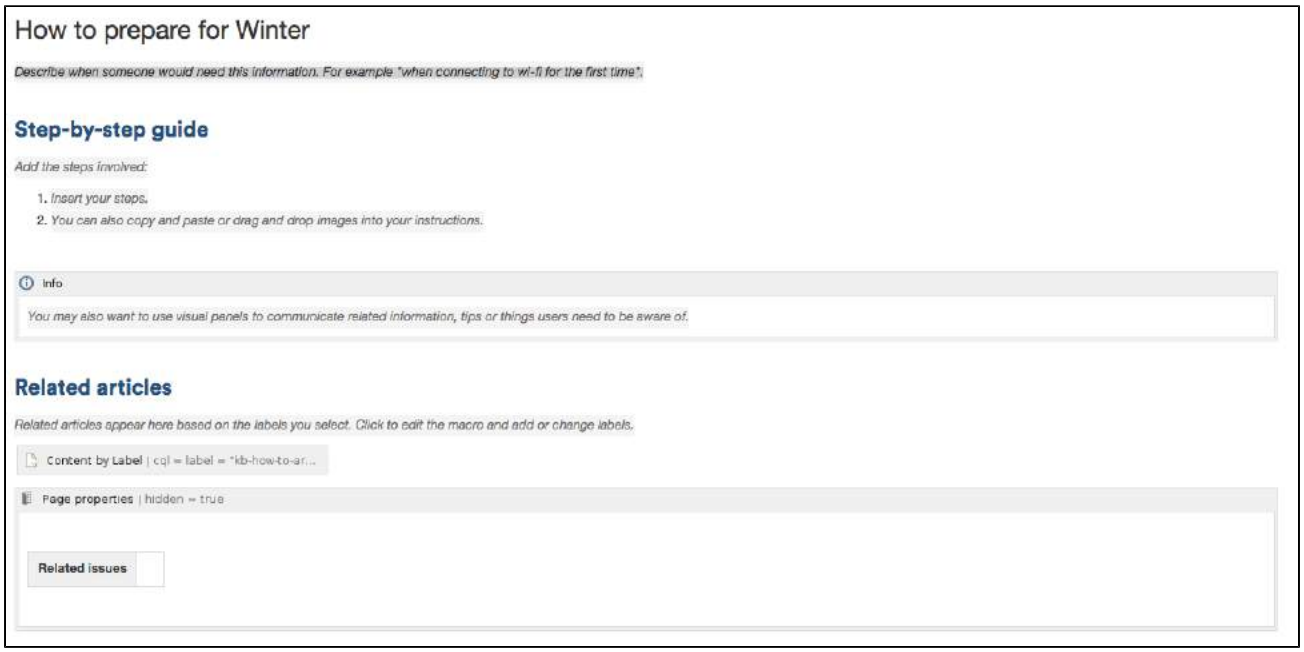

Once you save your page, Confluence will create an index page and add a shortcut to your space sidebar. The index lists all the How-To Article pages in that space, and displays a summary of the information on each page (such as creator and modified).

The How-To Article blueprint uses some cool Confluence features:

- Instructional text- Prompts you to enter information anddisappearswhen you start typing or view the page.
- Content by Label Macro -Displays lists of pages that have particular labels, to let you collect related pages together.
- Page Properties Macro - This works together with the Page Properties Report Macro to automatically create a list of 'related issues' on each article.

Customizingthis blueprint

You can customize the templates used by the How-To Article blueprint - seeCustomizing the blueprint templates. For example, you might choose to edit the How-to**indexpage**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 How-To Article was created.

SeeInstructional textto find out more about using instructional text in templates.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Troubleshooting Article Blueprint

The Troubleshooting Article blueprint helps you to provide solutions for commonly-encountered problems.

Related pages:

- Blueprints
- How-To Article Blueprint
- Create Technical and Onboarding Documentation
- Use Confluence as a Knowledge Base

To create a Troubleshooting Article page:

1. Choose**Create from template** in the Confluence header
2. Select**Troubleshooting Article**and hit**Next**
3. Enter the article name and some labels and hit**Create**

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.

1. **Create blank page**
2. **Create from template**

Screenshot: A blank Troubleshooting Article page showing instructional text

Once you save your page, Confluence will create an index page and add a shortcut on your space sidebar. The index lists all the Troubleshooting Article pages in the space, and displays a summary of the information on each page (such as creator and modified).

The Troubleshooting Article blueprint uses some cool Confluence features:

- Instructional text- Prompts you to enter information anddisappearswhen you start typing or view the page.
- Content by Label Macro -Displays lists of pages that have particular labels, to let you collect related pages together.
- Page Properties Macro - This works together with the Page Properties Report Macro to automatically create a list of 'related issues' on each article.

Customizingthis blueprint

You can customize the templates used by the Troubleshooting Article blueprint - seeCustomizing 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 Troubleshooting Article was created.

SeeInstructional textto find out more about using instructional text in templates.

You can also edit the Content Report Table macro used on the Index page to specify the number of pages you want to display.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

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 catalogs 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 alabelto the template (in this example, we'll add the label: 'status-update')

   Screenshot: Adding teams to our status update template

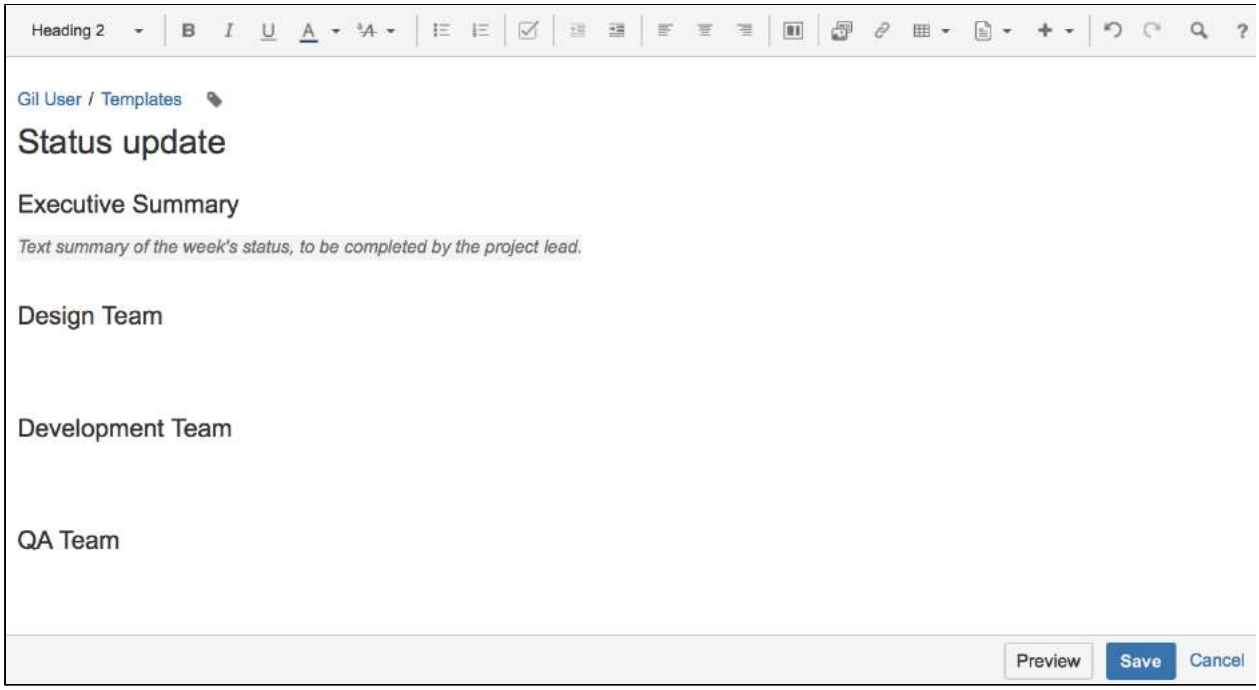

Now we'll add a Page Properties macro to record the status of the Design team.

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**Te mplate**>**Instructional Text**)
   We've also added a status macro.
10. **Edit**the Page Properties macro and enter a**Page PropertiesID**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 canenter a**Description**for your template - this appears in the Create dialog.

Screenshot: Our status update template

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

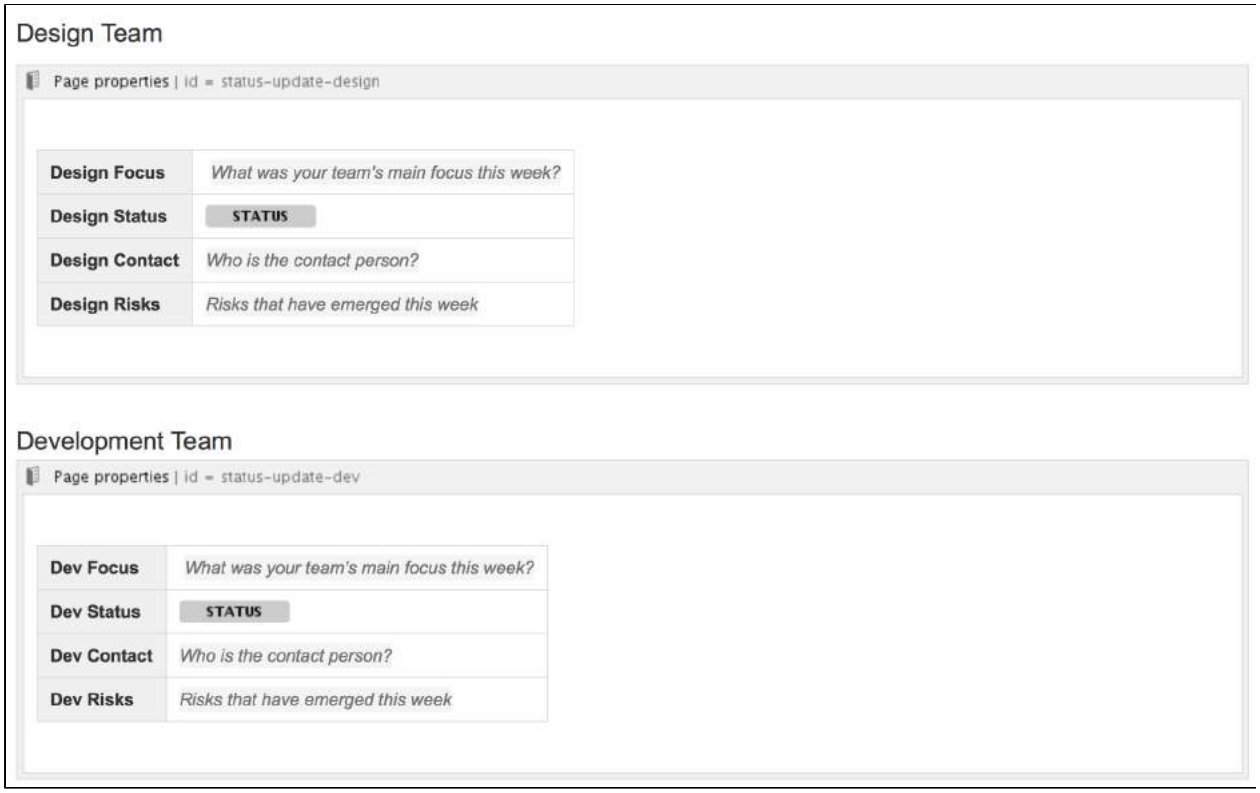

Part 2: Create a report showing the high level status of each team

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')

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

6. Choose **Save** to add the macro to the page

Screenshot:The page properties report macro on the'Status Report - all teams' page

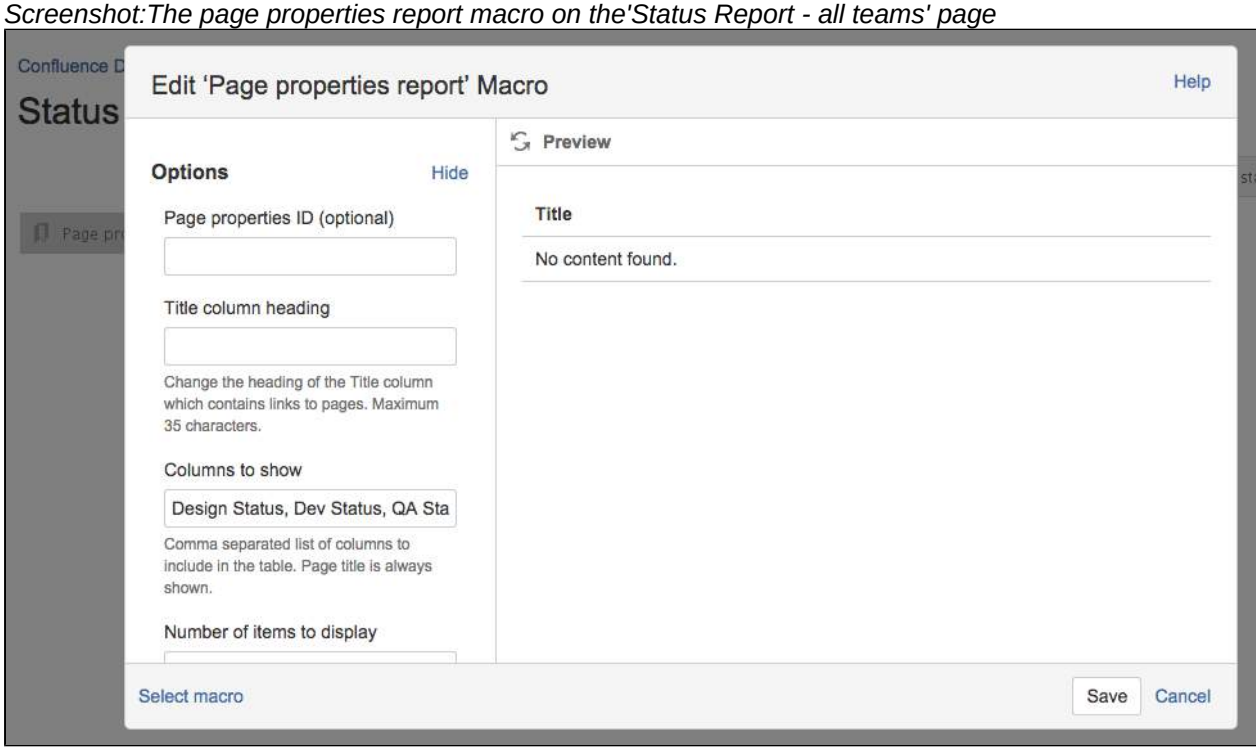

Now we'll add a button to the page to allow team leads to easily create new status update pages from the template we created earlier.

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)

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

6. Choose **Save**to add the macro to the page

Screenshot:The page properties report macro on the'Design Status Report' index page

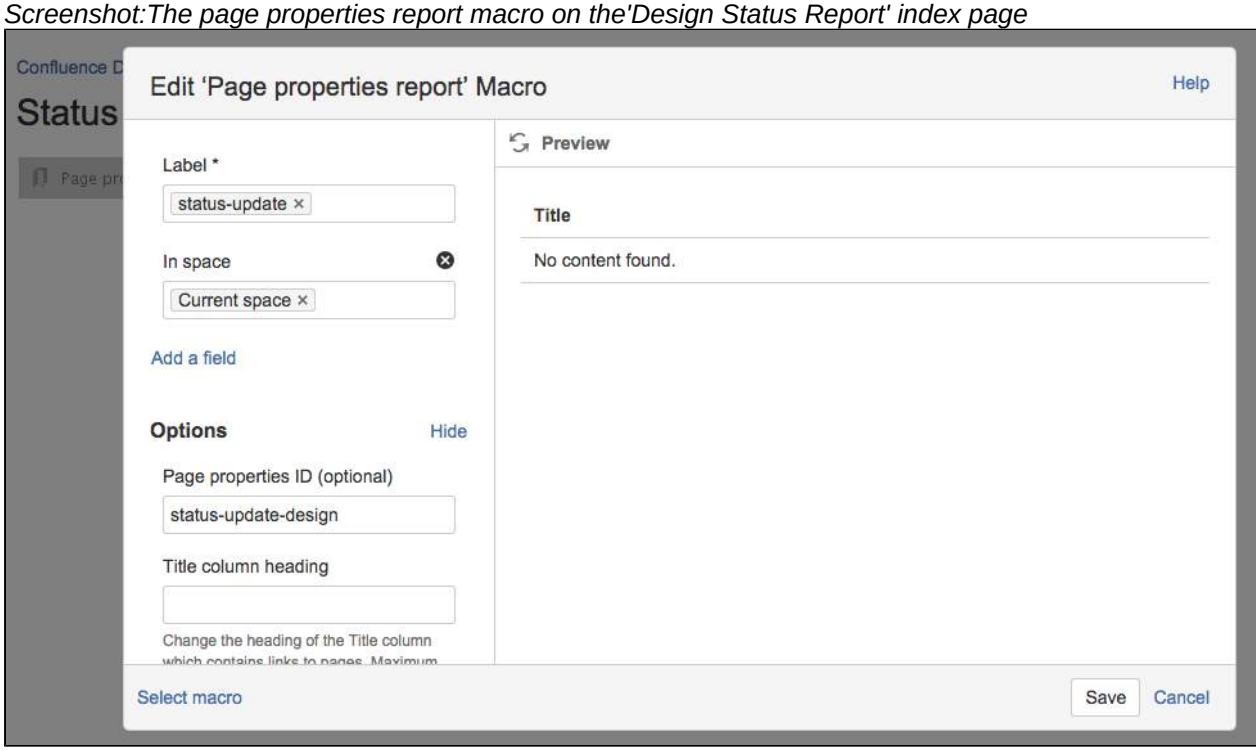

7. Add any other content, links or images to the page and**Save**
8. Choose**Space Tools**>**Configure Sidebar**>**Add Link**to add a shortcut to the page on the sidebar
9. Create a new page and repeat this process for each team
   Remember to specify a different Page Properties ID each time (in this 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

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

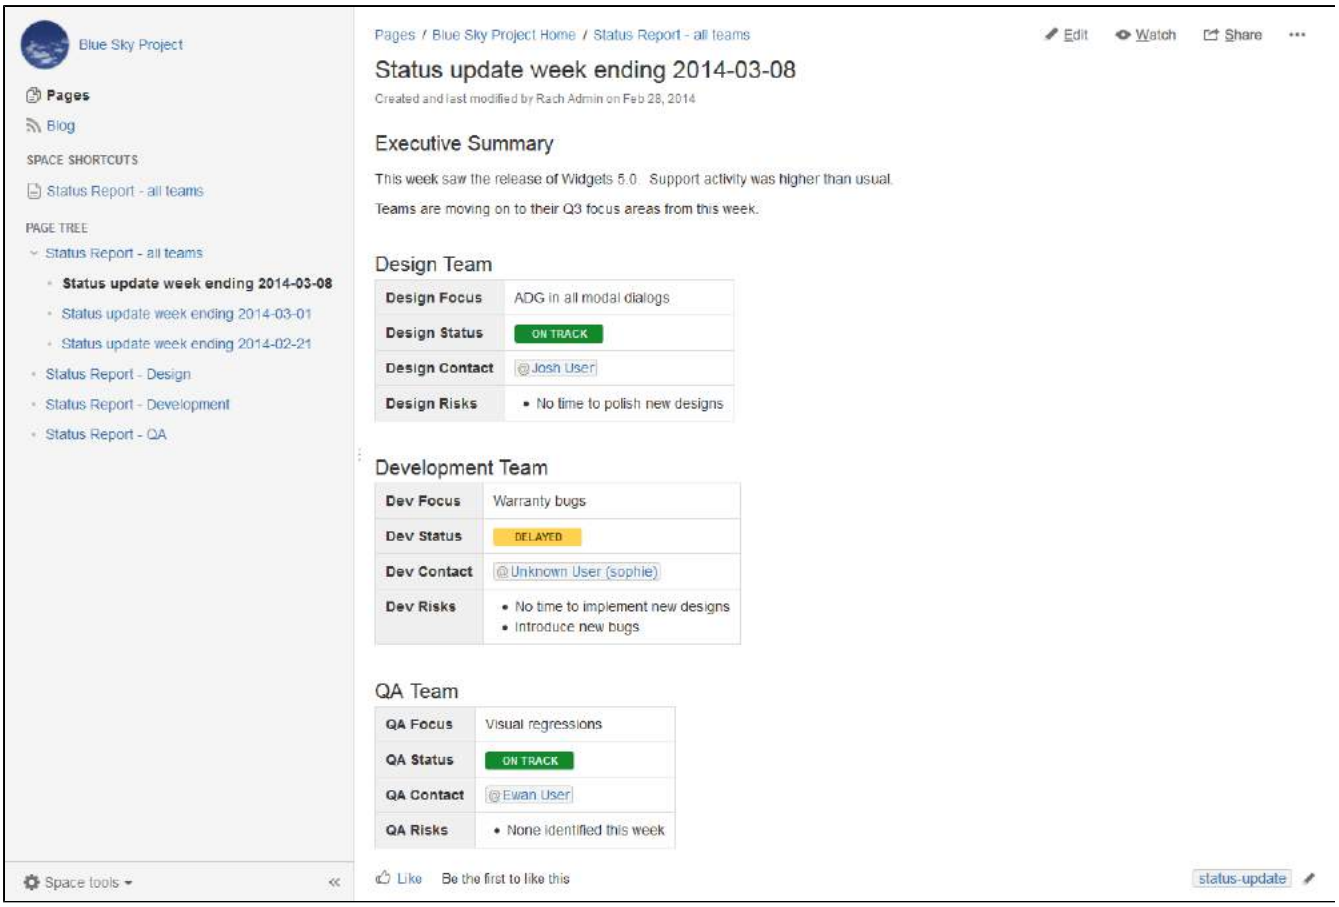

Screenshot:The Project Lead can see the status of each team, each week, at a glance in the All Teams status report.

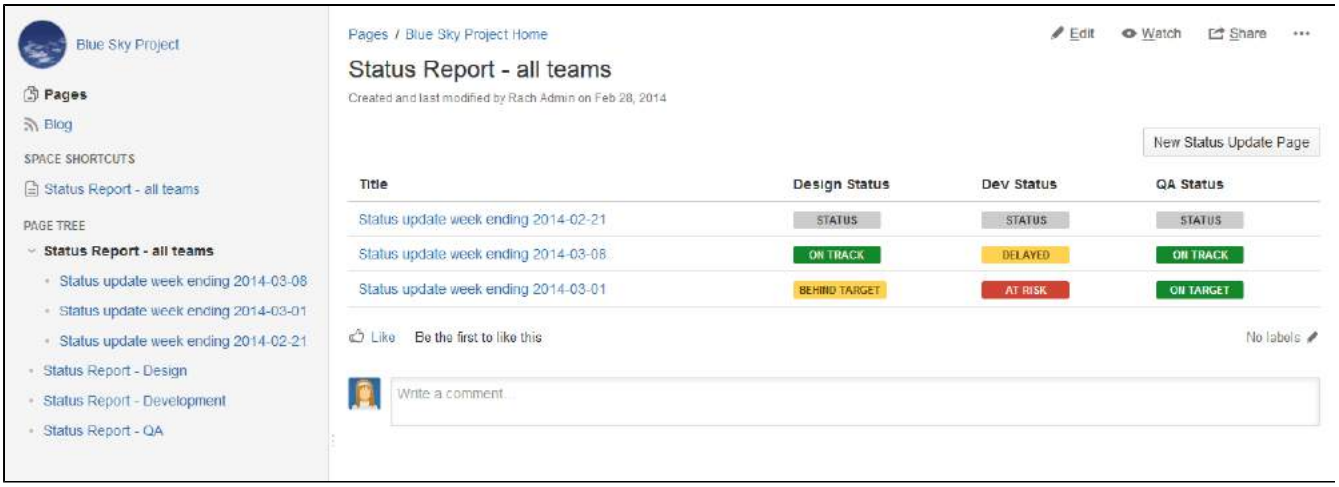

Screenshot:Each team can see their focus, risks and status at a glance in their status report.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

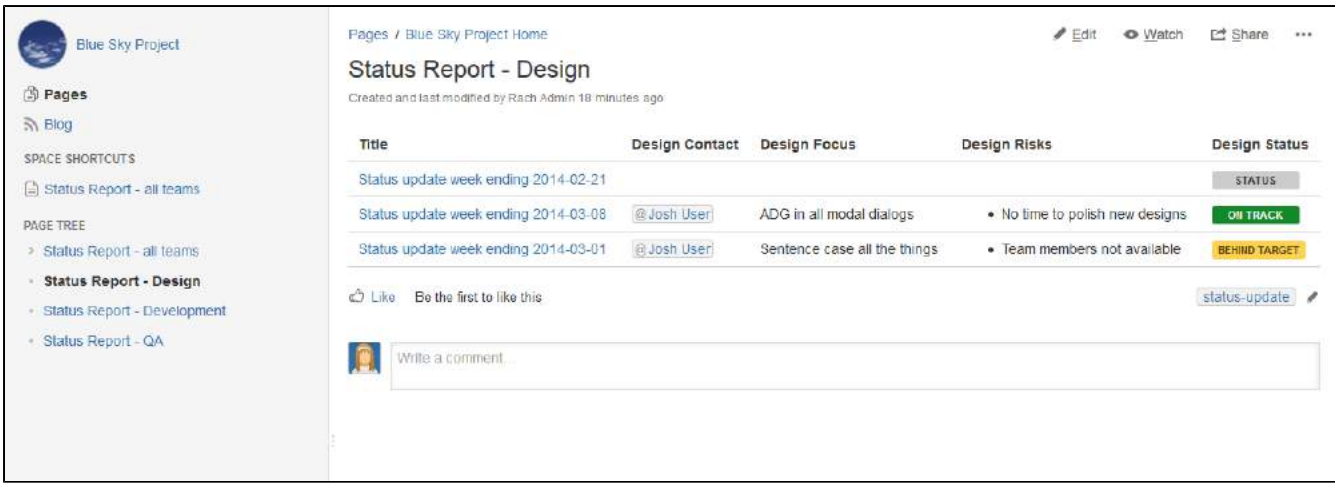

Remember, these concepts don't just apply to status updates you can use them for any purpose at all.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

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
- Importing content from another wiki
- Import other content
- Migrate to Confluence Cloud

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 Restore a Site and Restore a Space or multiple Spaces 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 embed 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 theHTML Include macro.
- Use HTML code in a page with theHTML macro.

Note: The HTML macro is not enabled in all sites. Talk to your Confluence Admin about whether you can use this macro.

Importing content from another wiki

Confluence does not provide a method for importing content from another wiki.

You may be able to build your own import solution using our REST APIs, as mentioned below, or work with anAtlassian Solution Partner to develop a custom solution.

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 beimported from disk.
- Files can be uploaded in bulk using theConfluence WebDav Plugin. See Use a WebDAV Client to Work with Pages.
- Build your own import solution using theConfluence APIs.

Migrate to Confluence Cloud

If you're migrating from Confluence Server to Confluence Cloud, you can use theConfluence Cloud Migration Assistantto migrate your content and spaces.

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.

On this page:

- Import a Word document
- Import options
- Supported file types
- Importing documents containing wiki markup
- Limitations

Related pages:

- Export Content to Word, PDF, HTML and XML

Import a Word document

To import a Word document in Confluence:

1. Create a page in Confluenceor 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 don't 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.

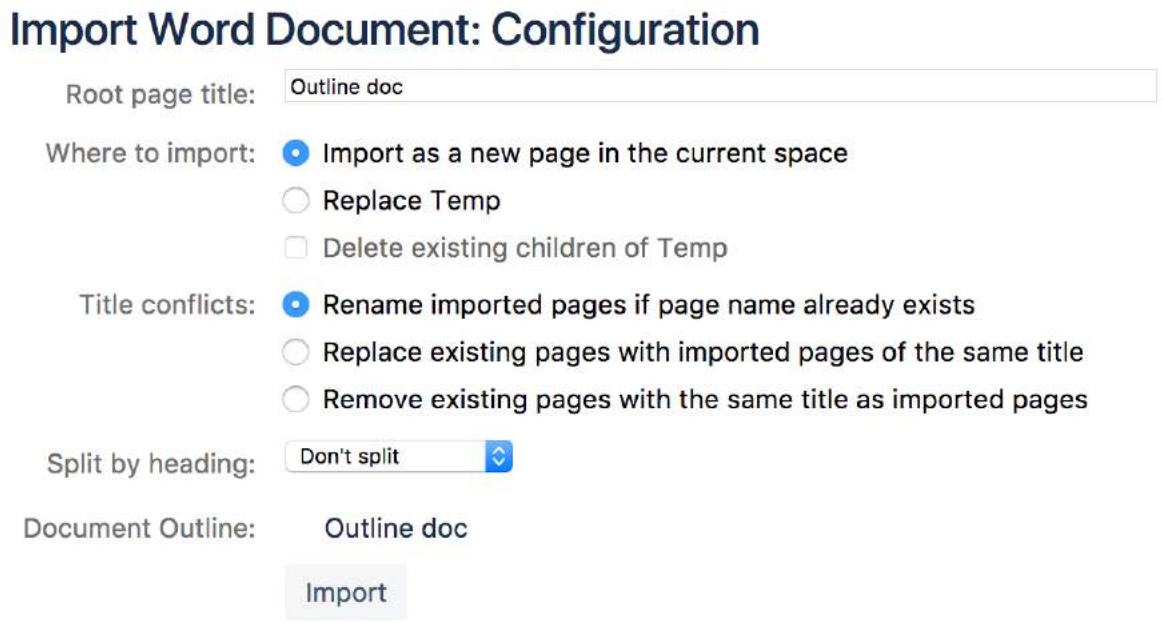

Supported file types

Confluence can import the content from Microsoft Word 97-2013 documents (.doc and .docx).

Importing documents containing wiki markup

If your document contains characters that are the same asConfluence wiki markup(for example square brackets, hashes etc) these characters may be automatically converted to their wiki markup value when you import your document, or wrapped in a wiki markup macro to preserve the formatting.

Limitations

In order to prevent out of memory errors, we limit the uncompressed size of the file you can import to 20 MB.

Your administrator can change this limit using the`confluence.word.import.maxsize`system property.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

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.

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**.

You can choose the link for an undefined page to create the page and add content to it.

Links to pages that are in the trash are not considered undefined links, and will not appear in this list.

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.
- **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

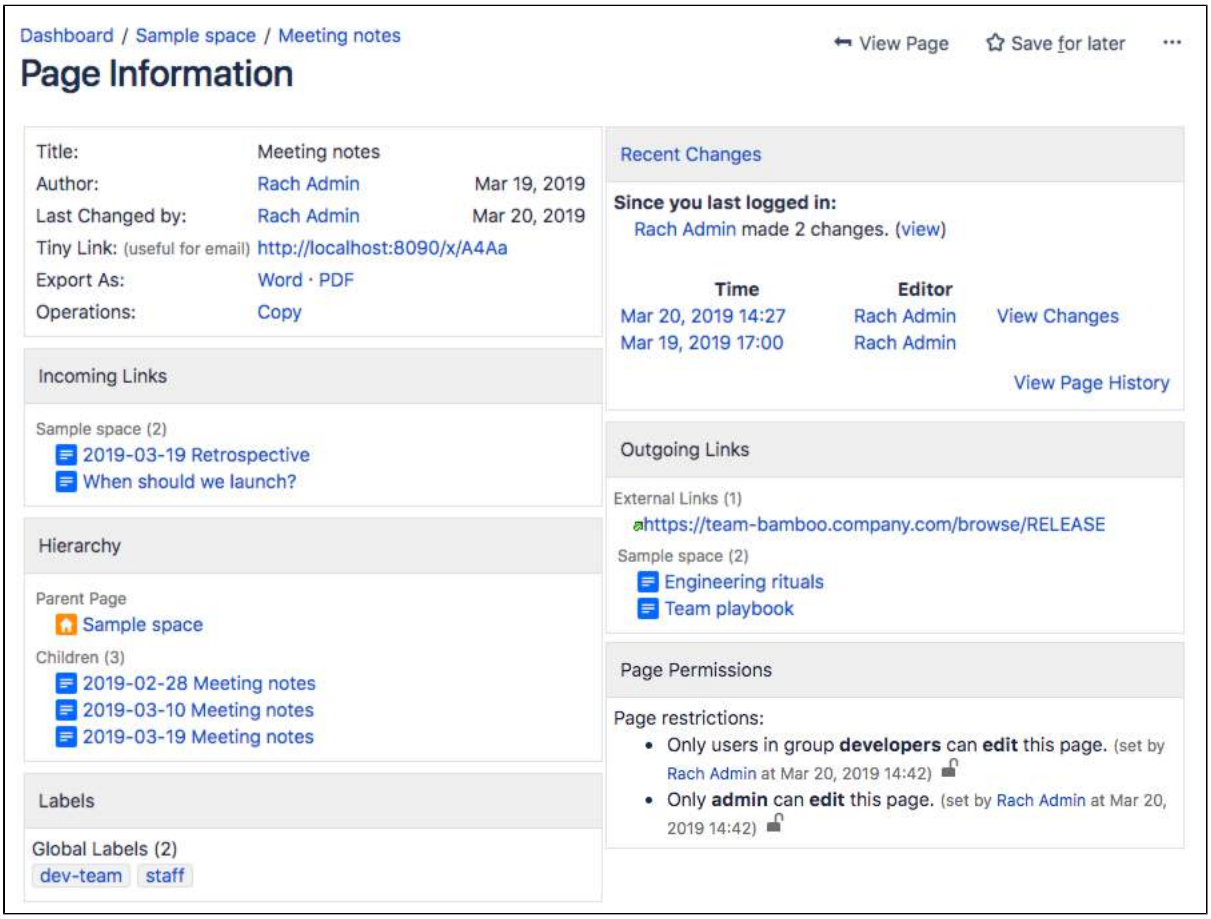

Page History and Page Comparison Views

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.

On this page:

- Access the page history
- View an older version
- Restore a previous version
- Delete a specific version
- View the changes made
- View unpublished changes
- Compare two versions
- More about the comparison view

Related pages:

- View Page Information
- Create and Edit Pages
- Watch Pages, Spaces and Blogs

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

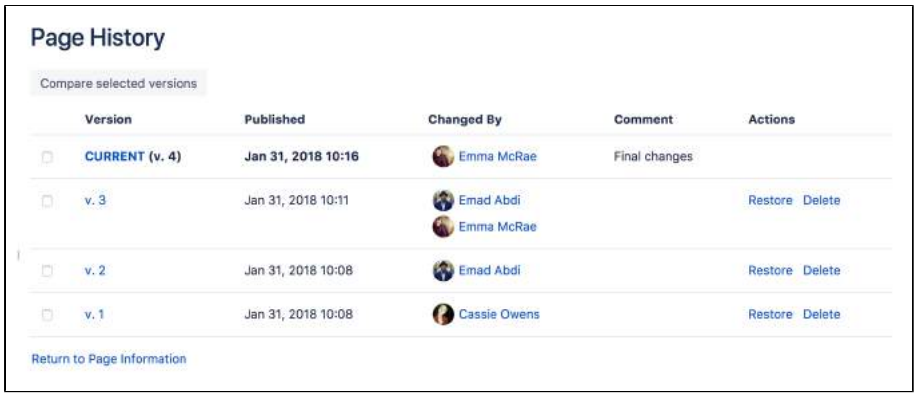

Hover over each avatar to see the names of people who contributed changes in that version. It is not possible to view the individual changes made by each person in a single page version.

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:

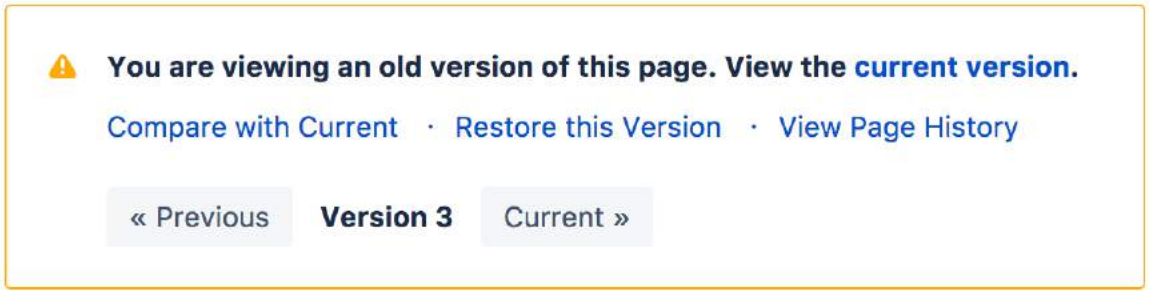

If you want to send this page version to someone, copy and paste the URL from your browser. The link will look something like this: `http://confluence.atlassian.com/pages/viewpage.action?pageId=12345`.

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**

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.

If the page has an unpublished draft, the content of the draft will be lost when you restore a previous version. We'll warn you if there is an unpublished draft.

Delete a specific version

Choose**Delete**next to a version in the page history, to remove that version.

Deleted versions do not go to the trash, and cannot be restored once deleted.

If you have **Confluence Data Center**, remaining versions will not be renumbered when a version is deleted. Older versions may also be deleted automatically, if your administrator has defined retention rules for the site or space.

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.

View unpublished changes

When you're in the editor, you can see all changes since the page was last published. Go to••• > **View changes**.

Changes are not attributed to individual people. The avatars of everyone who has contributed will be shown at the top of the editor.

Compare two versions

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

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 color	Meaning
Green	Added content
Red	Deleted content
Blue	Changed formatting

Screenshot: Comparing changes

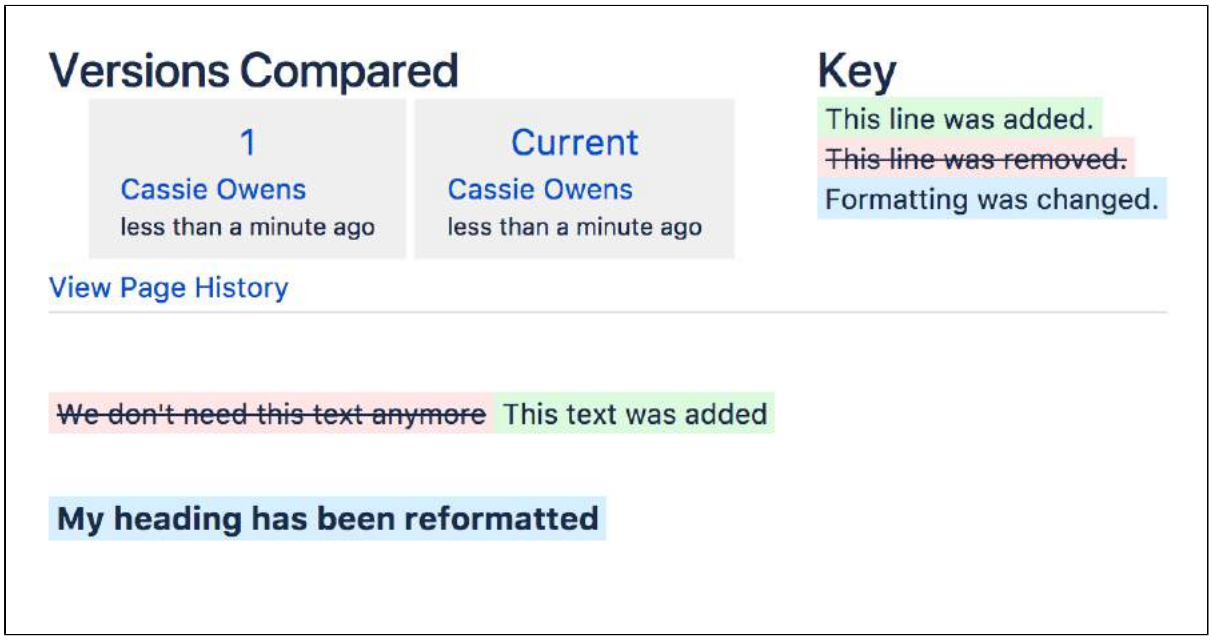

More about the comparison view

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 >>**

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Confluence Markup

This section describes two types of markup found in Confluence:

- **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.
- **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.

Wiki markup code examples for macros can be found in the documentation for each macro.

Related pages:

- Macros
- The Editor

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 theConfluence Source Editorplugin 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 theConfluence Source Editorplugin.
- 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:

- Headings
- Text effects
- Text breaks
- Lists
- Links
- Images
- Tables
- Page layouts
- Emojis
- Resource identifiers
- Template variables
- Instructional Text

Headings

Format type	In Confluence 4.0 and later	What you will get
Heading 1	`<h1>Heading 1</h1>`	**Heading 1**
Heading 2	`<h2>Heading 2</h2>`	**Heading 2**
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 4.0 and later	What you will get
strong/bold	`<strong>strong text</strong>`	**strong**
emphasis	`<em>Italics Text</em>`	*emphasis*

strikethrough	`<span style="text-decoration: line-through;"> strikethrough</span>`	~~strikethrough~~
underline	`<u>underline</u>`	underline
superscript	`<sup>superscript</sup>`	superscript
subscript	`<sub>subscript</sub>`	subscript
monospace	`<code>monospaced</code>`	`monospaced`
preformatted	`<pre>preformatted text</pre>`	`preformatted text`
block quotes	`<blockquote> <p>block quote</p> </blockquote>`	> *block quote*
text color	`<span style="color: rgb(255,0,0);">red text </span>`	red text
small	`<small>small text</small>`	small text
big	`<big>big text</big>`	big text
center-align	`<p style="text-align: center;">centered text </p>`	centered text
right-align	`<p style="text-align: right;">right aligned text </p>`	right aligned text

Text breaks

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Format type	In Confluence 4.0 and later	What you will get
New paragraph	`<p>Paragraph 1</p>` `<p>Paragraph 2</p>`	Paragraph 1 Paragraph 2
Line break	`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 4.0 and later	What you will get

Unordered list round bullets	`<ul>`
	`<li>round bullet list item</li>`
	`</ul>`
		• Round bullet list item

Ordered list (numbered list)	`<ol>`
	`<li>numbered list item</li>`
	`</ol>`
		1. Ordered list item

Task Lists	`<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>`
		☐ task list item

Links

Format type	In Confluence 4.0 and later	What you will get

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Link to another Confluence page	<ac:link> <ri:page ri:content-title="Page Title" /> <ac:plain-text-link-body> <![CDATA[Link to another Confluence Page]]> </ac:plain-text-link-body> </ac:link>	Link to another Confluence page
Link to an attachment	<ac:link> <ri:attachment ri:filename="atlassian_logo.gif" /> <ac:plain-text-link-body> <![CDATA[Link to a Confluence Attachment]]> </ac:plain-text-link-body> </ac:link>	Link to an attachment
Link to an external site	Atlassian	Atlassian
Anchor link (same page)	<ac:link ac:anchor="anchor"> <ac:plain-text-link-body> <![CDATA[Anchor Link]]> </ac:plain-text-link-body> </ac:link>	Anchor Link
Anchor link (another page)	<ac:link ac:anchor="anchor"> <ri:page ri:content-title="pagetitle"/> <ac:plain-text-link-body> <![CDATA[Anchor Link]]> </ac:plain-text-link-body> </ac:link>	Anchor Link
Link with an embedded image for the body	<ac:link ac:anchor="Anchor Link"> <ac:link-body> <ac:image><ri:url ri:value=" http://confluence.atlassian.com/ images/logo/confluence_48_trans.png" /></ac:image> </ac:link-body> </ac:link> ⚠ For rich content like *images*, you need to use `ac:link-body` to wrap the contents.	

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.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

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>, , , <i>, <code>, <tt>, <sub>, <sup>, <br> and .

Images

Format type	In Confluence 4.0 and later	What you will get
Attached image	`<ac:image> <ri:attachment ri:filename= "atlassian_logo.gif" /> </ac:image>`	
External image	`<ac:image> <ri:url ri:value="http://confluence.atlassian. com/ images/logo/confluence_48_trans.png" /></ac: image>`	

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

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

ac:hspace	the white space on the left and right of an image

Tables

Format type	In Confluence 4.0 and later	What you will get

Two column, two row (top header row)	`<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	`<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

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Confluence supportspage layoutsdirectly, as an alternative to macro-based layouts (using, for example, thesection and columnmacros). This section documents the storage format XML created when these layouts are used in a page.

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 recognized 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.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

```
<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>
```

Emojis

Format type	In Confluence 4.0 and later	What you will get
Emoticons	`<ac:emoticon ac:name="smile" />`	
	`<ac:emoticon ac:name="sad" />`	
	`<ac:emoticon ac:name="cheeky" />`	
	`<ac:emoticon ac:name="laugh" />`	
	`<ac:emoticon ac:name="wink" />`	
	`<ac:emoticon ac:name="thumbs-up" />`	
	`<ac:emoticon ac:name="thumbs-down" />`	
	`<ac:emoticon ac:name="information" />`	
	`<ac:emoticon ac:name="tick" />`	

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

	`<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.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

<table>
<tr>
<td>Attachment</td>
<td>
<pre>
<ri:attachment ri:filename>
        ... resource identifier for the container of the attachment ...
</ri:attachment>
</pre>
Notes:

<code>ri:filename</code>: (required) denotes the name of the attachment.
the body of the <code>ri:attachment</code> 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
<pre>
<ri:attachment ri:filename="happy.gif" />
</pre>
Absolute Attachment Reference
<pre>
<ri:attachment ri:filename="happy.gif">
        <ri:page ri:space-key="TST" ri:content-title="Test Page"/>
</ri:attachment>
</pre>
</td>
</tr>
<tr>
<td>URL</td>
<td>
<pre>
<ri:url ri:value="http://example.org/sample.gif"/>
</pre>
Notes:

<code>ri:value</code>: (required) denotes the actual URL value.

</td>
</tr>
<tr>
<td>Shortcut</td>
<td>
<pre>
<ri:shortcut ri:key="jira" ri:parameter="ABC-123">
</pre>
Notes:

<code>ri:key</code>: (required) represents the key of the Confluence shortcut.
<code>ri:parameter</code>: (required) represents the parameter to pass into the Confluence shortcut.
The example above is equivalent to <code>[ABC-123@jira]</code> in wiki markup.

</td>
</tr>
<tr>
<td>User</td>
<td>
<pre>
<ri:user ri:userkey="2c9680f7405147ee0140514c26120003"/>
</pre>
Notes:

<code>ri:userkey</code>: (required) denotes the unique identifier of the user.

</td>
</tr>
<tr>
<td>Space</td>
<td>
<pre>
<ri:space ri:space-key="TST"/>
</pre>
Notes:

<code>ri:space-key</code>: (required) denotes the key of the space.

</td>
</tr>
</table>

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

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:

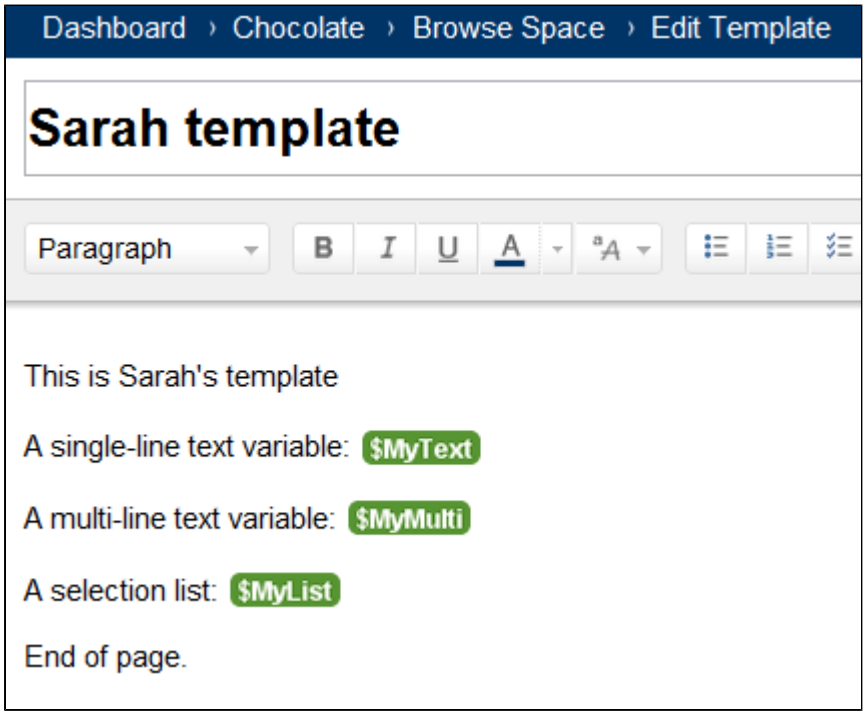

The template contains the following variables:

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:

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

```
<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.

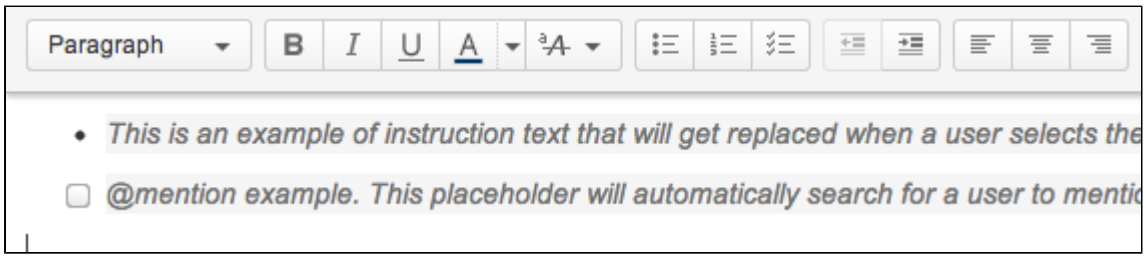


```
<ul>

<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>
```


Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

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:

- Type wiki markup directly into the editor. Confluence will convert it to the rich text editor format as you type.
- Createlinksusing the**Advanced**tab of the Links Browser.
- Add custom content to the **sidebar**, **header** or **footer** of a space.
- 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.

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

Can I type wiki markup into the editor?

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.)

Can I insert markdown?

Confluence supports inserting content in markdown. This is often used in ReadMe files. SeeMarkdown syntax guidefor 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'll 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.

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'll get

`- some`
`- bullet`
`- points`
	■ some
	■ bullet
	■ 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'll get

`* some`
`* bullet`
`** indented`
`** bullets`
`* points`
	• some
	• bullet
	  ○ indented
	  ○ bullets
	• 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'll get

`# a`
`# numbered`
`# list`
	1. a
	2. numbered
	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'll 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.
	  a. This is a sub-list point.
	  b. And a second sub-list point.
	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'll get

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

```
# 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.
    a. This is a sub-list point.
        i. Third list level.
        ii. Another point at the third level.
    b. And a second sub-list point.
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'll get

`# Here`
`#* is`
`#* an`
`# example`
`#* of`
`#* a`
`# mixed`
`# list`
	1. Here
	    • is
	    • an
	2. example
	    • of
	    • a
	3. mixed
	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. Once you've added this type of table, you can set the width of the columns using the table controls in the toolbar.

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'll 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:

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

```
||heading 1|col A1|col A2|col A3|
||heading 2|col B1|col B2|col B3|
```

What you'll get:

heading 1	col A1	col A2	col A3
heading 2	col B1	col B2	col B3

Table Type 2

This method uses section and column macros to create the table, and allows you to specify the width of the columns in the markup.

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'll 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

Color and Other Formatting

To add color 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'll get

Heading 1	**Heading 2**

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

• 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'll get
strong	**strong**
bold text	**bold text**
emphasis	*emphasis*
italics	*italics* *Hint:* To italicize 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+	inserted
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/m3
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.*

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

`{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 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.

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'll 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'll get

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

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'll 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.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

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'll 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 colors 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.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

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.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

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 canattach 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 Files
- Edit in Office using the Office Connector
- Install Atlassian Companion

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 aboutdisabling the download of attachmentsif 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
- What happens after a file is uploaded?
- 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:

- Dragthe file directly onto the page.
- Go to**Insert**>**Files and images**and upload a file.

When viewing a page you can:

- Dragthe file directly onto the page.
- Go to ••• > **Attachments**and upload a file.

You can attach multiple files at a time.

Copy and pasting a file from another application doesn't work reliably in many browsers. We recommend you use one of the methods above to make sure the file is uploaded successfully.

Accepted file types and size

Confluence allows you to attach most file types, butyou 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. SeeDisplay Files and Imagesto find out more.

The maximum file size you can upload to Confluence is set by your system administrator. By default it's 100mb, 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. SeeManage 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.

What happens after a file is uploaded?

Text extraction and indexing

When a file is uploaded, its text is extracted and indexed. This allows people to search for the content of a file, not just the filename.

SeeConfiguring Attachment Size for more information on how files are indexed.

Thumbnail and preview generation

When you insert an uploaded file into a page (for example a Word document, or Excel spreadsheet), Confluence will generate thumbnail images of the file contents, so it can be viewed inline in the page, or in the preview.

Because this process can be very memory intensive,a 30 second time limit applies when performing document conversion for complex image orpresentationfiles (such as PPT, PPTX, EMF, WMF). Your administrator can increase or decrease this timeout using the`confluence.document.conversion.imaging.convert.timeout`or `confluence.document.conversion.slides.convert.timeout`system properties.

Thumbnails are not generated for TIFF or PSD (Photoshop) files by default. Your administrator can override this behaviour using the`confluence.document.conversion.imaging.enabled.tif`or `confluence.document.conversion.imaging.enabled.psd`system properties.

In Confluence Server, document conversion is handled by Confluence. In Confluence Data Center this process is externalised, to minimise the impact on individual Confluence nodes. SeeDocument conversion for Confluence Data Centerto find out about how this affects thumbnail and preview generation.

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.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Display Files and Images

Files can be displayed on a page as a thumbnail or as a link. There are a few different ways toUpload Files.

You can control how the file appears on your page. The options available depend on the type of file.

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
- Edit Files

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
- Dragthe 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.

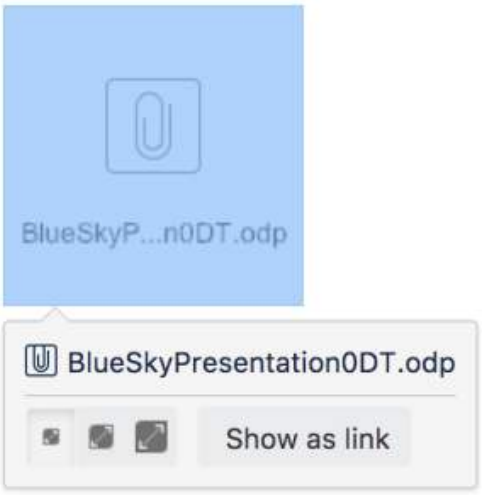

Insert a file attached to another page

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 attachmentcompletelyfromthe 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.
- Choose to edit the file with a desktop application.
- 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.

Images files	Office files	Other files
• JPEG • PNG • TIFF • PSD • WMF • EMF • ICO • ICNS	• DOC • DOCX • PPT • PPTX • XLS • XLSX	• PDF • MP3 • MP4

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

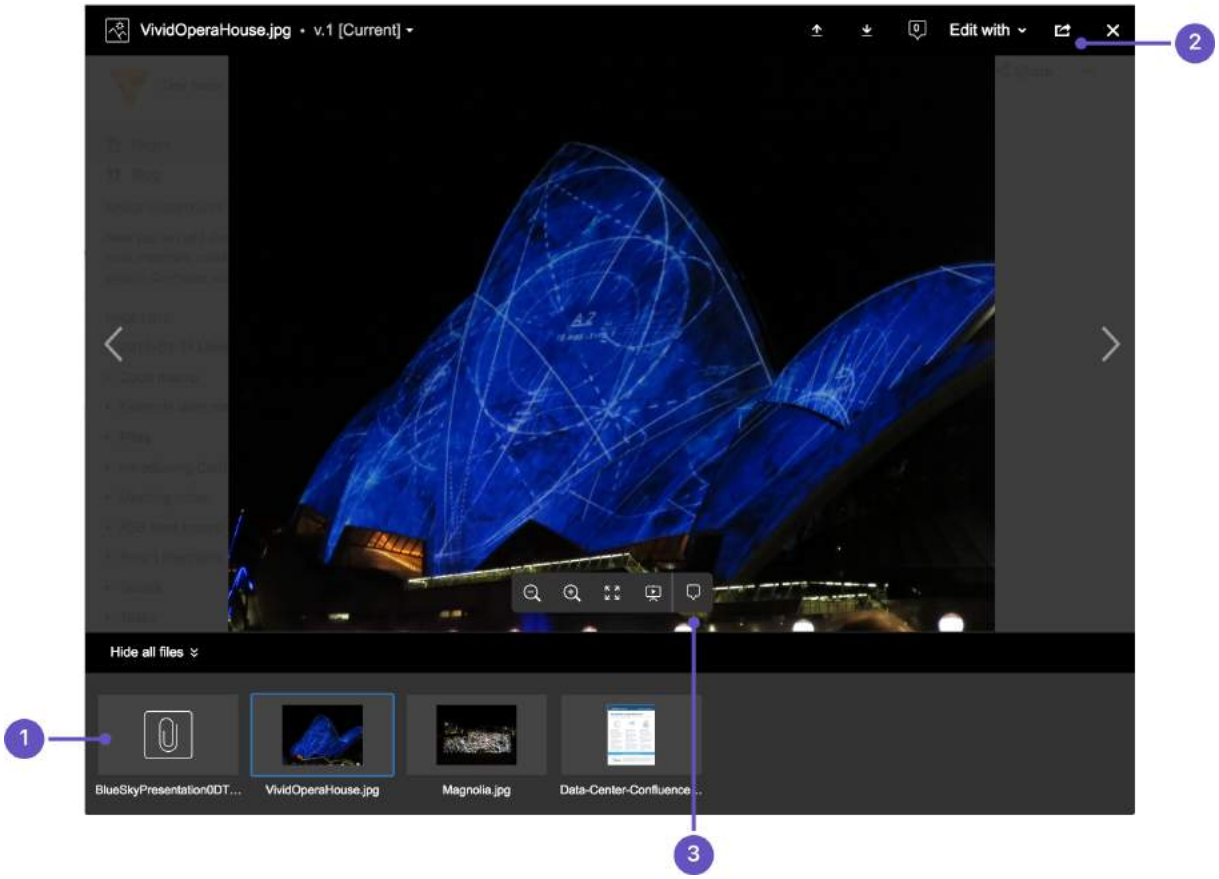

1. **See more files:**see other files also attached to this page.
2. **Manage this file:**download the file, upload a new version or share with your team.
3. **Add a comment:**drag the pin to comment on the file.

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 edit andcomment 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.

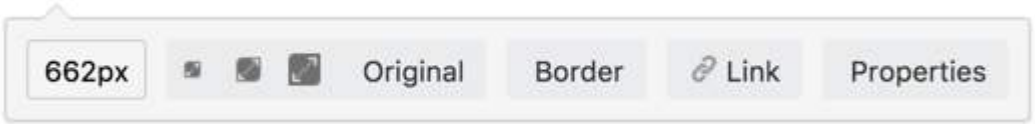

From the image properties panel you can:

- Choose a**preset size**for the image
- Enter a**width**for the image (between 16px and 900px)

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

- 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: TheInstant Cameraeffect 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 will prevent users from applying an image effect if the image is very large (based on data size and dimensions in pixels).

Confluence also 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. SeeImage effects are not displayed in Confluence 5.5 or laterif you need to adjust the number of threads.

Multimedia files

The file preview also supports MP3 audio and MP4 video files. It usesHTML5 to play attached audio and video files. This means the file types people can play in the preview depends on the audio and video formats their browser supports.

You can also display a wider range of multimedia files (video, audio and animation) using theMultimedia Macro.

Display online video (such as YouTube or Vimeo videos) using theWidget 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 theAttachments Macroto show files attached to the current page.
- Use theSpace Attachments Macroto show all files in a space.
- Use theGallery Macroto show thumbnails of images attached to a page.

You can also use theFile List blueprintfor uploading, viewing and managing lists of files.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Manage Files

Files are attached to Confluence pages.See Upload Filesto 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.

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

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 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

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 chooseto 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 exampleDOC: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

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

- 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 sidebarThen choose **Attachments**.

You can use the filters to onlyshow files with a particular label or file extension.

Screenshot: Space attachments macro

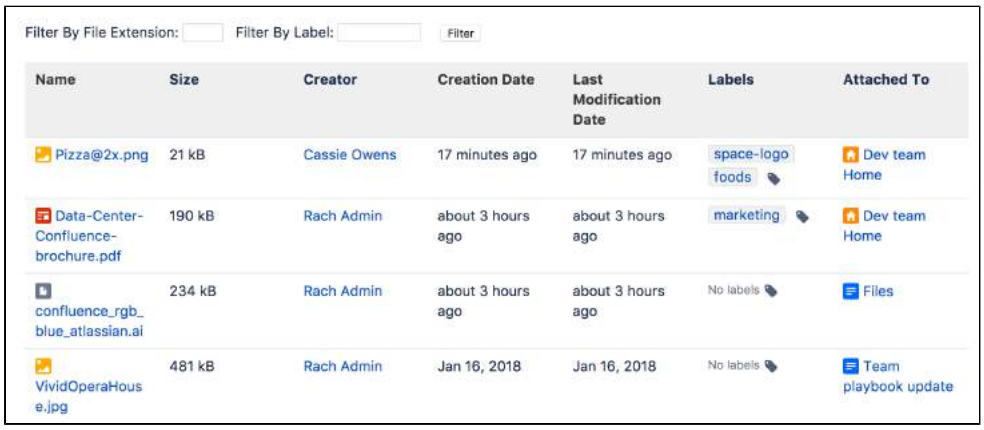

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

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.

On this page:

- Share a file
- Comment on a file

Related pages:

- Files
- Edit Files
- Display Files and Images
- Comment on pages and blog posts

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.

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 thepreview and drop it where you want to comment.
3. Add your comment and **Save**.

Pinned comments work just likeinline commentson pages. You can use @mentionsandlinks, and drop as many pins as you need on any part of the file. You can even add simplemacrossuch as the code macro using wiki markup autocomplete. 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 on that version of the file. 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.

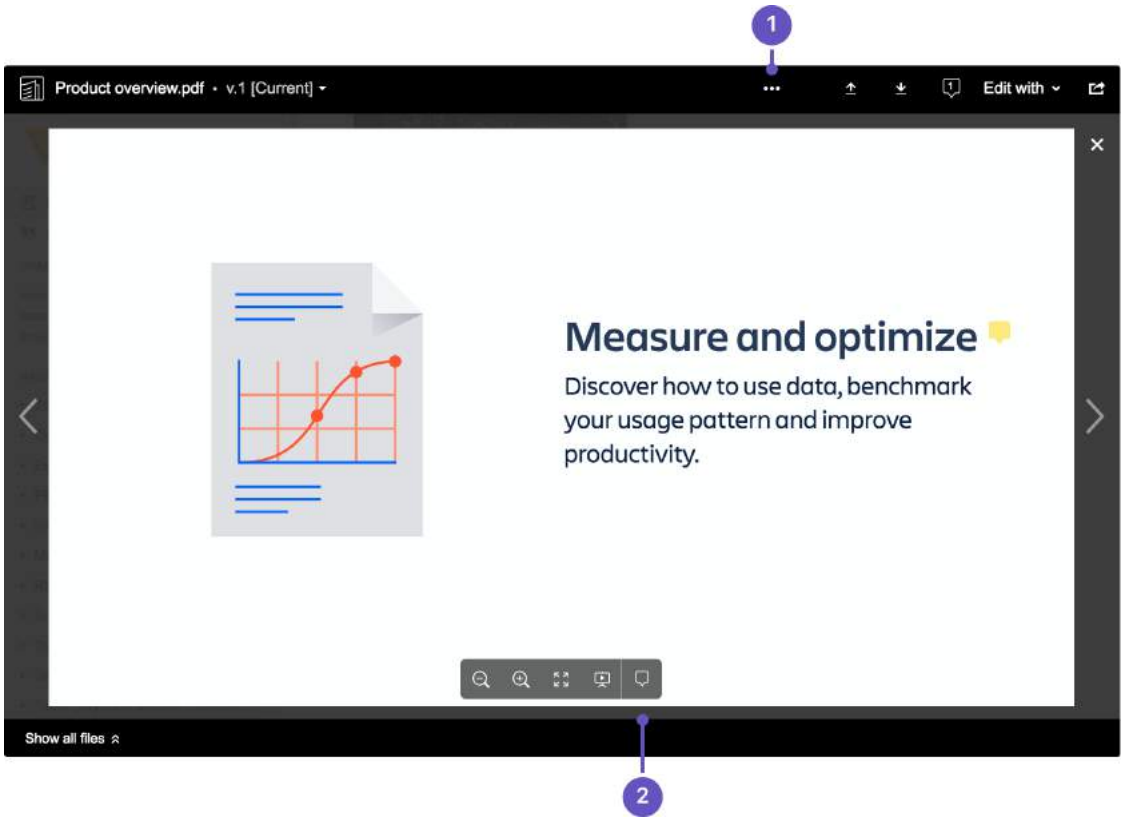

1. **Resolved comments:** Choose the 'more options' button to show or hide resolved comments.
2. **Comments:** Drag a pin onto a file to comment.

> 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).

What happens to comments when you upload a new version?

Comments are specific to the version of the file. This is to avoid confusion when the part of the document or image the comment is pinned to has changed significantly.

To see inline comments on a previous version of the file:

1. Click the thumbnail or link to preview the file.
2. Click the filename dropdown in the top left and select a previous version.
3. Comment pins will now be visible, for all comments made on that version.

How many comments can you add to one file?

While there is no limit to the number of comments that can be added to a file, Confluence can only display 100 comments. See CONFSERVER-43397 GATHERING INTEREST for more information.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Edit Files

This page covers how to edit files attached to a Confluence page.

On this page:

- Install and connect the Atlassian Companion app
- Change your default app
- Edit conflicts
- Manually upload an edited file
- Limitations
- Revoke trust between Confluence and the Companion app

Related pages:

- Administering the Atlassian Companion app

Confluence provides two methods for editing files:

- **Edit files using the Atlassian Companion app**
  Allows you to edit any type of file, if you have a compatible application installed. Requires you to install Companion on your computer.
- **Edit in Office using the Office Connector**
  Allows you to edit Microsoft Excel, PowerPoint, and Word files with a compatible browser and Microsoft Office application. Provided for organizations who can't use Companion.

Your Confluence administrator will decide which method is best for your organisation.

To check which method is available in your site, go to the file preview (click an image or file thumbnail). If you don't see the**Edit**button, your site is using theEdit in Office method, so the information on this page doesn't apply to you.

Download the latest Companion version

Download the Atlassian Companion app for Mac or Windows.

SeeInstall Atlassian Companionfor a step-by-step guide.

Edit an attached file

You can edit any file attached to a Confluence Server or Data Center page using your preferred desktop application, then when you're ready, upload the file back to Confluence. You can edit Office documents, Photoshop files, Keynote presentations any attached file with a compatible application installed on your computer.

To edit a file you'll need the 'Add Attachments' space permission.

To edit files, you'll also need to install the Atlassian Companion app and allow it to connect to your Confluence site. Once the Companion app is installed and running, you can start editing.

To edit a file in Confluence:

1. Go to the page containing the attached file.
2. Click the file to open it in preview.
3. Click **Edit**.
4. Your browser will prompt you open / launch the Companion app.
5. The Companion window will appear, and your file will open in the default application for that file type.
6. Make your changes and then save your file in the desktop application.
7. In the Companion window, click **Upload** to upload the edited file. It will be saved as a newfile version.

Screenshot: Editing an image file with Companion.

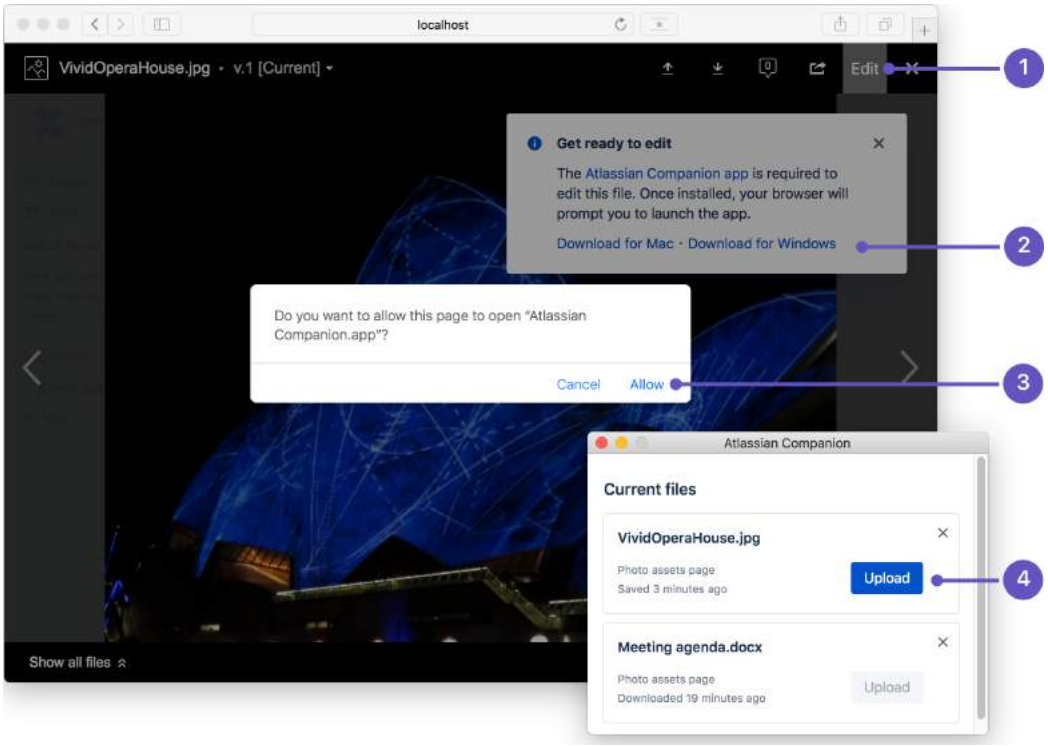

1. **Edit button** - if a file can be edited, you'll see an **Edit** button in the preview.
2. **Download Companion** - if you haven't already installed Companion, you can download it from this flag. The flag always appears, but you only need to download Companion once.
3. **Launch Companion**- your browser will ask for your permission to open Companion. The prompt is different in each browser, and will appear each time you edit a file.
4. **Upload or discard file** -once you have saved the changes to your file, click the Upload button toupload your edited file back to Confluence, or click the x icon to discard the changes, and remove the file information from the list.

> There is a known issue with editing attached files in some browsers in Confluence 6.11 to 7.2.SeeCa nt edit files in Confluence Server using Atlassian Companion app in Internet Explorer, Edge, Firefox, or Safari for more information and some workarounds.

Other ways to edit files

You can also edit a file from the attachments macro, view file macros(Word, Excel, PowerPoint, and PDF), or the attachments page.

To edit a file from the attachments page:

1. Go to the page that contains the attached file.
2. Go to ••• > **Attachments**
3. Next to the file name, click **Edit** , then follow the instructions above to edit and save your changes.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Install and connect the Atlassian Companion app

The Atlassian Companion app manages the download and re-upload of files you edit in Confluence.

The first time you edit a file, you'll be prompted to install the Companion app and allow it to connect to your Confluence site.

1. Click the file to open it in preview.
2. Click **Edit**.
3. Follow the prompts to download and install the Atlassian Companion app.
4. Launch the Companion app.
5. In Companion, click **Trust** to confirm you want to connect to Confluence.
6. Return to Confluence, and follow the steps above to edit the file.

You can also download and install the app manually for Mac or Windowsor use a Microsoft Installer (.msi file). See Administering the Atlassian Companion app for details.

For detailed installation instructions for your operating system, seeInstall Atlassian Companion.

Change your default app

Confluence allows you to edit files in your operating system's default app for that file type (for example, .psd files will open in Photoshop). To change the app your Confluence file opens in, change the default app in your operating system.

Edit conflicts

When you edit a file, Companion downloads a copy of the file to your computer. Once you're finished editing and you attempt to upload the file, Companion will warn you if a newer version of the file was uploaded after you downloaded the file. You can choose to click the **x** icon to discard your changes, or continue to upload your changes as a new version. Both versions will be available in the file history.

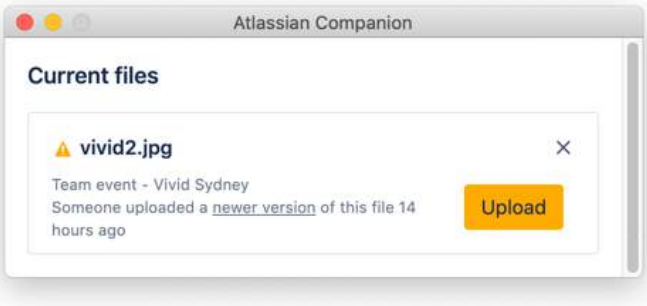

Confluence doesn't provide an option to check-out a file before you start editing. There are apps available on the Atlassian Marketplace that can provide this functionality.

Manually upload an edited file

You can only upload changes back to Confluence if those changes are saved in the original file. You **won't** be able to upload if:

- you edit the file and save it as a new version (save as)
- the application youre using saves the file in a different format to the original for example, from a PowerPoint file (.pptx) to a Keynote file (.key).

If this happens, you can upload your new version manually:

1. Click the original file in Confluence to open it in preview.
2. Click the**Upload a new version** button and select your new file version.
3. Click**Done**.

What to do if you lose your edited file

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Any files you've opened with Companion will remain listed in the Companion window. The status information below the file will indicate the time since the file was downloaded, edited, or uploaded.Files that haven't been edited in the last 60 daysare cleared when Companion app restarts.

You can also follow our guide to accessing Confluence files edited with the Atlassian Companion appif the file isn't listed or you're not able to re-open it in your desktop application.

Limitations

File size limits

Confluence won't allow you to upload your changes if the edited file is larger than your site's maximum file size limit. This limit is set by your system administrator. By default the limit is 10 MB, but your admin may have increased or reduced it. Check outUpload Filesfor more information.

Blocked file types

We block file types that might pose a risk to your security, including executable files such as .exe or .bat files.

It's possible to modify the file types companion can open using an environment variable. SeeHow to change the file types blocked by Companion in Windows.

Cross file links and references

If your file links or references to other files (for example if you link a worksheet in one Excel file, to another Excel file)these links will not work once the files have been uploaded to Confluence.

Revoke trust between Confluence and the Companion app

If you want to disconnect the Companion app from your Confluence site, you can remove it as a trusted site.

To revoke trust:

1. Click the Companion app icon in your system's toolbar.
2. Choose**Clear all trusted domains**.

Note: clearing trusted domains won't kill active connections. If you select**Clear all trusted domains**while editing a file, you'll still be able to upload those changes back to Confluence.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Edit in Office using the Office Connector

This page covers how to edit files using the legacy **E dit in Office** method.

- Editing Office files attached to a page
- Other ways to edit
- Office Connector pre-requisites
- Enable Edit in Office
- Troubleshooting

Confluence provides two methods for editing files:

- **Edit files using the Atlassian Companion app**
  Allows you to edit any type of file, if you have a compatible application installed. Requires you to install Companion on your computer.
- **Edit in Office using the Office Connector**
  Allows you to edit Microsoft Excel, PowerPoint, and Word files with a compatible browser and Microsoft Office application. Provided for organizations who can't use Companion.

Your Confluence administrator will decide which method is best for your organisation.

To check which method is available in your site, go to the file preview (click an image or file thumbnail). If you see the **Edit**button, your site is using the Companion app method, so the information on this page doesn't apply to you.

Editing Office files attached to a page

The Office Connector allows you to edit Office files that are attached to pages, if your site does not use the Companion App method. You'll need to use a browser, operating system and application (eitherMicrosoft OfficeorOpenOffice) as described in the compatibility matrix below.

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.

⚠Edit in Office will not work on files that have special characters (such as' #@ or) in the filename.

Screenshot: Edit in Office option on the attachments page

Other ways to edit

Editoptions also appear in the:

- Attachments macro(choose**Edit in Office**beside each attached office file)
- Office Word and Office Excel macros choose the**Edit**button above the content.
- Office PowerPoint macro choose the edit icon on the viewer.

Office Connector pre-requisites

You can only edit files using the Office Connector if your system administrator has enabled this feature on your site. By default, Confluence requires you to install the Atlassian Companion app to edit attached files. See Edit Files

Configuration matrix

Edit in Office is only compatiblewith desktop applications. Online versions of Office applications are not supported.

You need one of the following software combinations to edit Office files from your Confluence page.

Software	Operating system	Browser
Microsoft Office 2010 SP2, 2013, 2016, 2019	• Windows	• Chrome • Firefox[1] • Internet Explorer 11 • Edge
Microsoft Office 2011, 2016, 2019	• Mac OS	• Chrome • Firefox • Safari
Microsoft Office XP, 2003, 2007, 2010 SP1	• Windows	• Internet Explorer 11
OpenOffice 2.x 3.x	• Windows • Linux	• Chrome • Firefox • Safari

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

LibreOffice 6.x	• MacOS • Windows • Linux	• Chrome • Firefox • Safari

[1] Firefox only works in Windows with Confluence 7.1 and later.

Note

The known supported Office editors for Linux are OpenOffice & LibreOffice, 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) contact our support team, who can raise an issue about it. Tell us as much as you can about your operating system, application version, document version (if it's different to the version of Office / Open Office you're using to open the document) and browser.

Here are a few common issues:

- **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.
- **Special characters in the filename?** Edit in Office does not work for files with special characters (like ' # @) in the filename. See
  CONFSERVER-22403 - Attaching office documents with special characters stops the ability to edit from office GATHERING IMPACT
- **Not seeing the Office Connector options?** Your system administrator needs to enable this feature, and can control how it appears on your site. See Enable Edit in Office as a dark feature and Configuring the Office Connector.

Enable Edit in Office

You need System Administrator global permission to do this.

To enable the legacy Edit in Office functionality:

1. Go to > **General Configuration**>**Office Connector**.
2. Choose**Enable Edit in Office for all users**and save your changes.

This will disable Companion app functionality for all users in the site.

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 managing system and marketplace apps 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.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Install Atlassian Companion

This page applies to Confluence Server and Data Center

The Companion app was removed from Confluence Cloud at the end of March 2022. Read announcement

We have no plans to end support for the Companion app in Confluence Data Center and Server.

If your site uses Atlassian Companion method to edit files, you will need to install the Atlassian Companion app before you can start editing. Learn more aboutEditing Files.

The way you install Companion will depend on your operating system, and your organisation's environment.

Download the latest Companion version

Download the Atlassian Companion app for Mac or Windows.

Windows

These instructions are for Windows 10. The process is similar in earlier, supported Windows versions.

1. Download the Companion app fromWindows download
2. Double click to run the Atlassian Companion-1.X.X Setup.exe file you just downloaded.
3. That's it. You may briefly see a progress indicator like the one below.

Once installed, Companion runs in the background. You may need to click the**Show hidden items**arrow in the system tray to see it. Right click the Companion icon to see the version you have installed.

The first time you edit a file, you'll be prompted to trust your Confluence site. Once that's done, Companion is ready to use.

Mac

These instructions are for MacOS 10. Your version may look slightly different.

1. Download the Companion app fromMac download
2. Double click the Atlassian Companion.dmgfile you just downloaded.
3. The install dialog will appear. Drag the Atlassian Companion icon to the Applications folder icon.

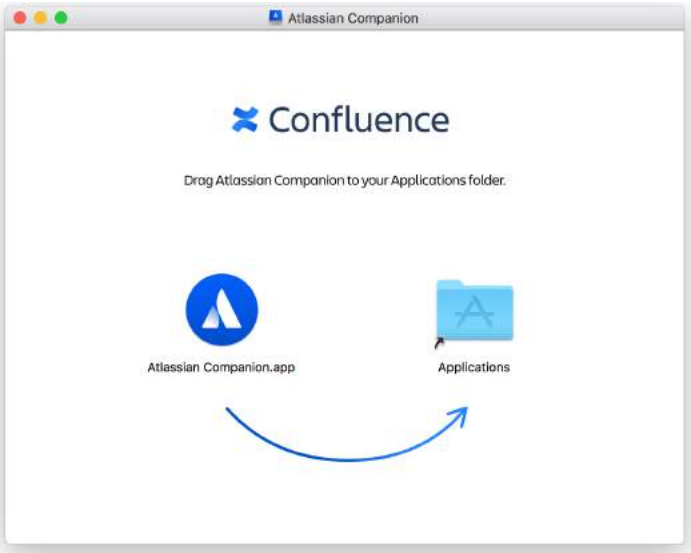

4. Launch the companion app:
    a. In Finder, go to your Applications folder, and click Atlassian Companion.app, or
    b. In Launchpad, select Atlassian Companion, or
    c. In Spotlight, search for Atlassian Companion.

Once launched, Companion runs in the background. Click the Companion icon in the status area of the menu bar to see the version you have installed.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

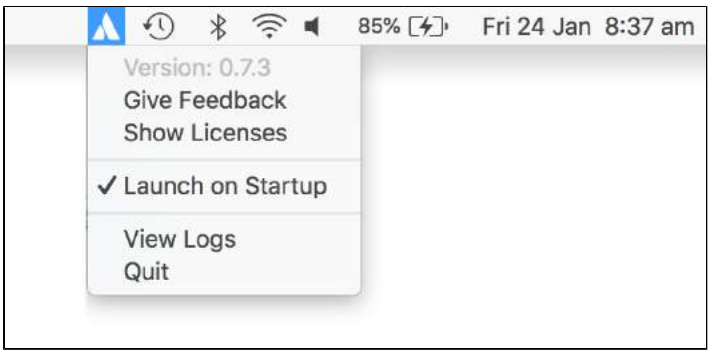

The first time you edit a file, you'll be prompted to trust your Confluence site. Once that's done, Companion is ready to use.

Linux

Atlassian Companion is not currently available for Linux.

Managed or restricted environments

If you're not able to install applications on your computer, your administrator or IT team may need to do this for you.

Send them a link toAdministering the Atlassian Companion App for information about MSI installation.

Troubleshooting

Companion is installed twice in Windows

It's possible to have two versions of Companion installed simultaneously if you have installed Companion using the installer (.exe) and the MSI (.msi). If this happens, you should uninstall both versions from Add/Remove programs and then re-install Companion (either using the `.exe`or the `.msi`depending on how you want to be able to update Companion in future).

Remove residual files after uninstalling Companion in Windows

If you need to completley uninstall Companion you may want to check if any residual files remain.

The default installation directories are:

- EXE: %LOCALAPPDATA%/atlassian-desktop-companion
- MSI: Program Files (x86)/Atlassian Companion

The app stores temporary data and edited files in:

- %APPDATA%/Atlassian Companion
- C:\Users\<UserName>\.atlassian-companion

Companion creates the following registry entries:

- Computer\HKEY_CLASSES_ROOT\atlassian-companion
- Computer\HKEY_CURRENT_USER\Software\Classes\atlassian-companion
- Computer\HKEY_CURRENT_USER\Software\Microsoft\atlassian-desktop-companion

You can do a registry search for "atlassian" to locate them. These are used for both the MSI and EXE.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Confluence Mobile

There are two ways to stay connected to your team's work in Confluence while you're out and about - theConfluence Data Center and Server mobile app for iOS and Android, or via your device's browser.

If you're a Confluence user, check out the two ways you can use Confluence on your device.

- Using the Confluence Data Center and Server mobile app
- Using Confluence via your mobile browser

On this page:

- Considerations for administrators
- Limitations and known issues

What you'll need

Confluence requirements

In order for your users to try the Confluence Data Center and Servermobile app, you will need to:

- upgrade your Confluence siteto**Confluence 6.8** or later
- allow users to access your site on their device (if your site isn't accessible on the public internet, people will need be connected to your network or use aVPN)

Device requirements

In order to use the app, your users will need a device with either:

- Android4.4 (KitKat) or later, or
- iOS 15 or later (iPhone, iPad or iPod Touch)

Users will need to log in to use the app, even if your site allows anonymous users.

Considerations for administrators

Here are some things to consider when determining whether your users will be able to use the app.

VPN and firewalls

If your Confluence site is not accessible on the public internet, users will need to connect their device to your network or virtual private network (VPN) in order to use the app.

We recommend providing your users with step-by-step instructions on how to connect to your VPN when you let them know the mobile app is available, as this is something Atlassian Support will not be able to help them with.

The mobile app will also attempt to check the compatibility of your siteprior to presenting the login screen.If you've configured a custom filter to prevent unauthenticatedrequeststo your server,you will need to change it to allow`<confluence-base-url>/rest/nativemobile/1.0/info/login` to pass through without authentication.

HTTPS and certificate requirements

In the latest version of the iOS and Android apps, you can connect to the app using either HTTP or HTTPS.

If you're using HTTPS your proxy must allow TLS 1.2 traffic.This is an iOS requirementthat we'vechosen to implement for both the iOS and Android apps to prevent confusion (for example where one device can log in, and another cannot).

Ideally, your certificate should be from a trusted Certificate Authority.If you have certificate that is self-signed, or from an unknown Certificate Authority (for example, you are your own CA), users may still be able to use the app by manually installing your certificate on their device. See ourKnowledge base articlefor more information on how to do this.

iOS 13 introduced a number of other requirements that your certificate will need to meet if your users will be using the app on iOS devices. SeeRequirements for trusted certificates in iOS 13.

Login and authentication

The app supports all common Confluence user management configurations, including external user directories and SAML single sign-on. Users will need to sign in to use the app, even if your site allows anonymous access.

User-Agent header format

The app uses anRFC-compliant user-agent header for making the requests during login. This allows usto effectively identify the app and apply further processing based on the users hardware and software.

Given the wide variety of supported configurations we decided to adopt the most versatile format. For both iOS and Android we compose the User-Agent header value by concatenating the standard WebView User-Agent and unique app identifier.

iOS User-Agent

Format

```
Mozilla/5.0 (<system-information>) <platform> (<platform-details>) Safari/<technical-version>
AtlassianMobileApp
```

Example

```
Mozilla/5.0 (iPhone; CPU iPhone OS 14_7 like Mac OS X) AppleWebKit/605.1.15 (KHTML, like Gecko)
Safari/604.1 AtlassianMobileApp
```

Android User-Agent

Format

```
Mozilla/5.0 (<system-information>) <platform> (<platform-details>) AtlassianMobileApp
```

Example

```
Mozilla/5.0 (Linux; Android 9; Mi A1 Build/PKQ1.180917.001) AppleWebKit/537.36 (KHTML, like Gecko)
AtlassianMobileApp
```

Storage and encryption

The iOS and Android apps cache some content (spaces, pages, attachments) locally on the user's device. This helps keep the app responsive when navigating around pages and spaces. We don't use any application-level encryption when storing cached data, but the device's internal storage may be encrypted by the operating system.

When a user logs out, all cached data is deleted.

We don't store passwords in the app. Instead we use session cookies, which are encrypted by default.

Session timeout

Mobile apps rely only on the 401 response status code to recognize a session timeout event. If your environment has a proxy, make sure that it doesnt override the response status code (e.g. in case of ADFS you may run into a redirect to login.microsoft.com, a 3xx status code, when your session expires). Otherwise, you might see errors like `No connection` or `Can't show the issue/page` when your session expires, and you should log in again explicitly to resolve them.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Mobile Device Management (MDM)

You can distribute the Confluence Data Center and Server app to people in your organisation using your MDM solution. For more info on how to do this, seeMobile Device Management.

Marketplace apps, themes, and visual customizations

The mobile app provides a simple, lightweight way for users to view, create, edit and collaborate on pages. Complex interactions, including those provided by Marketplace apps, such as blueprints, calendars, workflows will not be available in the app. Some third party macros may be available, depending on whether the third-party app supports rendering these macros on mobile.

Any theming or look and feel customizations you've made to your site will not be reflected in the mobile app.

Cloud services

In order to provide push notifications to users' devices, we have developed a cloud-based notification service. This service is developed and maintained by Atlassian, and is hosted on our AWS infrastructure (AWS SNS). See Push notifications service below for more information.

This is the only cloud-based service used by the app.

Push notifications service

TheConfluence Data Center and Server mobile app can push notifications directly to users' devices. Users choose whether they'd like to receive push notifications from the app, and can opt out at any time. This feature uses a cloud-based notifications service developed and maintained by Atlassian and hosted on our AWS infrastructure. No user or message content is sent to the service, only notification IDs, and we don't store any data.

If you need to avoid using any cloud-based services you can choose to disable push notifications entirely. Head to > **General Configuration**>**Mobile apps**.

If you're using restrictive firewall or proxy server settings, you'll need to allow (whitelist) `https://mobile-server-push-notification.atlassian.com` to ensure push notifications work as expected.

For sites that are not accessible on the public internet (for example users need to be connected via VPN to use the app) we adapt the push notification message as follows:

- If the user is connected to your network or VPN, we'll show the full notification, for example "Sara Leung shared 'End of year party' with you"
- If the user is not currently connected to your network or VPN, we'll show a shorter notification, for example "1 new notification".

Mobile web and linking directly to pages

It is not possible to go directly from a link, for example in an email notification, to the app. To help with this limitation, when people land ona Confluence page in their device's browser, they'll see the **Open in app** button. Tapping this prompt will open the app, if they have it installed, or take them to the App or Play store to download it.

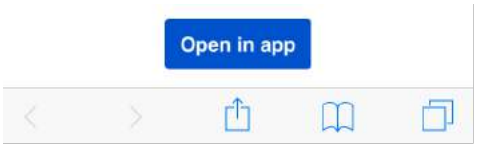

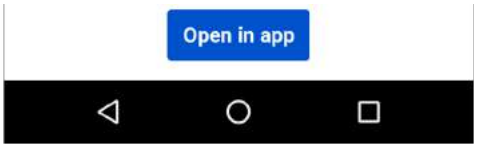

If you don't want this button to display in mobile web, you'll need to disable the entire Confluence Mobile plugin, which is required to use the mobile app. There is currently a known issue with this workaround. See

~~CONFSERVER-57423~~ - Open in app button doesn't hide when disabling Confluence Mobile Plugin while in Mobile View CLOSED

for details.

Limitations and known issues

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Not all macros are available

Not all macros will display in the app or mobile web. If a macro can't be displayed, you'll see the message below, and have the option to tap through to the desktop version of the page, in your device's browser.

Screenshot: Error that appears when a macro is not rendered in Confluence mobile

Administrators can disable Confluence mobile on your site

If you're not able to use the mobile app or mobile web, it may be because your administrator has disabled one or both of the following system apps:

- **Mobile plugin for Confluence Data Center and Server** (formerly called**Confluence mobile plugin**). This is required to use the mobile app.
- **Confluence mobile web plugin.**This is required to use mobile web.

Disabling the Mobile plugin for Confluence Data Center and Server (formerly known as'Confluence Mobile Plugin')will also disable all the modules of the Workbox - Host Plugin plugin. This issue is being tracked at

CONFSERVER-40782 - Disabling the Confluence Mobile Plugin also disables the Workbox - Host Plugin in Confluence LONG TERM BACKLOG

Seperate Cloud and Server mobile apps

While the functionality of the two mobile apps is very similar, you will need to download the **Confluence Data Center and Server mobile app** to be able to authenticate with a server site. You can't use the Confluence Cloudapp with a Confluence Server site or vice versa.

Troubleshooting

Having other problems with the app? SeeConfluence Server and Data Center mobile appsin the knowledge base.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Using the Confluence Data Center and Server mobile app

Stay connected to your team's work with the Confluence mobile app.No matter whether you have an iPhone, iPad, Android phone, or Android tablet, we've got you covered when you're away from your desk.

If you're the first in your organisation to try the app, ask your Confluence admin to have a read throughConfluence Mobile - Considerations for administrators, so they can make sure you're able to connect.

On this page:

- Confluence Data Center and Server for Android
- Confluence Data Center/Server for iOS
- Customizing push notifications

Which app do I need?

Confluence Data Center and Server app - If you're using Confluence 6.8 or later, download the Confluence Data Center and Server app from Play Store or App Store.

Confluence Cloud app - If you're using Confluence Cloud (your Confluence version number is 1000 or higher) head over to our Cloud documentation to find out about the Confluence Cloud app for iOS and Android.

Confluence Data Center and Server for Android

Here's what you'll get, and what you can do in the app:

- **Create and edit pages**
  Create a quick page when you're on site with a customer, or fix that typo on an existing page before anyone notices.Create and edit lets you do the important stuff, wherever you are.
- **Notifications when you need them most**
  Get push notifications for mentions,comment replies,page shares, and tasks assigned to you, so you can act quickly when it really matters.
- **Collaborate on the go**
  Comment on pages to keep work moving, wherever you are. Like pages to show your support, and share themvia email and other apps.
- **Get back to your work, fast**
  The**Recents**tab lets you quickly find pages you've viewed recently. Find those meeting notes you added yesterday, or the blog post you were reading earlier.
- **Visit any space, and browse using the page tree**
  The**Spaces**tab lets you visit your My Spaces, and any other space on your site. Pick a space and browse using the familiar Confluence page tree.

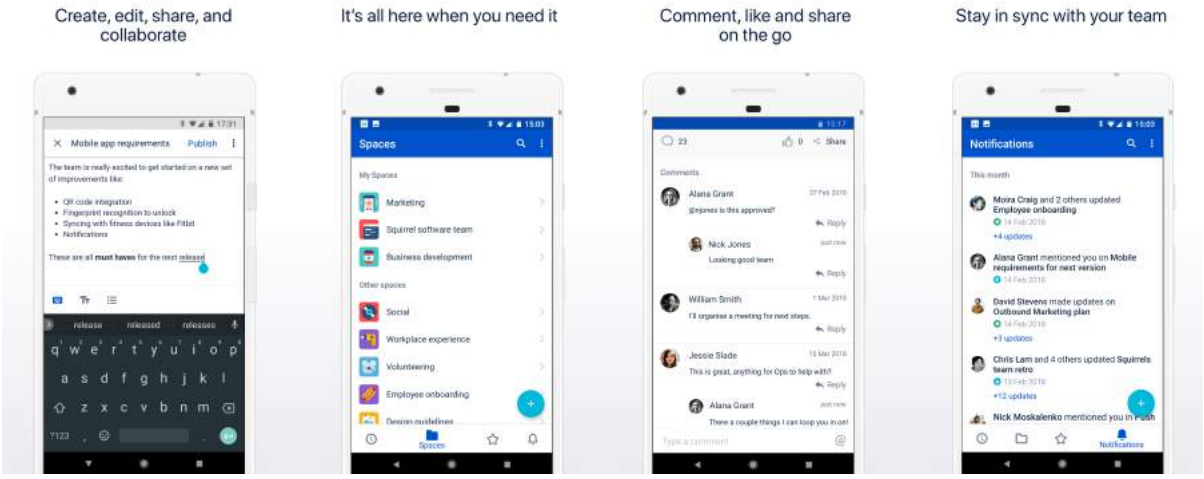

Have ideas on how to make the app even more useful? We want your feedback! Shake your phone (or head to**Settings**>**Feedback**) to drop us a note.

Limitations and known issues

- Some page macros won't display in the app or mobile web. You'll need to view the page in your browser (or switch to full desktop mode on your device).
- Image and file upload is not currently available in the app.
- Links to Confluence pages (from emails or other apps) don't automatically open in the app.
- Admins can disable push notifications for your entire site.
- We introduced a new editor in Android 0.11. If you have a problem editing a particular page tap the ellipsis (...) to temporarily switch to the legacy editor. This change doesn't apply to the iOS app.

Confluence Data Center/Server for iOS

Confluence Data Center and Server for iOS is a universal app for iPhone and iPad, so you can choose the best device for the job. If you'd rather not create and edit pages on your iPhone, switch over to your iPad and take advantage of the larger screen and keyboard. You can also use split view on iPad and work side by side with other apps.

Here's what you'll get, and what you can do in the app:

- **Create and edit pages**
  Create a quick page when you're on site with a customer, or fix that typo on an existing page before anyone notices. Create and edit lets you do the important stuff, wherever you are.
- **Notifications when you need them most**
  Keep up with what your team is doing with push notifications for new pages and posts, comments, mentions,page shares, likes, and tasks assigned to you.
- **Get back to your work, fast**
  The**Your work**tab serves up pages you've viewed or worked on recently. Find those meeting notes you added yesterday, or the blog post you were reading earlier.
- **Visit any space, and browse using the page tree**
  The**Spaces**tab lets you visit your My Spaces, and any other space on your site. Pick a space and browse using the familiar Confluence page tree. While you're in the spaces list, you canadd spaces toMy Spacesor remove the ones that aren't important any more.
- **Stay connected to your team**
  The activity feed lets you seeallactivityon your site, or filter it by space.Likeand comment on pages, orshare a link to any page,right from the app.

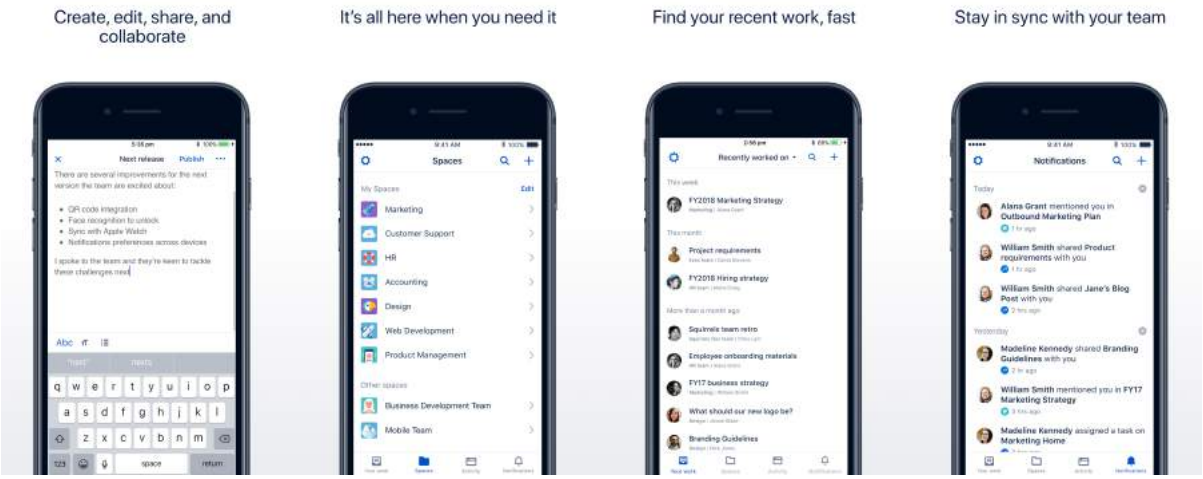

Have ideas on how to make the app even more useful? We want your feedback! Shake your phone (or head to**Settings**>**Feedback**) to drop us a note.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Download on the App Store

Limitations and known issues

- Some page macros won't display in the app or mobile web. You'll need to view the page in your browser (or switch to full desktop mode on your device).
- Image and file upload is not currently available in the app.
- Links to Confluence pages (from emails or other apps) don't automatically open in the app.
- Admins can disable push notifications for your entire site.

Did you participate in our beta?

First of all, thank you! Your feedback was invaluable to us. To keep using the app, however, you'll need to:

- **Upgrade your site** to Confluence 6.8 or later, and
- **Update your app** to version 1.x or later:
  - For Android, head to the Play Store and hit Update.
  - For iOS, head to the App Store to download the official app.

If you update your app without upgrading Confluence, the app will still work for a while, but you won't get access to new features, like push notifications. Also, once you log out, you won't be able to log back in, as the app checks that you have the required Confluence version.

Customizing push notifications

Push notifications are a great way to stay in the loop, as they appear on your device, even when you're not using the app. Tap the notification, and be taken straight into the app.

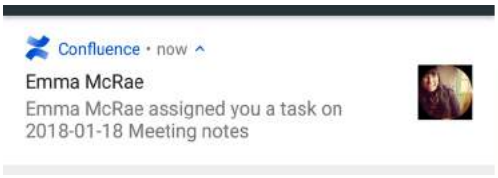

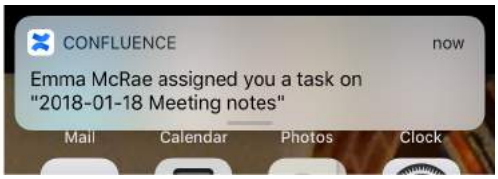

There are three notification levels, 'All activity', 'Activity for me', and 'None'. iOS users also have an additional 'Custom' option, where they can turn off individual notifications.

To change your push notification settings:

- For Android tap > **Settings** > **Push**
- For iOS tap >**Settings** > **Push**

Here's a summary of common Confluence actions, and whether a push notification is sent.

	Push notification setting		
Action	**All activity**	**Activity for me**	**None**
Someone mentions you	✓	✓	✗
Someone shares a page / blog post with you	✓	✓	✗
Someone assigns a task to you	✓	✓	✗

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Someone comments on a page / blog post you're watching	✓	✗	✗
Someone comments on a page / blog post you created	✓	✗	✗
Someone replies to your comment	✓	✓	✗
Someone likes a page / blog post you created	✓	✓	✗
Someone likes your comment or your reply	✓	✓	✗
Someone likes a page / blog post you're watching	✗	✗	✗
Someone invites you to edit a page / blog post	✗	✗	✗
Someone edits a page / blog post / comment	✗	✗	✗

If you're using the iOS app, choose 'Custom' to further tailor your notifications, and turn off any of the following notifications individually:

- Shares
- Mentions
- Tasks
- Comments on pages / blog posts you created
- Likes on pages / blog posts / comments you created
- Comments on pages / blog posts you're watching.

Good to know

- If your site isn't accessible on the public internet (for example you need to be connected to your office wifi, or use a VPN to access it from home) we adapt the push notification message, so that you get a shorter version when you're not connected to your network.

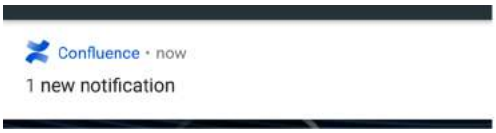

- Your admin can disable push notifications for the entire site. If this is the case, you'll see a message when you go to the Push settings screen in the app.
- On iOS,when you first install the app, you'll be prompted to allow the app to send notifications to your device. We recommend you choose **Allow**, as you can very easily mute the notifications in the app later. If you do choose **Don't allow**, and change your mind, you'll need to go to **Settings** > **Notificatio ns** > **Confluence** thenmake sure **Allow notifications** is enabled.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Using Confluence via your mobile browser

This page is about **Confluence mobile web**, which is accessed using the browser your your device.

For a richer experience, try the Confluence Data Center and Server mobile app for iOS and Android.

When you access Confluence using the browser on a mobile device, you'll see a version of Confluence which is optimized 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**.

On this page:

- The dashboard the first thing you see
- Searching for content and people
- Viewing pages, blog posts and comments
- Viewing people's profiles
- Following up on notifications
- Viewing tasks

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, oredit 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.

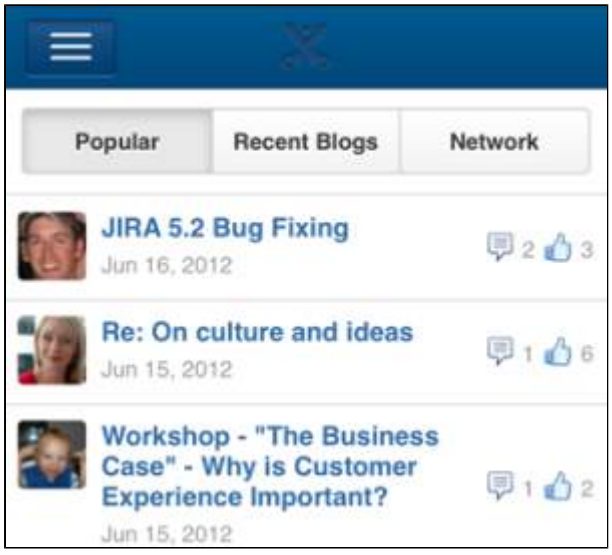

Searching for content and people

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 search, which returns matches on page title only.To use advanced search, switch to desktop mode.

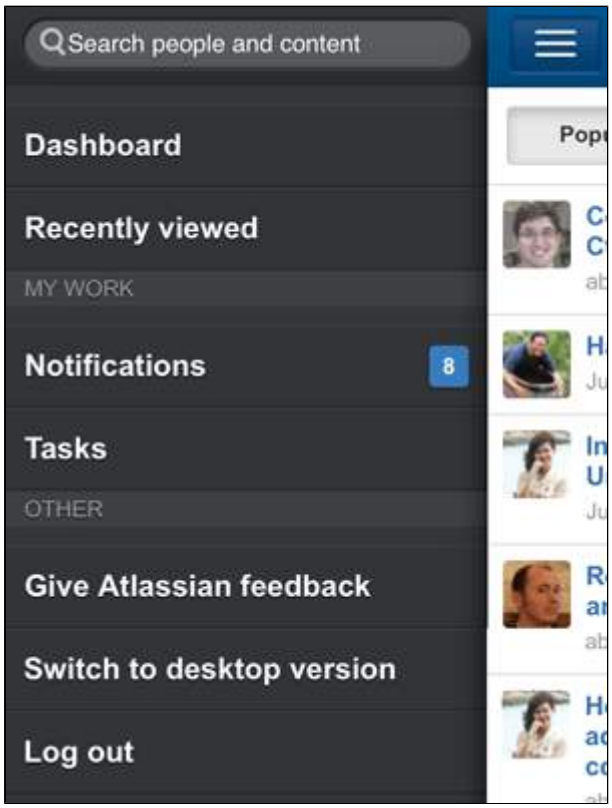

Viewing pages, blog posts and comments

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, seeWorkbox Notifications.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

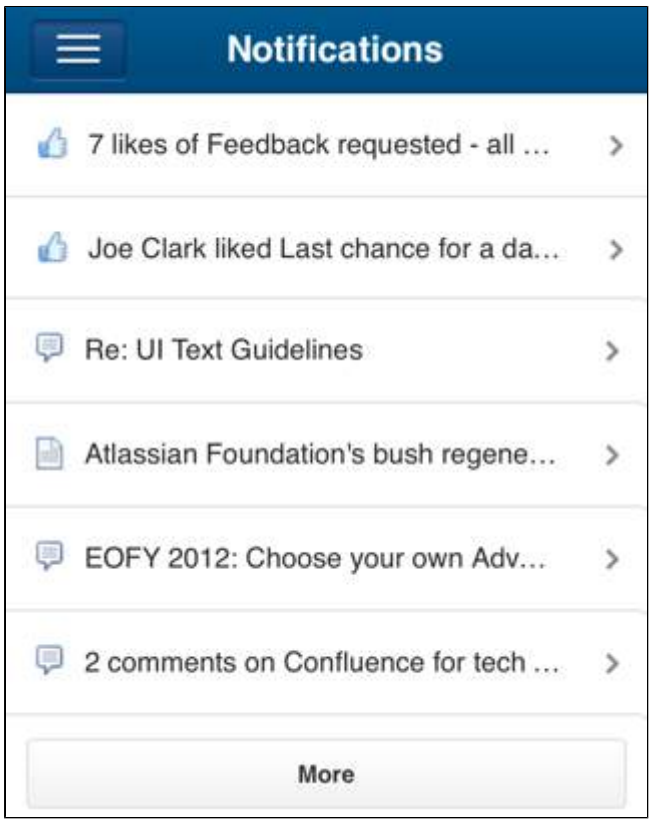

Viewing tasks

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**thentap a task to see its details.

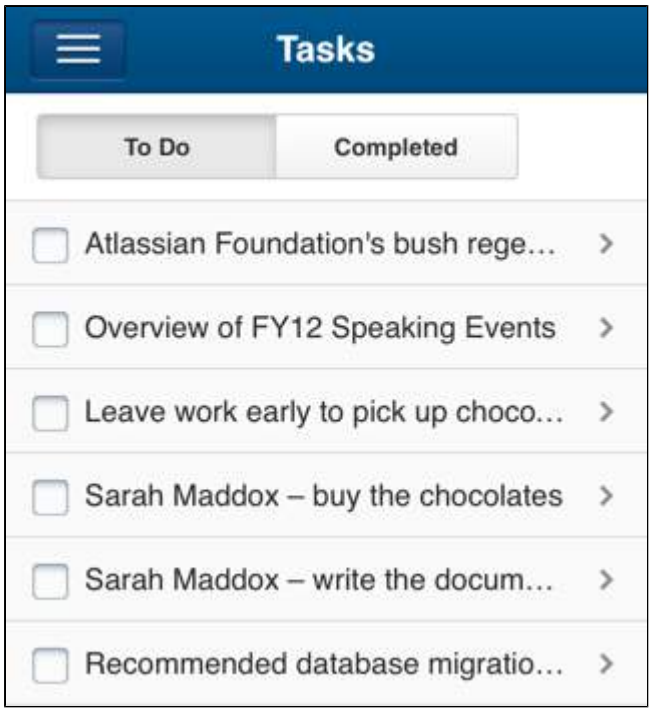

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Mobile Device Management (MDM)

You can distribute the Confluence Data Center and Server app to people in your organisation using your MDM solution. This allows you to:

- Deploy the Confluence Data Center and Server app to company-approved iOS and Android devices.
- Apply security policies to require Face ID, Touch ID, or passcode to unlock the app, or prevent users from copying, pasting or sharing text. Requires Confluence Data Center and Server iOS app version 1.25.0 or later.
- Pre-populate your Confluence site URLs (just the URL, we don't pass login credentials).People will still be able to enter a URL, which is useful if you have some rarely used sites, and don't want to pre-populate them all. Requires Confluence Data Center and Server iOS app version 1.8.0 / Android app version 0.4.3 or later.

MDM Function	Supported for Confluence Data Center and Server app
App deployment	Yes
App configuration	Yes - to pre-populate the site URL only
App tunnel	Yes
Single sign-on (SSO)	No
Security policies & access control	Yes

Read more about AppConfig and what we currently support in our AppConfig Technical Capabilities white paper.

Using the Confluence Cloud app? Learn how to configure MDM for cloud apps

Distribute the app to managed devices

The way you do this will depend on your particular MDM provider. In most cases you will:

- Add the app to your MDM app catalog.
  - For iOS you will likely bulk purchase through the Apple VPP store, then link the app in your MDM solution (this might be by providing an sToken, or could happen automatically)
- Choose which devices should be able to install the app (this might be called something like distributing or assigning the app).

Refer to the documentation for your MDM provider for more information.

Apply security policies (iOS only)

The keys below are only applicable to iOS because Android supports security policies by MDM providers out of the box.

Use security policies to control things like copy / paste, sharing text, and unlocking the app. Requires Confluence Data Center and Server iOS app version 1.25.0 or later.

Key	Description
Prevent cut, copy, and paste	

`atlassian.devicePolicyClipboardRestriction`	**Value type:** Integer (`1` or `0`, any other value will result in an error) **Description:** Set this value to `1` to restrict clipboard use. This prevents users from cutting, copying, and pasting data from the app to other apps. Set this value to `0` to allow clipboard use. **Note:** Users will be able to copy and paste between Jira and Confluence Data Center or Cloud apps if those apps also set this value to `1`.
Prevent sharing text to other devices	
`atlassian.devicePolicyDataExportRestricted`	**Value type:** Boolean (true/false) **Description:** Set this value to `true` to prevent text from being shared to other devices via the iOS share sheet.
Require Face ID, Touch ID, or password to unlock the app	
`atlassian.devicePolicyLocalAuthRequired`	**Value type:** Boolean (true/false) **Description:** Set this value to `true` to require users to unlock the app using Touch ID, Face ID, or a passcode.
`atlassian.devicePolicyLocalAuthTimeout`	**Value type:** Integer (number greater than zero) **Description:** Enter a number in seconds (for example `100`) to define the time that the app can be open but not in use (in the background), before a passcode is required again. If you dont set a value, the default timeout is `120` seconds.

Pre-populate site URLs on the login screen

If your MDM solution supports theAppConfigstandard, you can save your users time and prevent mistakes by pre-populating your site URL in the app.

To pre-populate the mobile app login screen with one site URL:

1. In your MDM, navigate to the App Config section. Check the documentation for your MDM for how to do this.
2. Add a new key called "sites"
3. In the Value field for the key, enter your site title and URL in JSON format as shown in the examples below. Replace the title and base URL with your own site details.

   For a single URL:

   ```
   [
   { "title": "My Confluence Site", "baseURL": "https://conf.example.com"}
   ]
   ```

   For multiple URLs:

   ```
   [
    {"title": "My Docs Site", "baseURL": "https://docs.example.com"} ,
    {"title": "Intranet", "baseURL": "https://team.example.com/confluence", "skipInfo": false}
   ]
   ```

   Copy the JSON carefully, including the[and].
4. Save your changes.We recommendyou access the Confluence Data Center and Server app on a device to check you have passed the URLs correctly, before distributing the app config changes to your users. You'll see an error if the app can't display your list of sites.

Parameters

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

The following parameters are accepted for the`sites` key in app config.

Parameter	Required	Description
title	Yes	Descriptive name of your Confluence instance. This will appear above the URL on the app login screen.
baseURL	Yes	The base URL of your Confluence site. To check it, go to > **General Configuration**.
skipInfo	No	Use this parameter if you have log in problems in because your proxy server blocks any unauthenticated requests. When set to `true`, the mobile app wont make login/server-info calls that require authorization and instead will redirect your users to your single sign-on page. In general, the Confluence Mobile plugin is a prerequisite for using the mobile app. However, if the `skipInfo` parameter is set to true, and the plugin is disabled, your users will see a desktop version of Confluence and be able to log into it, which is not the correct behavior, but might not look like an obvious error. (Requires iOS app 1.14.0 / Android app 0.6.2)

Here's an example of how the AppConfig looks in AirWatch and MobileIron, two popular MDMs. These screenshots are for the Jira Data Center and Server app, but the process is the same for Confluence.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

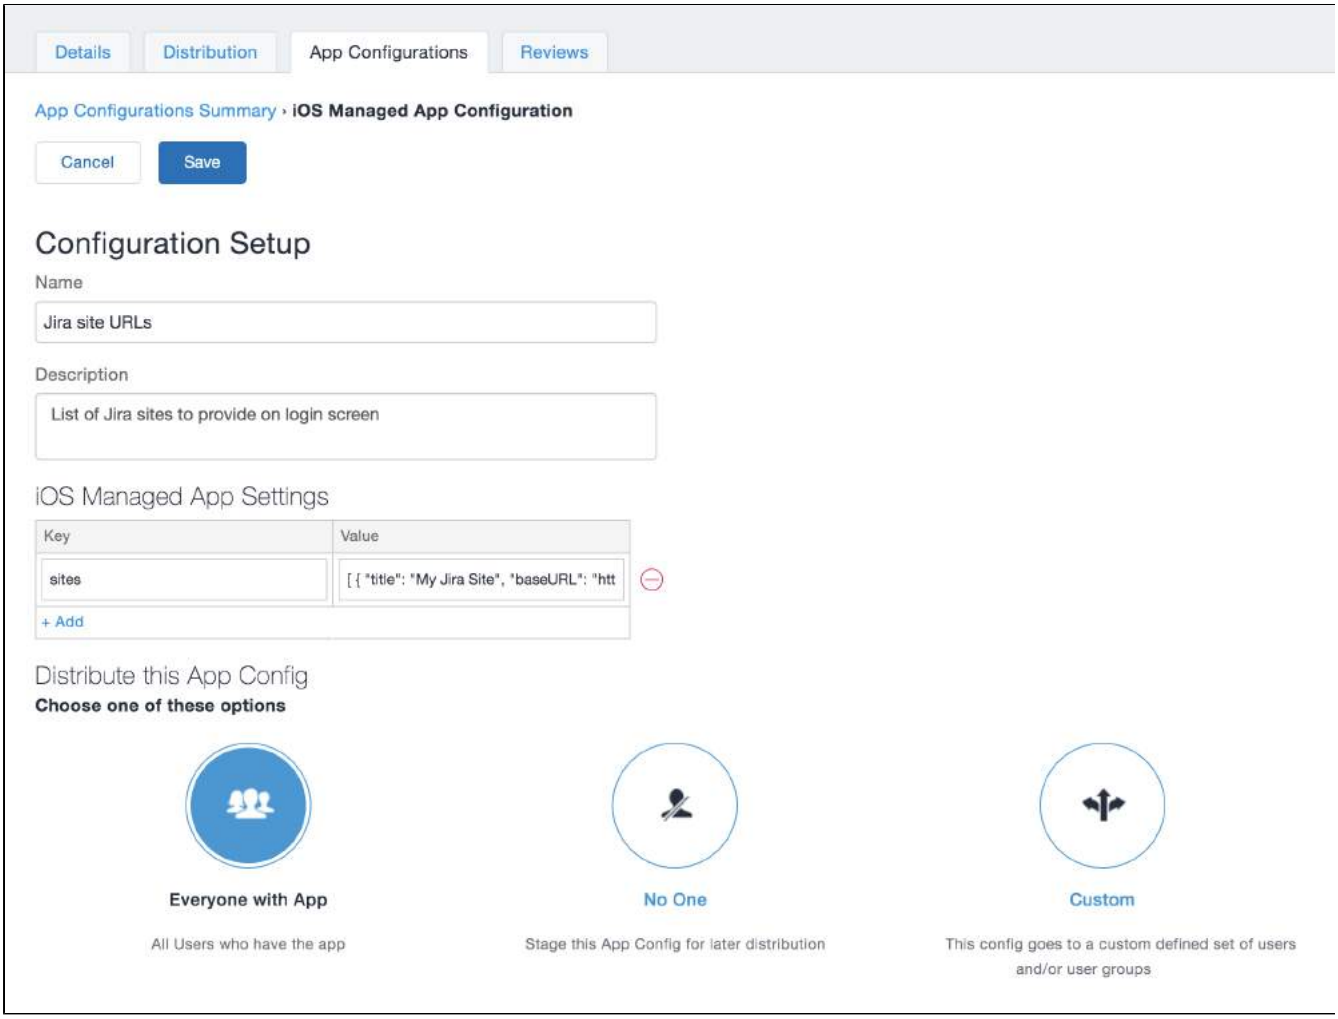

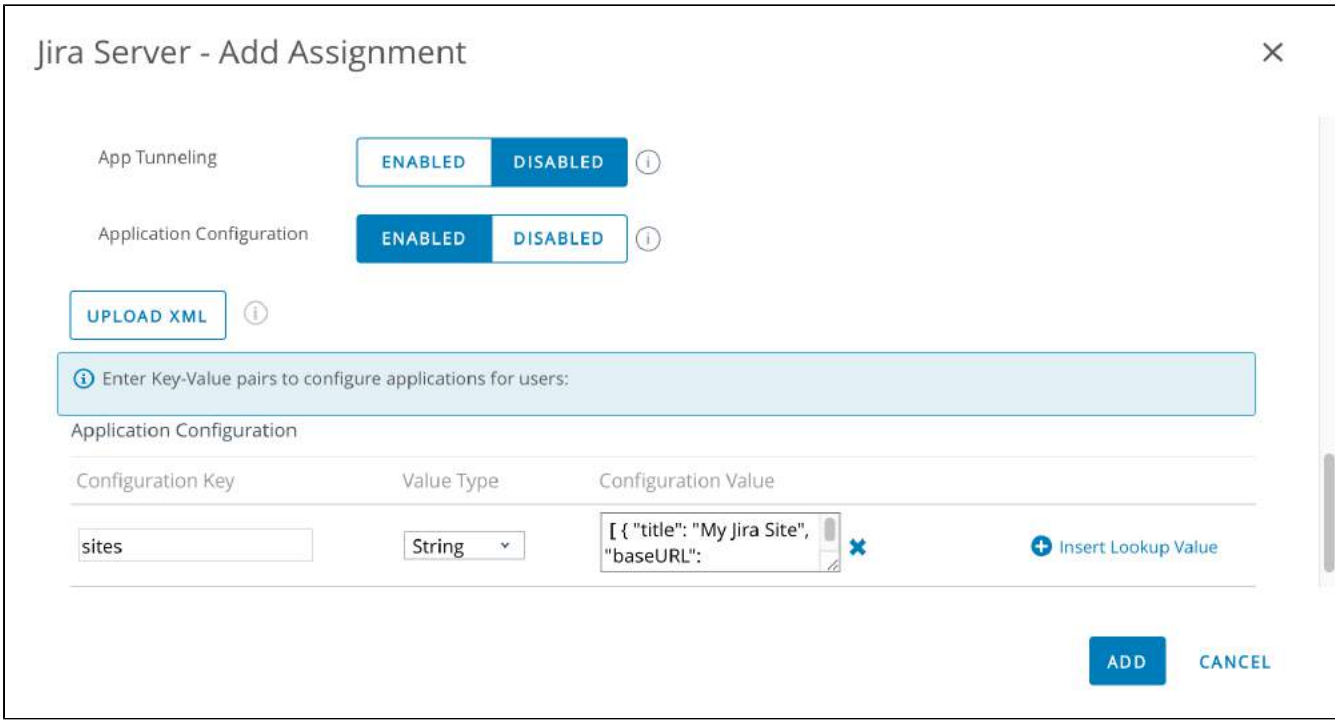

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Macros

Macros are how you supercharge your Confluence pages.

You can use macros to:

- change the format and layout of your page
- display media like video, audio, and social media content
- collate and organise Confluence pages, blogs, and files
- perform actions from a page, such as creating a page from a template.

Take your Confluence space to the next level using macros.

On this page:

- Macro basics
- Confluence macros
- Get more macros
- Macro usage statistics
- Unknown macro

Screenshot: Page containingStatus, Page Properties Report, Livesearch, and Profile Picture macros tohelp people find information about particular projects.

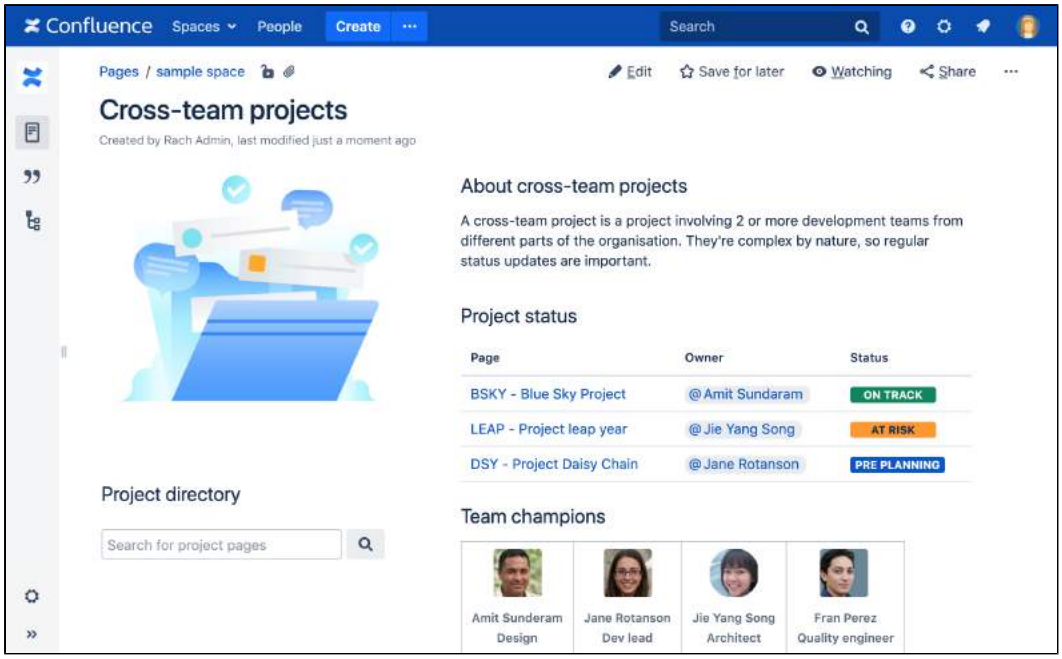

Macro basics

Add a macro to your page

To add a macro to your page:

1. From the editor toolbar, choose **Insert** + > **Other Macros**.
2. Select a macro from the list.
3. Enter any required parameters.
4. Choose**Insert**.

In the editor you'll see a placeholder that represents the macro. Once you publish your page, you'll see the macro in its full glory.

Edit a macro

Macro parameters are used to change the behaviour of a macro.

To change the macro parameters:

1. In the editor, click the macro placeholder and choose **Edit**.

2. Update the parameters as required then choose **Insert**.

Head to the documentation for specific macros below for full details of the parameters available in each macro.

Confluence macros

Here's a list of macros currently available with Confluence Server and Data Center.

Click a macro name for details of the usage, including optional parameters and examples.

- Advanced Roadmaps 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
- Favorite Pages 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
- 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

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

- 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
- Spaces List Macro
- Status Macro
- Table of Contents Macro
- Table of Content Zone Macro
- Task Report Macro
- Team Calendar Macro
- User List Macro
- User Profile Macro
- View File Macro
- Widget Connector Macro

Get more macros

This documentation provides information on all the macros that are provided with Confluence by default. There may be other macros in your site, from Marketplace Apps, or perhaps developed in-house.

Third-party macros from the Atlassian Marketplace

You can find a wide range of Atlassian and third party macros atThe Marketplace. These are distributed as apps and can be installed by a Confluence Administrator.

Create your own macros

Users with System Administrator permissions can create user macros - seeWriting User Macros.

If you want to create something more complex, you can develop your own plugin - seeWriting Confluence Plugins.

Macro usage statistics

Confluence administrators can check how often each macro is used in your site.

To see how often a macro is used, go to > **General Configuration**>**Macro Usage**. This listshow often each macro is used in current spacesbut doesn'tinclude any macros used on pages in archived spaces or macros provided by disabled apps.

Unknown macro

In view page mode, disabled macros show a placeholder image:

When macros with body content are disabled, the content will be preserved and display in the editor.

To display the content in view mode:

1. In the editor, select the macro and choose **Edit**.
2. In the macro browser, select **Back.**
3. Search for a new macro and select it.
4. Check your content displays in the **Preview** pane.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

5. Select **Insert**.

In the editor, the disabled macro will switch to the new macro and display your content. Update the page to display the content in view mode.

Do more with Confluence

Extend Confluence with one of the hundreds of other macros in theAtlassian Marketplace, such as:

- Composition Tabs & Page Layouts: Toggle or expand the visible of portions of your pages with the Toggle and Cloak macro
- Content Formatting for Confluence: Over 30 easy-to-use Confluence macros gives you the ability to create better, more engaging content

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Advanced Roadmaps Macro

Add the Advanced Roadmaps for Jira in Confluence macro to a page to share your team's plans and progress with stakeholders.

This is great for:

- reports and retrospectives
- project status updates
- sharing updates with people in your organization who don't use Jira regularly.

The macro can display issues from any connected Jira Server, Data Center, or Cloud application, including Jira Software and Jira Service Management.

On this page:

- Connect Confluence and Jira
- Macro basics

Screenshot: An Advanced Roadmaps plan on Jira displayed on a Confluence page.

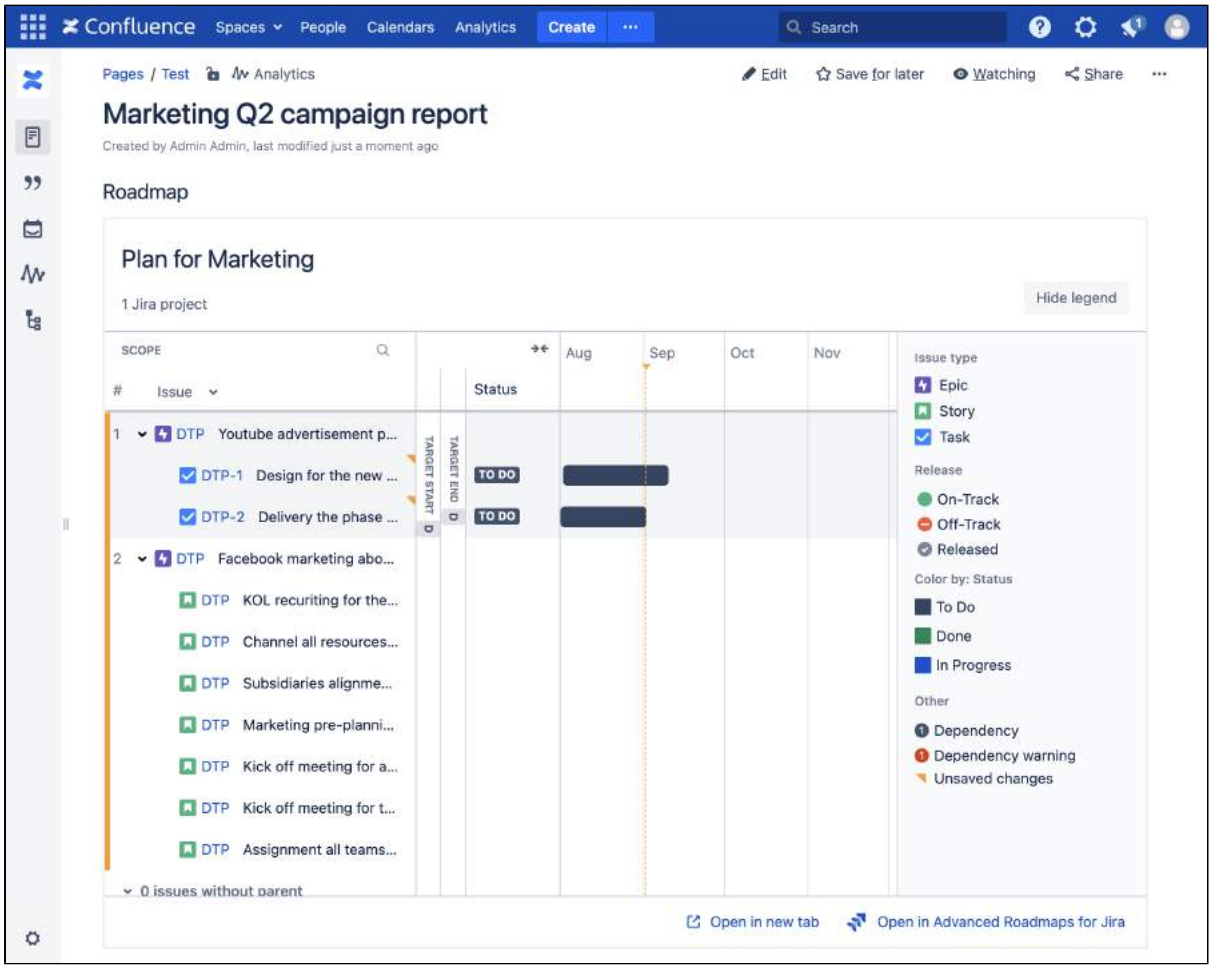

Connect Confluence and Jira

Before you can use this macro, your Confluence and Jira application must be connected viaApplication 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.

SeeUse Jira applications and Confluence togetherfor more information.

Macro basics

Learn how to add this macro to your page

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Anchor Macro

Add the Anchor macro to a page to link to a specific part of a page.

This is useful on long pages, where you want to link to specific parts of the page.

> **Don't see the Anchor macro?** This macro isn't available in the new Confluence Cloud editor. See We're cleaning up our macros for alternative ways to link.

The example below shows an example of an Anchor macro as it appears in the editor, and as it would appear to someone viewing the page.

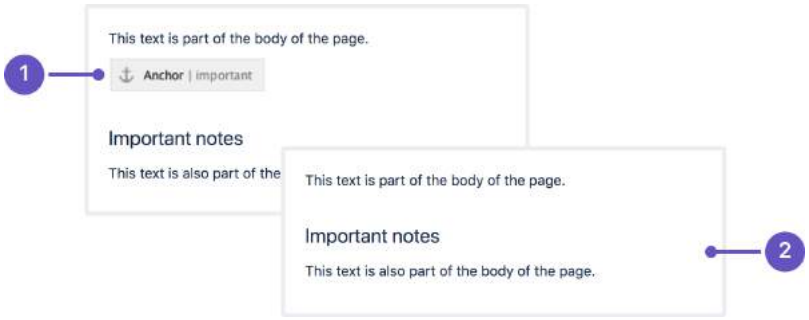

1. Anchor macro as it appears in the **editor**
2. Anchor macro as it appears when **viewing a page** (it isn't visible).

On this page:

- Add the anchor macro to your page
- Linking to an anchor
- Change the macro parameters
- Other ways to add this macro

Add the anchor macro to your page

To add the Anchor macro to a page:

1. From the editor toolbar, choose **Insert** + > **Other Macros**.
2. Choose**Anchor**from the**Confluence content**category.
3. Enter the anchor name - this will form part of your link.
4. Choose**Insert**.

You can then add a link to your macro from this page or another page.

Linking to an anchor

To link to an anchor on the same page:

1. Click the **Insert Link** icon.
2. Choose the **Advanced** tab.
3. Enter # followed by your **anchor name** in the **Link** field. For example`#top`.
4. Choose **Insert**.

You can also link to an anchor on another page. SeeAnchorsfor more information on the different link syntax you can use.

Animation: Adding the Anchor macro to a page, then linking to that macro.

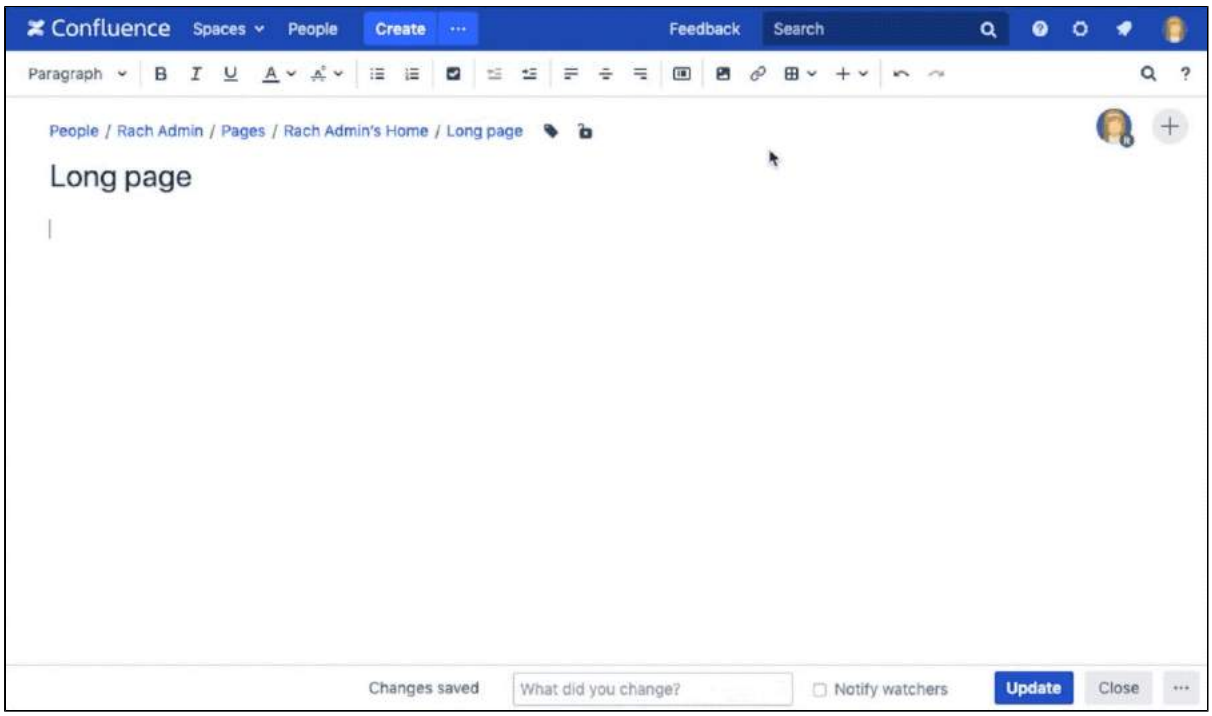

Change the macro parameters

Macro parameters are used to change the behaviour of a macro.

To change the macro parameters:

1. In the editor, click the macro placeholder and choose**Edit**.
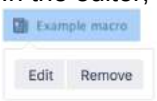

2. Update the parameters as required then choose**Insert**.

Here's a list of the parameters available in this macro.

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.

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`).

Other ways to add this macro

Add this macro as you type

Type**{**followed by the start ofthe macro name, to see a list of macros.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

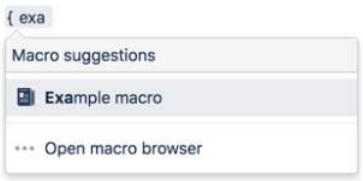

Add this macro using wiki markup

This is useful when you want to add a macro outside the editor, for example as custom content in the sidebar, header or footer of a space.

Macro name:`anchor`

Macro body: None.

```
{anchor:here}
```


Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Attachments Macro

This macro is available in **Confluence Server and Data Center**. Learn about the macros available in Confluence Cloud.

Add an Attachments macro to your page to display a list of files attached to your current page, or another page in your site.

This macro is great for providing quick access to:

- project files
- forms and downloadable templates
- images and diagrams.

Because you can display files attached to any page, you can use this macro to avoid duplication if you need to provide quick access to the same file on multiple pages.

On this page:

- Add this macro to your page
- Change the macro parameters
- Edit a file displayed by this macro
- Other ways to add this macro

Related pages:

- Display Files and Images

Screenshot: The Attachments macro, showing details of an attachment.

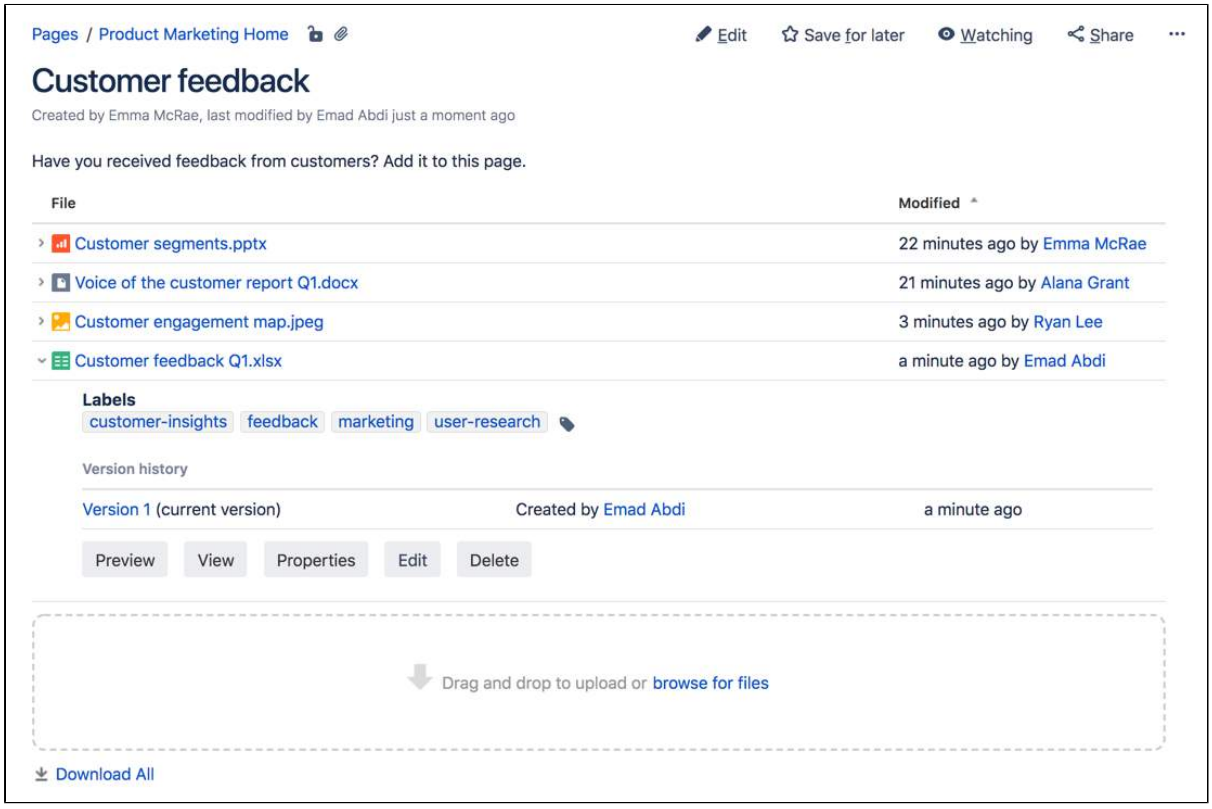

Once added to a page, people with appropriate space permissions can:

- view a list of attached files
- 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 files
- download all files attached to the page as a zip file.

For general information about working with files in Confluence, seeFiles.

Add this macro to your page

To add the Attachments macro to a page:

1. From the editor toolbar, choose **Insert** + > **Other Macros**.
2. Choose**Attachments** from the **Confluence content** category.
3. Set any parameters. These are all optional.
4. Choose **Insert**.

You can then publish your page to see the macro in action.

Screenshot: Entering parameters and changing the sort order in the Blog Posts macro.

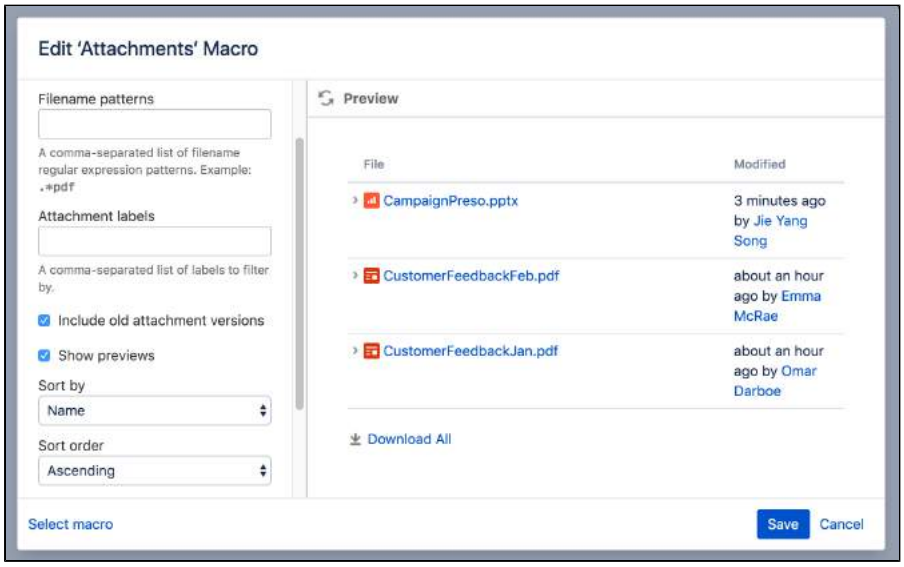

Change the macro parameters

Macro parameters are used to change the behaviour of a macro.

To change the macro parameters:

1. In the editor, click the macro placeholder and choose**Edit**.
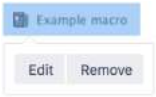

2. Update the parameters as required then choose**Insert**.

Here's a list of the parameters available in this macro.

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.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

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. It can be useful to disable previews if you have very large attachments.

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`).

Edit a file displayed by this macro

There are a few ways to edit attachments in Confluence.

To edit a file from the attachments macro list:

1. Click the arrow next to the file name to view its version history.
2. Click**Edit**.
3. Atlassian Companion will open the file in your desktop application.
4. Make your changes and then save your file. When you're ready, click **Upload** in Companion to send the file back to Confluence.

Learn more aboutediting files in Confluence.

Other ways to add this macro

Add this macro as you type

Type**{**followed by the start ofthe macro name, to see a list of macros.

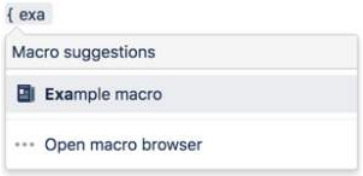

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Add this macro using wiki markup

This is useful when you want to add a macro outside the editor, for example as custom content in the sidebar, header or footer of a space.

Macro name: `attachments`

Macro body: None.

```
{attachments:old=false|patterns=.*png,.*jpg|sortby=name|page=My page about
chocolate|sortorder=descending|labels=chocolate,cookies|upload=false|preview=false}
```


Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Blog Posts Macro

This macro is available in **Confluence Server and Data Center**. Learn about the macros available in Confluence Cloud.

On this page:

- Add the Blog Posts macro to your page
- Change the macro parameters
- Other ways to add this macro

Add the Blog Posts macro to a page to display a curated list of blog posts. You can choose to show the just the title, an excerpt from the blog, or the entire contents of each blog post.

This macro is great when you want to present a curated list of blogs for:

- company announcements
- new team member introductions
- point-in-time project updates
- change management communications.

Because you can display blog posts from any space, with any label or author, you can display the same blogs on multiple pages. This reduces duplication and helps people in your team find information when they need it.

Screenshot: The Blog Posts macro, configured to show an excerpt of each blog post.

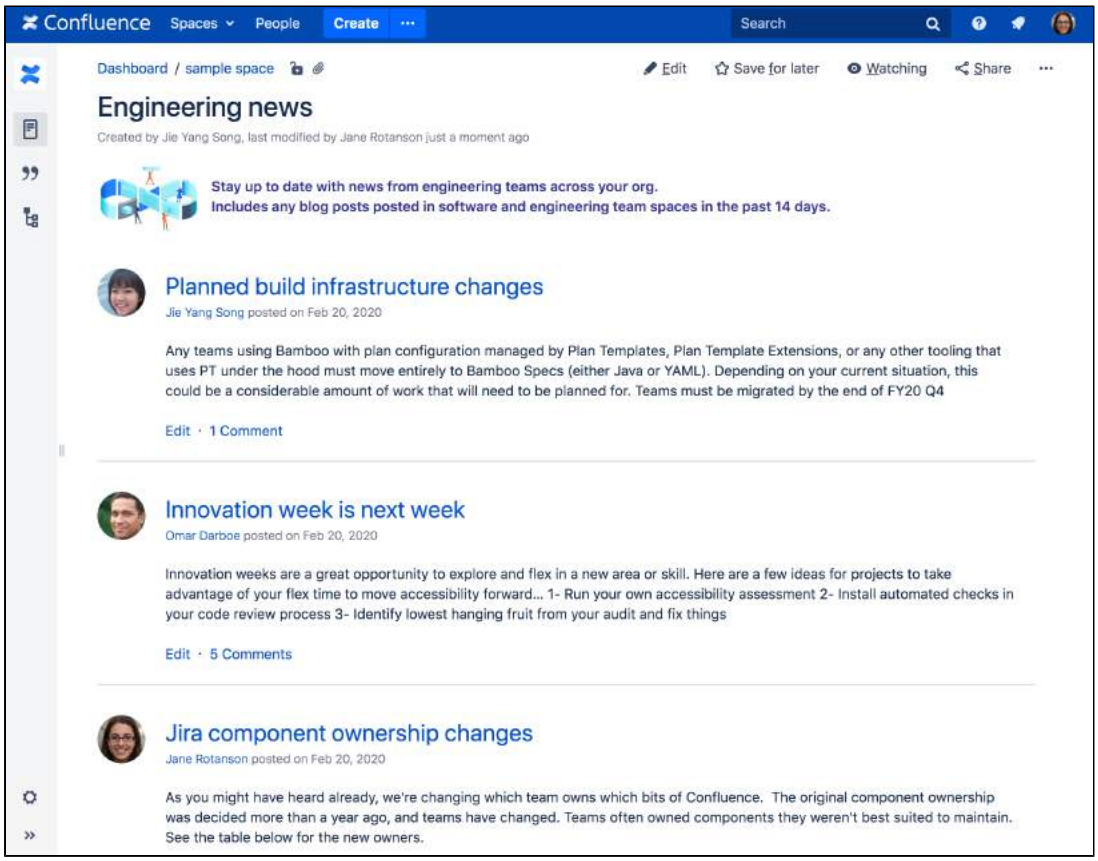

For general information about blogging in Confluence, see Blog Posts.

Add the Blog Posts macro to your page

To add the Blog Posts macro to a page:

1. From the editor toolbar, choose **Insert** + > **Other Macros**.
2. Choose**Blog Posts**from the**Confluence content**category.

3. Use the parameters below to determine how you want the blog posts to display, and to narrow your query by time frame, space, author, or label.
4. Choose**Insert**.

You can then publish your page to see the macro in action.

Screenshot: Entering display type, time frame, and label parameters in the Blog Posts macro.

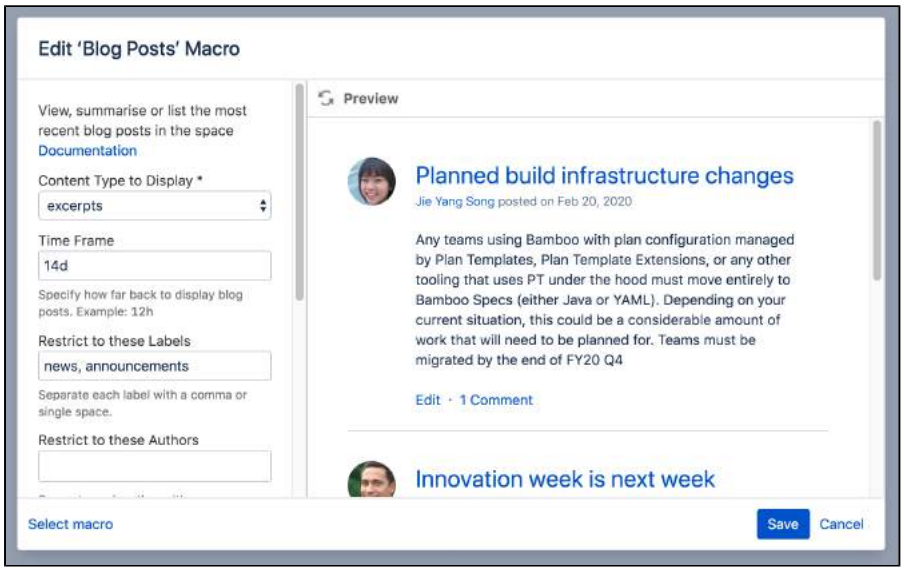

Change the macro parameters

Macro parameters are used to change the behaviour of a macro.

To change the macro parameters:

1. In the editor, click the macro placeholder and choose**Edit**.
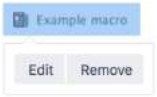

2. Update the parameters as required then choose**Insert**.

Here's a list of the parameters available in this macro.

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.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

			• 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 labeled 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 Authors (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 favorite. • @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.
	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.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Maximum Number of Blog Posts (`max`)			
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.

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`).

Other ways to add this macro

Add this macro as you type

Type**{**followed by the start ofthe macro name, to see a list of macros.

Add this macro using wiki markup

This is useful when you want to add a macro outside the editor, for example as custom content in the sidebar, header or footer of a space.

Macro name:`blog-posts`

Macro body: None.

```
{blog-posts:content=titles|spaces=@self,
ds|author=jsmith|time=4w|reverse=true|sort=creation|max=10|label=chocolate,cookies}
```


Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Change History Macro

This macro is available in **Confluence Server and Data Center**. Learn about the macros available in Confluence Cloud.

On this page:

- Add the Change History macro to your page
- Change the macro parameters
- Other ways to add this macro

Add the Change History macro to a page todisplay a table of recent updates to that page including version number, author, date and version comment.

This macro is great for:

- document control
- change management
- wiki gardening and keeping your pages fresh.

Screenshot: The Change History macro in Confluence showing two history versions.

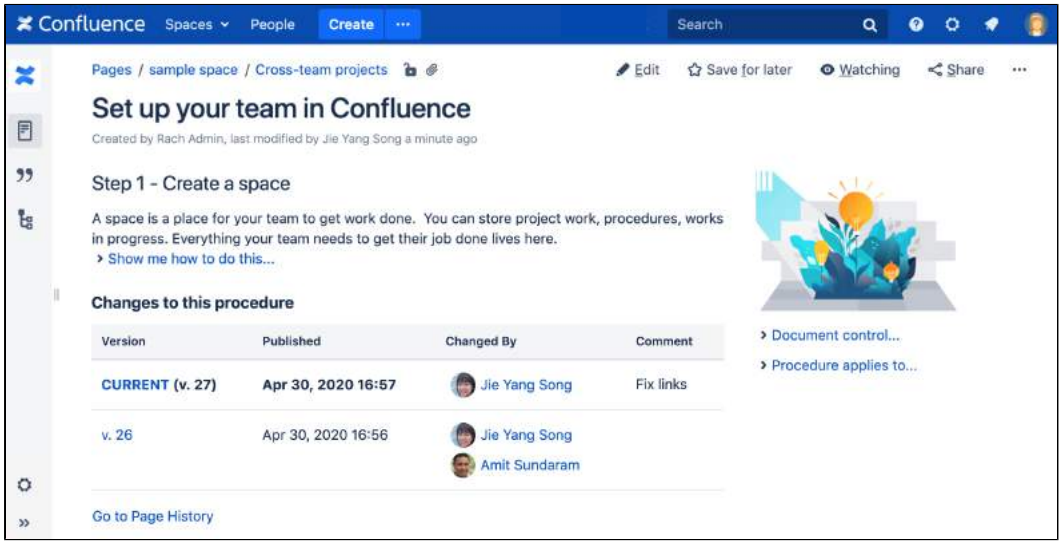

For general information about version control in Confluence, seePage History and Page Comparison Views.

Add the Change History macro to your page

To add the Change History macro to a page:

1. From the editor toolbar, choose **Insert** + > **Other Macros**.
2. Choose**Change History**from the**Confluence content**category.
3. Choose the number of page versions to display.
4. Choose**Insert**.

You can then publish your page to see the macro in action.

Screenshot: configuring the Change History macro in the macro browser.

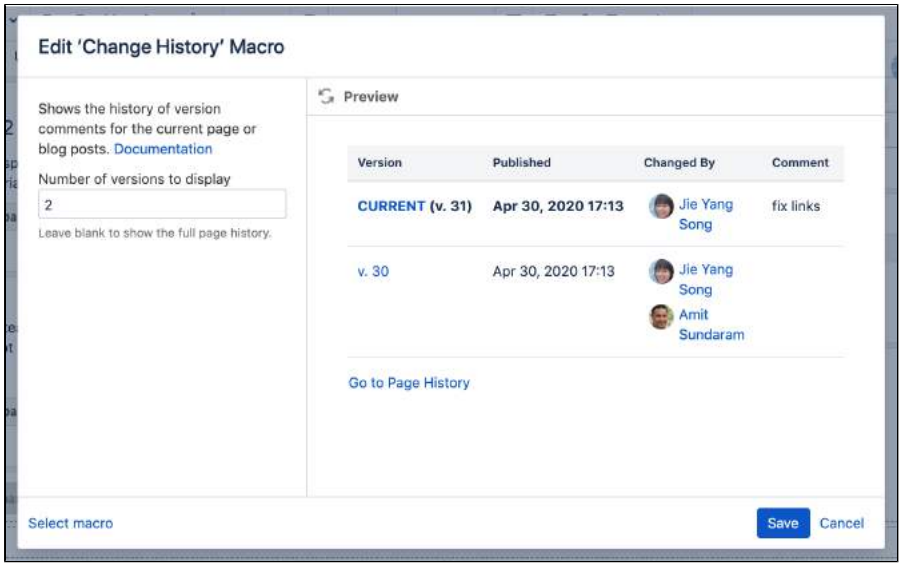

Change the macro parameters

Macro parameters are used to change the behaviour of a macro.

To change the macro parameters:

1. In the editor, click the macro placeholder and choose**Edit**.

2. Update the parameters as required then choose**Insert**.

Here's a list of the parameters available in this macro.

Parameter	Required	Default	Description
Number of versions to display (`limit`)	No	blank	Limit the amount of page history to display. Leave blank to display all versions in the page history.

Other ways to add this macro

Add this macro as you type

Type**{**followed by the start ofthe macro name, to see a list of macros.

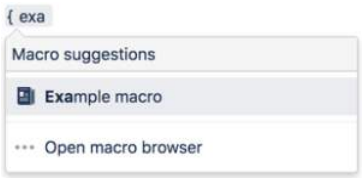

Add this macro using wiki markup

This is useful when you want to add a macro outside the editor, for example as custom content in the sidebar, header or footer of a space.

Macro name: `change-history`

Macro body: None.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

```
{change-history:limit=2}
```


Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Chart Macro

This macro is available in **Confluence Server and Data Center**. Learn about the macros available in Confluence Cloud.

Add the Chart macro to a page to display a chart based on data in a table on the same page, or from an attached file.

This is great for showing a simple visualisation of data on the page.

On this page:

- Add the Chart macro to your page
- Change the macro parameters
- Pie chart
- Bar chart
- 3D Bar chart
- Time series chart
- XY Line Chart
- XY Area Chart
- Area chart
- Stacked area chart
- Gantt chart
- Other ways to add this macro
  - Add this macro as you type
  - Add this macro using wiki markup
  - Chart type parameters
  - Chart display parameters
  - Chart title and label parameters
  - Chart data parameters
  - Chart color parameters
  - Chart axis parameters
  - Pie chart Parameters
  - Chart attachment parameters

Screenshot: Page with two chart macros.

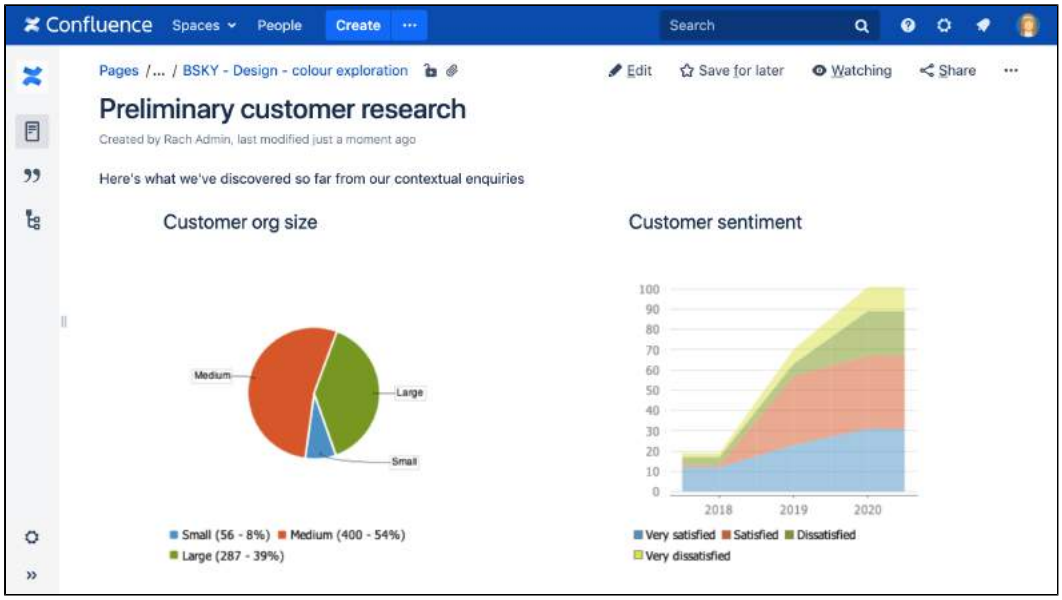

Want to display information from Jira on your page? Check out the Jira Chart Macro.

Add the Chart macro to your page

To add the Chart macro to a page:

1. From the editor toolbar, choose **Insert** > **Other Macros**.
2. Choose**Chart**from the**Visuals and images**category.
3. Choose**Insert**.

4. Enter your chart data as one or more tables in the body of the macro placeholder. See the examples later in this page for more info.
5. Click the macro placeholder and choose**Edit**.
6. Select a chart type using the**Type**parameter (see below).
7. Choose other parameter settings in the macro browser, as described below.
8. Click**Refresh**in the 'Preview' area, to check that the chart appears as you expect.
9. Click**Save**to add the chart to your page.

You can then publish your page to see the macro in action.

Screenshot: Two Chart macros in the editor, containing the data for a pie chart and stacked chart.

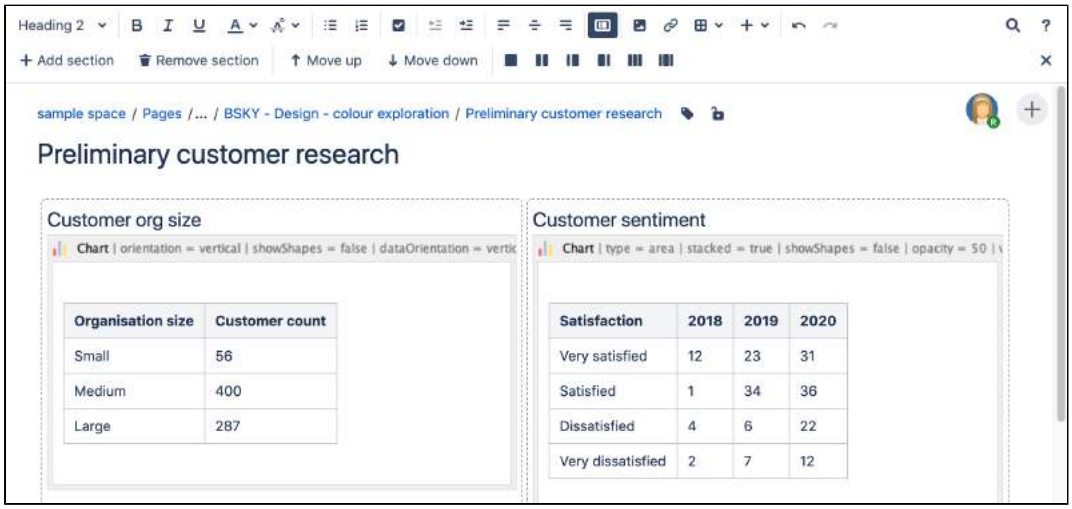

Change the macro parameters

Macro parameters are used to change the behaviour of a macro.

To change the macro parameters:

1. In the editor, click the macro placeholder and choose**Edit**.
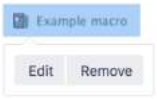

2. Update the parameters as required then choose**Insert**.

Here's a list of the parameters available in this macro.

Chart Type Parameters | Display Control Parameters | Title and Label Parameters | Data Specification Parameters | Color Parameters | Axis Parameters | Pie Chart Parameters | Attachment 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

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

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. The maximum width is limited by the`confluence.chart.macro.width.max` system property.
Height	300	The height of the chart in pixels. The maximum height is limited by the`confluence.chart.macro.height.max` system property.
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.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

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. Theseadditional 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.

Color Parameters

Colors are specified usinghexadecimal notation or HTML color names.

Parameter	Default	Description

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Background Color	White	Background of the chart.
Border Color	no border	Border around the chart.
Colors		Comma-separated list of colors used to customize category, sections, and series colors.

Axis Parameters

Depending on the chart type, the range and domain axis may be customized. 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

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

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 authorized 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.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

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 | Color Parameters | Axis Parameters | Pie Chart Parameters | Attachment Parameters

Pie chart

Here's an example of a pie chart.

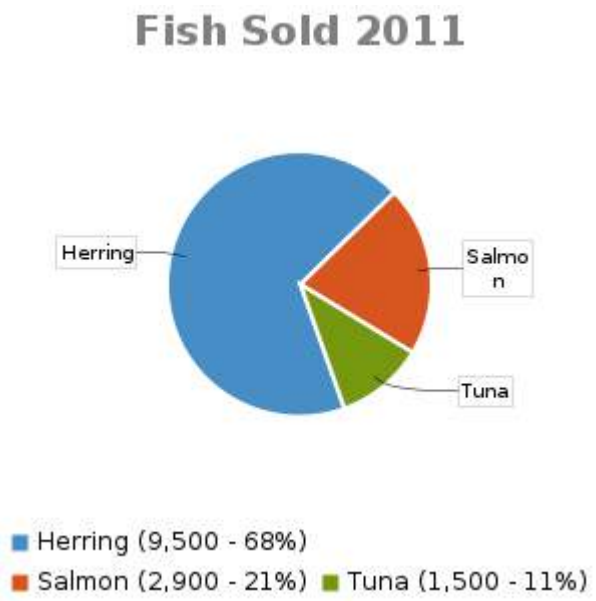

To create this chart, we set these parameters in the macro browser:

- **Type**: pie
- **Chart title**: Fish sold in 2011
- **Show legend**: true
- **Content orientation**: vertical

and added this table in the macro body:

Fish Type	2011
Herring	9,500
Salmon	2,900
Tuna	1,500

Bar chart

Here's an example of a bar chart.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

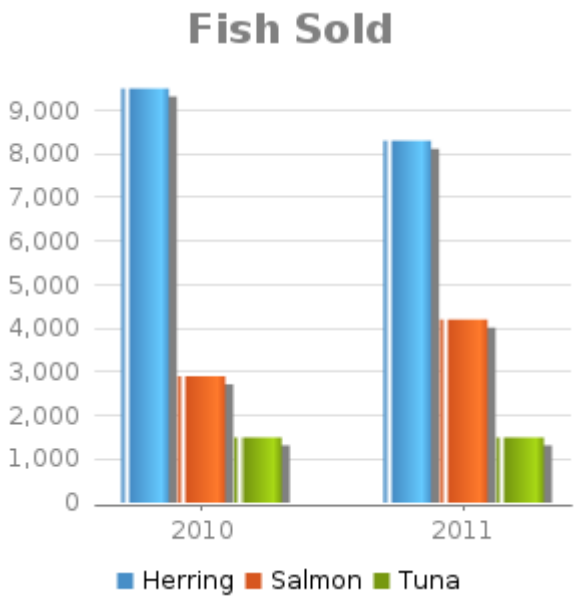

To create this chart, we set these parameters in the macro browser:

- **Type**: bar
- **Chart title**: Fish sold
- **Show legend**: True

and added this table in the macro body:

Fish Type	2010	2011
Herring	9,500	8,300
Salmon	2,900	4,200
Tuna	1,500	1,500

3D Bar chart

Here's an example of a 3D bar chart.

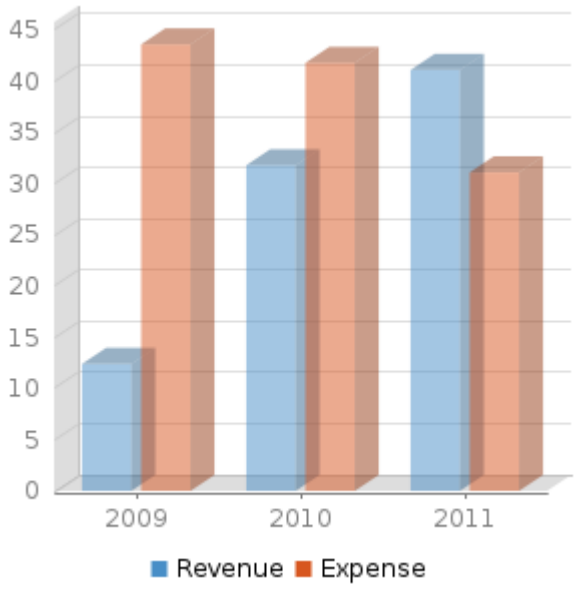

To create this chart, we set these parameters in the macro browser:

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

- **Type**: bar
- **Show in 3D**: true
- **Opacity**: 50
- **Show legend**: true

and added this table in the macro body:

2009	2010	2011
Revenue	12.4	31.8
Expense	43.6	41.8

Time series chart

Here's an example of a Time series chart.

To create this chart, we set these parameters in the macro browser:

- **Type**: Time series
- **Date format**:MM/yyyy
- **Time period**: Month
- **Content orientation**: vertical
- **Range axis lower bound**: 0
- **Show legend**: true

and added these two tables in the macro body:

Month	Revenue
1/2011	31.8
2/2011	41.8
3/2011	51.3
4/2011	33.8
5/2011	27.6

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

6/2011	49.8
7/2011	51.8
8/2011	77.3
9/2011	73.8
10/2011	97.6
11/2011	101.2
12/2011	113.7

Month	Expenses
1/2011	41.1
2/2011	43.8
3/2011	45.3
4/2011	45.0
5/2011	44.6
6/2011	43.8
7/2011	51.8
8/2011	52.3
9/2011	53.8
10/2011	55.6
11/2011	61.2
12/2011	63.7

XY Line Chart

Here's an example of a XY Line chart.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

To create this chart, we set these parameters in the macro browser:

- **Type**: xyLine
- **Show legend**: true

and added this table in the macro body:

	12	14	23
Revenue	41.1	31.8	12.4
Expense	31.1	41.8	43.6

XY Area Chart

Here's an example of an XY Area chart.

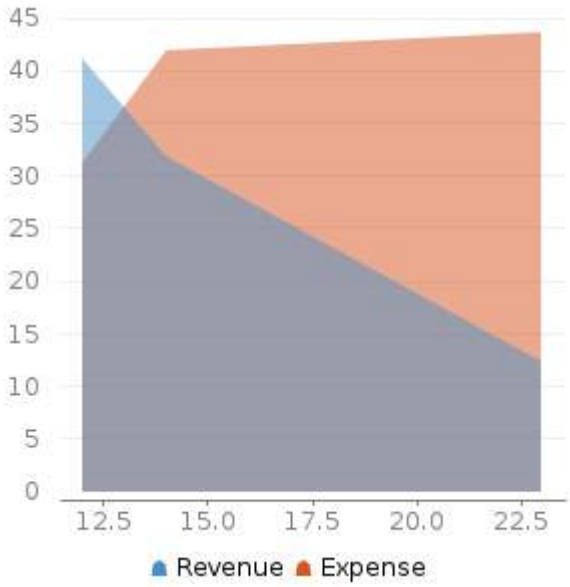

To create this chart, we set these parameters in the macro browser:

- **Type**: xyArea
- **Show legend**: true

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

and added this table in the macro body:

	12	14	23
Revenue	41.1	31.8	12.4
Expense	31.1	41.8	43.6

Area chart

Here's an example of an area chart.

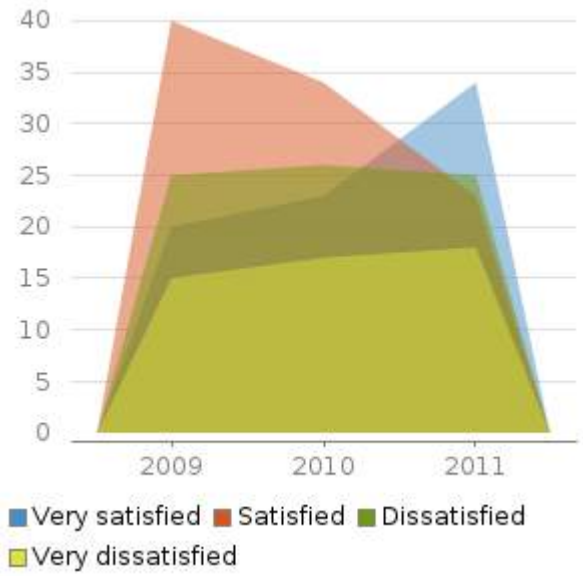

To create this chart, we set these parameters in the macro browser:

- **Type**: area
- **Show legend**: true
- **Width**: 300
- **Height**: 300
- **Opacity**: 50

and added this table in the macro body:

Satisfaction	2009	2010	2011
Very satisfied	20	23	34
Satisfied	40	34	23
Dissatisfied	25	26	25
Very dissatisfied	15	17	18

Stacked area chart

Here's another example of an Area chart.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

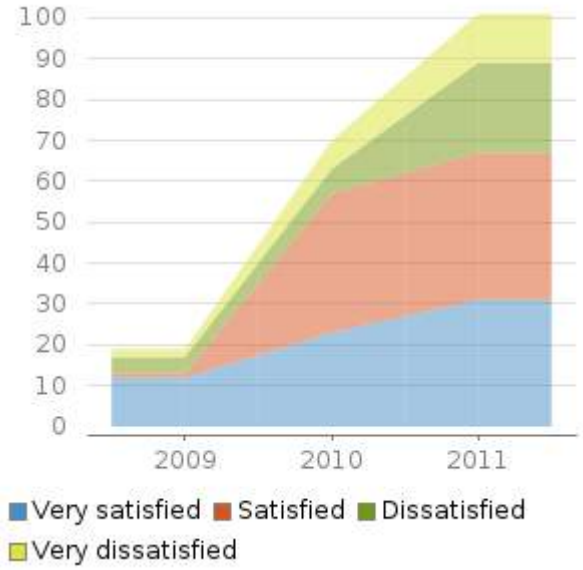

To create this chart, we set these parameters in the macro browser:

- **Type**: area
- **Show legend**: true
- **Width**: 300
- **Height**: 300
- **Stacked values**: true

and added this table in the macro body:

Satisfaction	2009	2010	2011
Very satisfied	12	23	31
Satisfied	1	34	36
Dissatisfied	4	6	22
Very dissatisfied	2	7	12

Gantt chart

Here's an example of a Gantt chart.

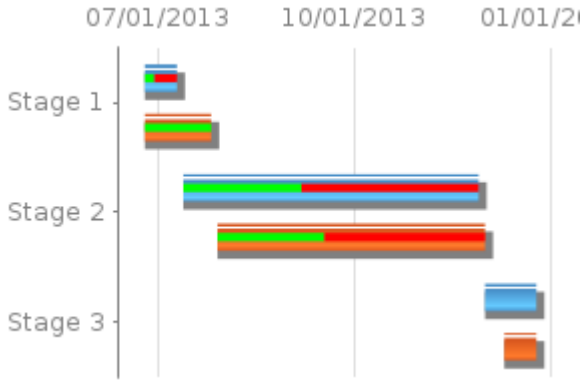

To create this chart, we set these parameters in the macro browser:

- **Type**: gantt
- **Width**: 300

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

- **Height**: 200
- **Columns**:,,1,2,3,4 (note the two commas to start)
- **Date format**:MM/dd/yyyy

and added these two tables in the macro body:

Plan	Start	End	Status
Stage 1	6/25/2013	7/10/2013	30%
Stage 2	7/13/2013	11/28/2013	40%
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	

You must include the two leading commas in the column parameter (for example `,,1,2,3,4`)for the chart to be created correctly.

Other ways to add this macro

Add this macro as you type

Type**{**followed by the start ofthe macro name, to see a list of macros.

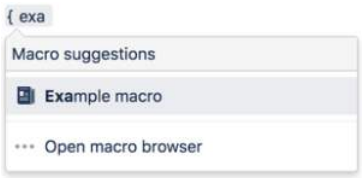

Add this macro using wiki markup

This is useful when you want to add a macro outside the editor, for example as custom content in the sidebar, header or footer of a space.

Macro name:`chart`

Macro body: Accepts rich text, consisting of tables that hold the chart's data.

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}
```

This macro recognizes a large number of parameters, listed here by type.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

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.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

`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-separatedlist 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.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

`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 customization 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. Theseadditional 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 arethe 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 color parameters

Colors are specified usinghexadecimal notation or HTML color names.

Parameter	Required	Default	Description
`bgColor`	No	White	Background color of the chart.
`borderColor`	No	No border	Color of the border around the chart.
`colors`	No		A comma-separated list of colors used to customize the colors of categories, sections, and series.

Chart axis parameters

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Depending on the chart type, the range and domain axis may be customized. 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

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

`pieSectionLabel`	No	Show only the pie section key value	Formatof 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 authorized 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. • `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 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.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

`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.

Do more with Confluence

Create interesting and engaging charts for your Confluence pages with these top charts and diagrams apps on the Atlassian Marketplace.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Wiki Markup Examples for Chart Macro

This page is an extension of the documentation for theChart 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

Unknown macro: 'chart'

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

Unknown macro: '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

Unknown macro: '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

Unknown macro: '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

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Unknown macro: '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

Unknown macro: '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

Unknown macro: 'chart'

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

Unknown macro: 'chart'

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Cheese Macro

This macro is available in **Confluence Server and Data Center**. Learn about the macros available in Confluence Cloud.

On this page:

- Add the Cheese macro to your page
- Change the macro parameters
- Other ways to add this macro

Add a Cheese macro to your page to display the words "I like cheese!".

This macro is great for:

- testing Confluence macro functionality
- sharing your love of cheese.

Seriously though, this macro is just for testing purposes. It's provided by the Basic Macros system plugin, and you can disable the Cheese module if cheese isn't your thing.

Add the Cheese macro to your page

To add the Cheese macro to a page:

1. From the editor toolbar, choose **Insert** > **Other Macros**.
2. Choose**Cheese**from the**Confluence content**category.
3. Choose**Insert**.

You can then publish your page to see the macro in action.

Change the macro parameters

This macro has no parameters.

Other ways to add this macro

Add this macro as you type

Type**{**followed by the start ofthe macro name, to see a list of macros.

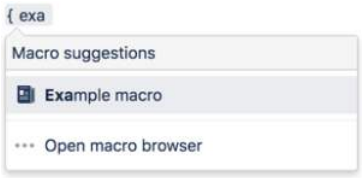

Add this macro using wiki markup

This is useful when you want to add a macro outside the editor, for example as custom content in the sidebar, header or footer of a space.

Macro name:`cheese`

Macro body: None.

```
{cheese}
```

Children Display Macro

This macro is available in **Confluence Server and Data Center**. Learn about the macros available in Confluence Cloud.

On this page:

- Add the Children Display macro to your page
- Change the macro parameters
- Other ways to add this macro

Add the Children Display macro to a page to display a list of pages from a specific part of the page hierarchy. You can choose to display pages that are a child of the current page, or a child of any other page in a space.

This macro is great for providing quick access to:

- pages related to a project
- procedures and how-to pages.

Because it relies on the page hierarchy, the list of pages is automatically updated when pages are added, deleted, or moved. You can even show an excerpt from the page for extra context.

Screenshot: The Children Display macro, showing a list of pages about printers.

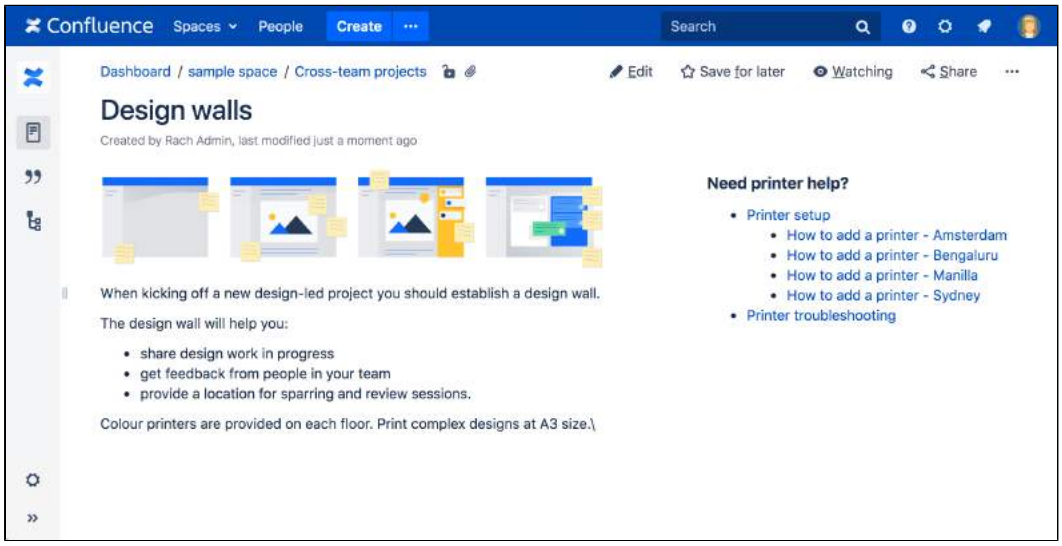

Want to create a table of contents from headings on your page? See Table of Contents Macro

Add the Children Display macro to your page

To add the Children Display macro to a page:

1. From the editor toolbar, choose **Insert** + > **Other Macros**.
2. Choose**Children Display**from the**Confluence content**or **Navigation** category.
3. Use the parameters below to specify which pages to display, and how you want them to look.
4. Choose**Insert**.

You can then publish your page to see the macro in action.

Screenshot:*Specifying the parent page and display options in the Children Display macro.*

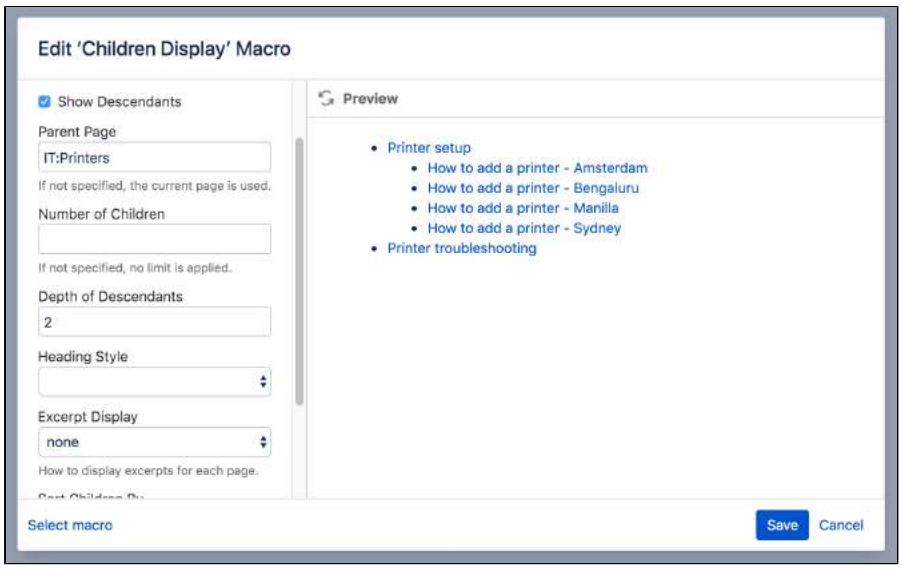

Change the macro parameters

Macro parameters are used to change the behaviour of a macro.

To change the macro parameters:

1. In the editor, click the macro placeholder and choose**Edit**.

2. Update the parameters as required then choose**Insert**.

Here's a list of the parameters available in this macro.

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. Your administrator may enforce a maximum number for this parameter.
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 if **Show Descendants** is enabled. Your administrator may enforce a maximum number for this parameter.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

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 on 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. Note that your administrator may have set this parameter to `none` for it to apply to the whole site. The JVM paramater `confluence.child-macro.disable-excerpt` can be set to "true" to disable excerpts in children display macros. See Recognized System Properties for more details on this parameter and when it can be used.
Sort Children By (`sort`)	Manual if manually ordered, otherwise alphabetical	Leave blank to display pages in the order they currently appear in the page tree. Alternatively, 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.

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`).

Other ways to add this macro

Add this macro as you type

Type **{** followed by the start ofthe macro name, to see a list of macros.

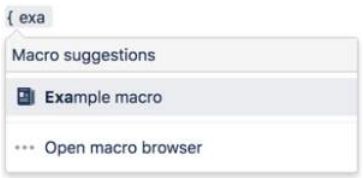

Add this macro using wiki markup

This is useful when you want to add a macro outside the editor, for example as custom content in the sidebar, header or footer of a space.

Macro name: `children`

Macro body: None.

```
{children:reverse=true|sort=creation|style=h4|page=Home|excerpt=none|first=99|depth=2|all=true}
```


Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Do more with Confluence

Extend Confluence with one of the hundreds of other macros in the Atlassian Marketplace. Here are a couple for organizing your Confluence page:

- Navitabs: Create tabs to group content to improve navigation between Confluence pages
- Advanced Children Display for Confluence: combine Confluence's built-in children display and table of contents macros
- Subspace Navigation for Confluence: Use the navigation macro to create overviews of the menu within a Confluence page

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Code Block Macro

This macro is available in **Confluence Server and Data Center**. Learn about the macros available in Confluence Cloud.

On this page:

- Add the Code Block macro to your page
- Change the macro parameters
- Administer the Code Block macro
- Other ways to add this macro

Add a Code Block macro to your page to display code examples with syntax highlighting.

This is great for sharing code snippets such as:

- sample code
- terminal commands
- excerpts from application logs.

Screenshot: code sample in the Code Block macro, with syntax highlighting and a dark theme.

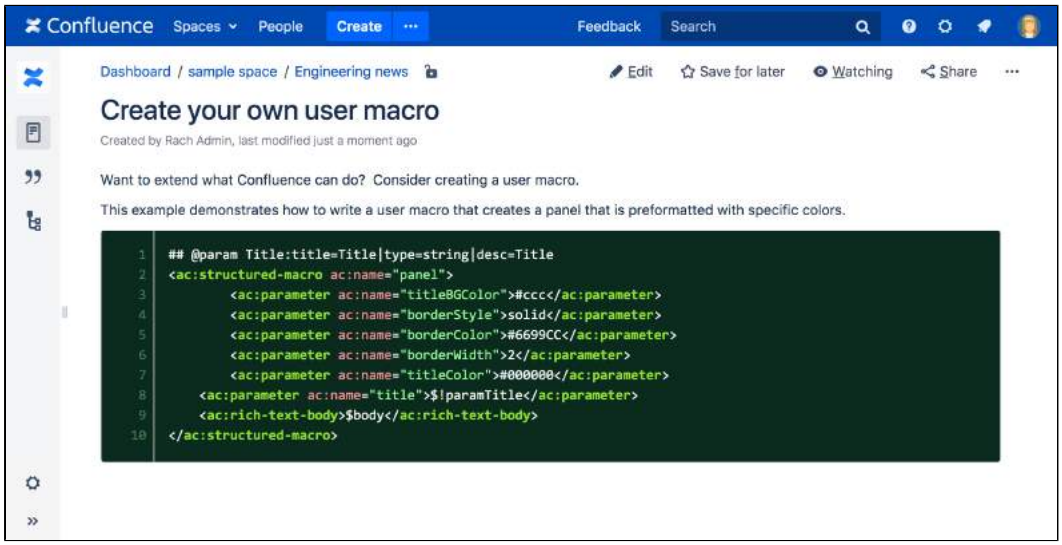

Add the Code Block macro to your page

To add the Code Block macro to a page:

1. From the editor toolbar, choose **Insert** > **Other Macros**.
2. Choose **Code Block**from the**Formatting**category.
3. Choose a language for syntax highlighting.
4. Use the parameters below to customise how the code block should appear on your page.
5. Choose**Insert**.
6. Type or paste your code into the macro placeholder.

You can then publish your page to see the macro in action.

Screenshot: Choosing syntax highlighting language andtheme in the Code Block macro

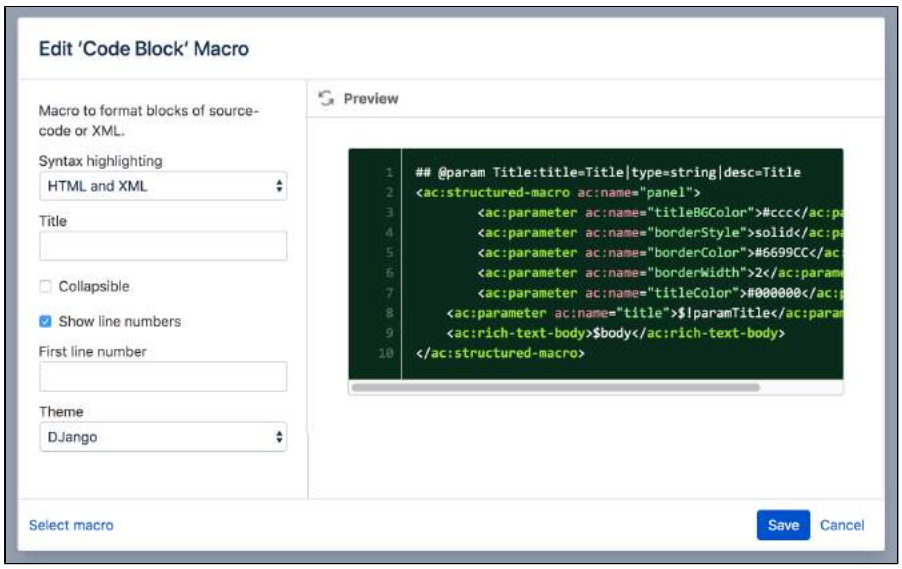

Note: You type the code block directly into the macro placeholder in the editor. Note that any white space contained in the placeholder *is not* manipulated in any way by the Code Block macro. This is to provide the writer with flexibility over code indentation.

Change the macro parameters

Macro parameters are used to change the behaviour of a macro.

To change the macro parameters:

1. In the editor, click the macro placeholder and choose**Edit**.
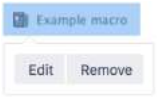

2. Update the parameters as required then choose**Insert**.

Here's a list of the parameters available in this macro.

Parameter	Default	Description

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

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: • **ActionScript** • **AppleScript** • **Bash** • **C#** • **C++** • **CSS** • **ColdFusion** • **Delphi** • **Diff** • **Erlang** • **Groovy** • **HTML and XML** • **Java** • **Java FX** • **JavaScript** • **PHP** • **Plain Text** • **PowerShell** • **Python** • **Ruby** • **SQL** • **Sass** • **Scala** • **Visual Basic** • **YAML**
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.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Theme	Default	Specifies the color scheme used for displaying your code block. Many of these themes are based on the default color schemes of popular integrated development environments (IDEs). The default theme is **Confluence** (also known as **Default**), which is typically black and colored 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**

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`).

Administer the Code Block macro

You can configure the Code Block macro to use a specific language and theme by default and also upload new languages.You need Confluence Administratorpermissions to change the default theme and language andSystem Administratorpermissions to upload new languages.

To set the default appearance of code blocks in your site:

1. Go to > **General Configuration**>**Configure Code Macro**.
2. Select a**Default Theme** and**Default Language**.
3. 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. Go to > **General Configuration**>**Configure Code Macro**.
2. Choose**Add a new language**.
3. Locate your language file and enter a**Name**for the new language (this will appear when selecting the language).
4. Choose**Add**.

Language files must be correctly formatted JavaScript files and adhere to theCustom Brush syntax. You can see examples in`<install-directory>/confluence/includes/js/third-party`.

To disable or remove a user-installed language:

1. Go to > **Manage apps**.
2. Go to**User-installed Apps**and locate the app for your uploaded language - it will appear like this 'Custom Code Macro Highlighting for...'
3. Choose**Uninstall**or**Disable**.

The language will no longer appear in the Code Macro.

Other ways to add this macro

Add this macro as you type

Type**{**followed by the start ofthe macro name, to see a list of macros.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

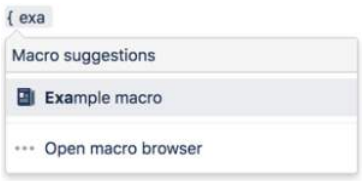

Add this macro using wiki markup

This is useful when you want to add a macro outside the editor, for example as custom content in the sidebar, header or footer of a space.

Macro name:code

Macro body: Accepts plain text.

```
{code:title=This is my
title|theme=FadeToGrey|linenumbers=true|language=java|firstline=0001|collapse=true}
This is my code
{code}
```

Do more with Confluence

Extend Confluence with one of the hundreds of other macros in theAtlassian Marketplace. Some of our most popular include:

- Code Pro for Confluence- Get a real-time view of your code from any source in Confluence.
- Include Bitbucket Server for Confluence- Easily include code snippets in Confluence that sync automatically to Bitbucket Server.
- Markdown Extension for Confluence- Embed markdown from private and public Github & Bitbucket repositories in Confluence.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Column Macro

Add the column macro to a page to organise your content in columns. This macro is used in conjunction with the Section macro, and provides more flexibility than page layouts.

This macro is great for situations where:

- you need more than three columns, or
- you need your columns to be a specific width.

On this page:

- Add this macro to your page
- Change the macro parameters
- Other ways to add this macro

Screenshot: page with a four column layout using the Section and Column macros.

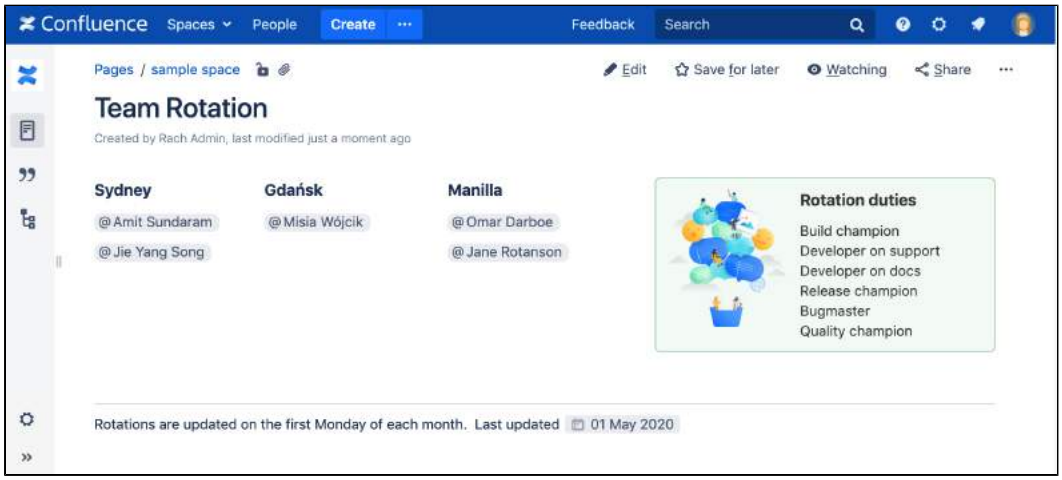

Want a simpler way to lay out your page? Try a page layout instead.

Add this macro to your page

Column macros must be added inside a page layout section, or within a Section macro.

To add the Column macro to a page:

1. Position your cursor inside the body of a Section macro, or page layout section.
2. From the editor toolbar, choose **Insert** + > **Other Macros**.
3. Choose**Column**from the**Formatting**category.
4. Choose**Insert**.

You can then start typing into the macro body, then publish your page to see the macro in action.

Screenshot: section and column macros in the editor

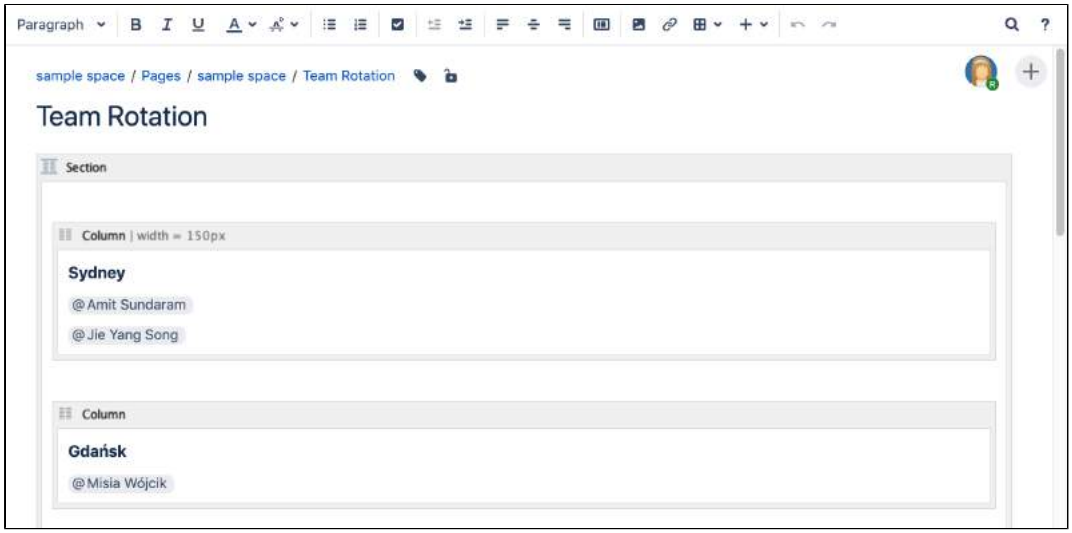

Change the macro parameters

Macro parameters are used to change the behaviour of a macro.

To change the macro parameters:

1. In the editor, click the macro placeholder and choose**Edit**.

2. Update the parameters as required then choose**Insert**.

Here's a list of the parameters available in this macro.

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%`).

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`).

Other ways to add this macro

Add this macro as you type

Type**{**followed by the start ofthe macro name, to see a list of macros.

Add this macro using wiki markup

This is useful when you want to add a macro outside the editor, for example as custom content in the sidebar, header or footer of a space.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Macro name: column

Macro body: Accepts rich text.

```
{column:width=100px}
This is the content of *column 1*.
{column}
```


Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Content by Label Macro

Add a Content by Label macro to your page todisplay a list of pages, blog posts or attachments that have particular labels.

This macro is great forfor collecting related pages together and filtering out content that you don't want to see. You could:

- list of all pages that have the label 'feature-shipped' and include the word 'Blueprint'
- list any pageswith the label 'meeting-notes' that you've been mentioned in.

On this page:

- Add this macro to your page
- CQL filters
- Change the macro parameters
- Other ways to add this macro

Screenshot: The Content by Label macro showing all pages that contain the label "printer".

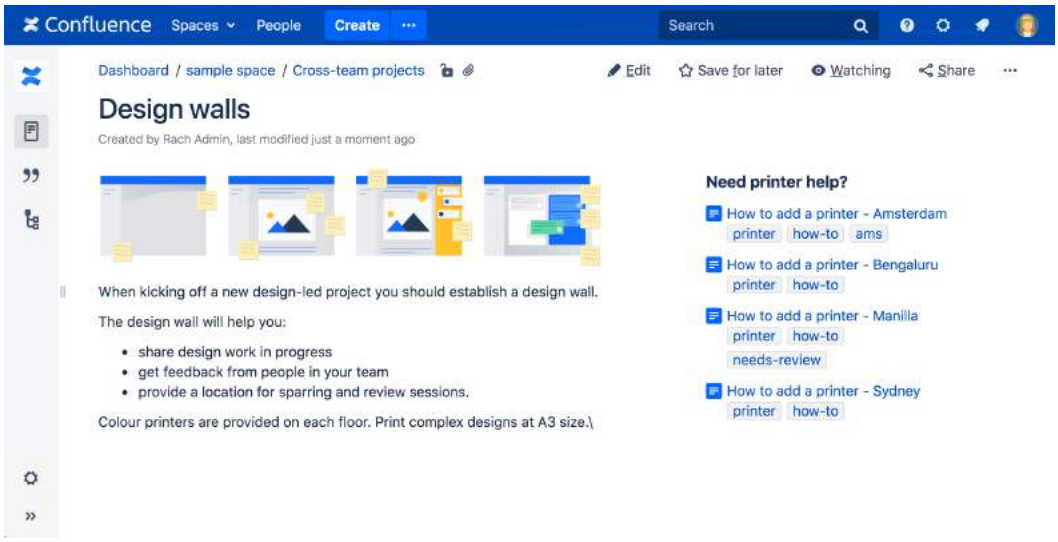

For general information about using labels in Confluence, see Add, Remove and Search for Labels.

Add this macro to your page

To add the Content by Label macro to a page:

1. From the editor toolbar, choose **Insert** + > **Other Macros**.
2. Choose**Content by Label**from the**Confluence content**category.
3. Enter the labels you want to use as the basis for your query.
4. Add additional filters to further narrow your query. These filters use CQL.
5. Choose Show to change the macro parameters. These are optional.
6. Choose**Insert**.

You can then publish your page to see the macro in action.

Screenshot: Using CQL to search for content with a particular label in two spaces in the Content by Label macro

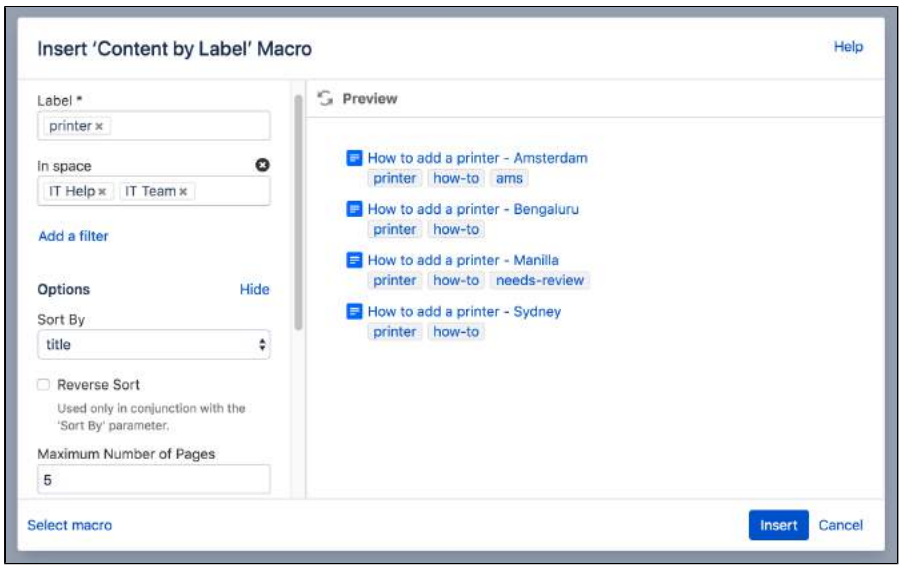

CQL filters

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.

- 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

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

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.

Change the macro parameters

Macro parameters are used to change the behaviour of a macro.

To change the macro parameters:

1. In the editor, click the macro placeholder and choose**Edit**.
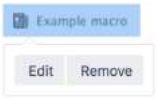

2. Update the parameters as required then choose**Insert**.

Here's a list of the parameters available in this macro.

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. This can be any value up to 500 pages.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

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.
Exclude current page	False	Allows you to exclude the page the macro appears on from the list. This is useful when the current page contains the same labels as the pages you want to include in the list.

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`).

Other ways to add this macro

Add this macro as you type

Type **{** followed by the start ofthe macro name, to see a list of macros.

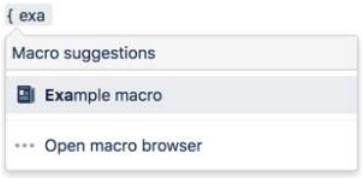

Add this macro using wiki markup

You can't use wiki markup to add this macro.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Content by User Macro

Add the Content by User macro to a page to display a list of all the things created by a particular user including spaces, pages, blog posts, comments, attached files, and even user accounts.

This is a **legacy macro**, and doesn't provide any way to limit the amount of information displayed. This means it can cause timeouts if the user specified has created a lot of content.

Advanced search is a much better way to find all content created by a specific person.

On this page:

- Add this macro to your page
- Change the macro parameters
- Other ways to add this macro

Screenshot: The Content by User macro displaying everything created by a particular user.

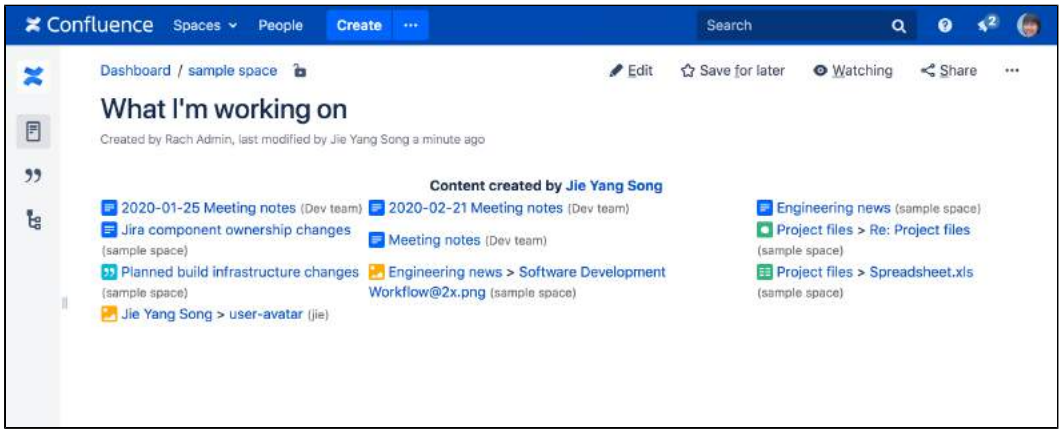

Add this macro to your page

To add the Content by User macro to a page:

1. From the editor toolbar, choose **Insert** > **Other Macros**.
2. Choose**Content by User**from the**Confluence content**category.
3. Enter the username of the person you want to activity for.
4. Choose**Insert**.

You can then publish your page to see the macro in action.

Screenshot: specifying a user in the Content by User macro

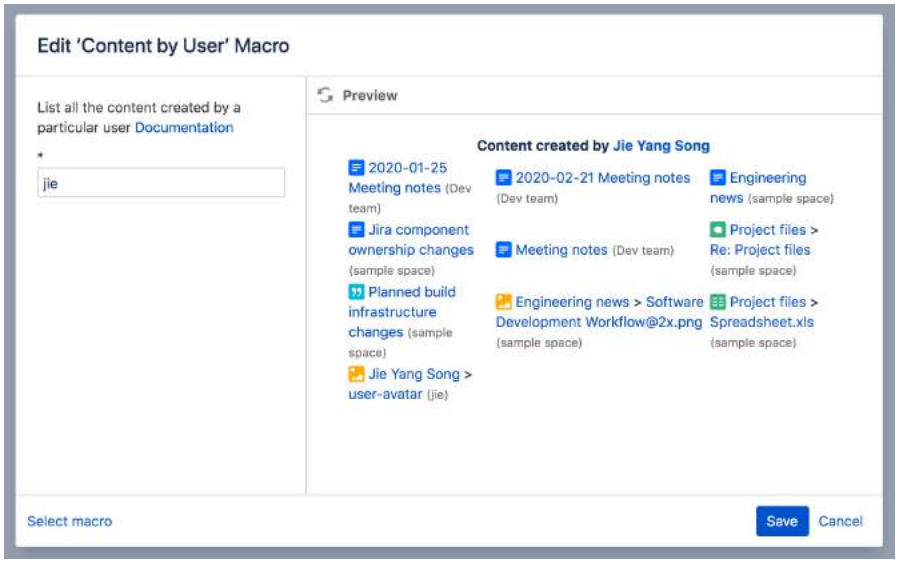

Change the macro parameters

Macro parameters are used to change the behaviour of a macro.

To change the macro parameters:

1. In the editor, click the macro placeholder and choose**Edit**.

2. Update the parameters as required then choose**Insert**.

Here's a list of the parameters available in this macro.

Parameter	Required	Default	Description
Username **(name)**	yes	none	The Confluence username for a person who has created content. This parameter is unnamed in wiki markup.

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`).

Other ways to add this macro

Add this macro as you type

Type**{**followed by the start ofthe macro name, to see a list of macros.

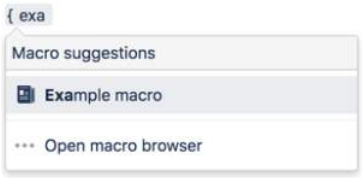

Add this macro using wiki markup

This is useful when you want to add a macro outside the editor, for example as custom content in the sidebar, header or footer of a space.

Macro name:`content-by-user`

Macro body: None.

```
{content-by-user:jsmith}
```


Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Content Report Table Macro

Add the Content Report Table to a page todisplay a table of pages and blog posts with a specific label along with the creator and modified date.

This macro is great for situations where you need to see information about a set of pages at a glance, such as:

- project pages
- document control
- change management
- process and procedure documentation.

On this page:

- Add this macro to your page
- Change the macro parameters
- Other ways to add this macro

You can see this macro in action in the Meeting Notes blueprint.

Screenshot: The Content Report Table macro, configured to show pages with the label "printer".

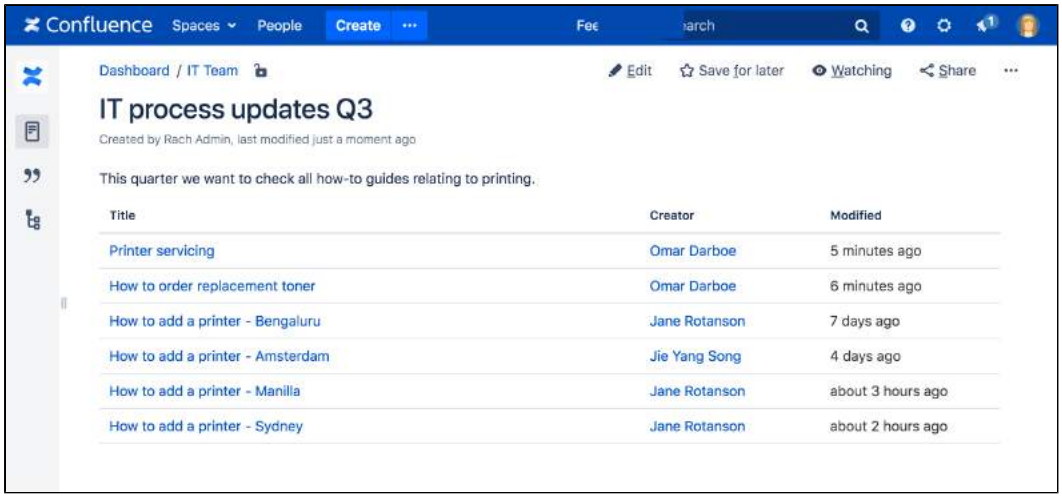

For general information about using labels in Confluence, see Add, Remove and Search for Labels.

Add this macro to your page

To add the Content Report Table macro to a page:

1. From the editor toolbar, choose **Insert** > **Other Macros**.
2. Choose**Content Report Table**from the**Confluence content**category.
3. Enter the labels you want to display pages for.
4. Use the parameters below to narrow your query.
5. Choose**Insert**.

You can then publish your page to see the macro in action.

Screenshot: Entering the label parameter in the Content Report Table macro.

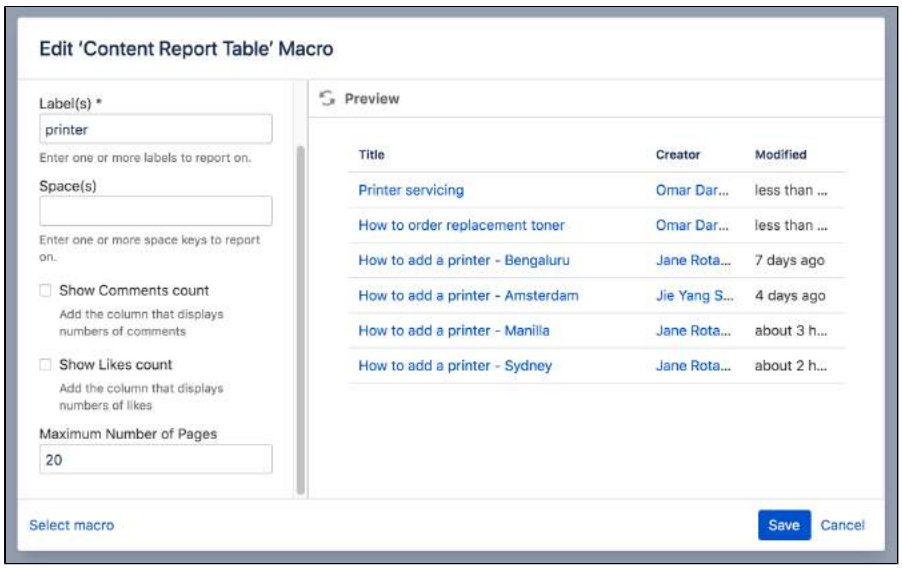

Change the macro parameters

Macro parameters are used to change the behaviour of a macro.

To change the macro parameters:

1. In the editor, click the macro placeholder and choose**Edit**.

2. Update the parameters as required then choose**Insert**.

Here's a list of the parameters available in this macro.

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 labeled '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.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

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`).

Other ways to add this macro

Add this macro as you type

Type **{** followed by the start ofthe macro name, to see a list of macros.

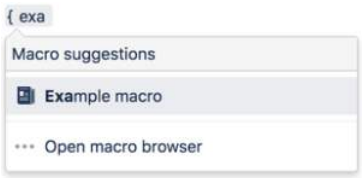

Add this macro using wiki markup

You can't use wiki markup to add this macro.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Contributors Macro

Add the Contributors macro to a page to display a list of Confluence users who have contributed to this page, another page, or set of pages. Contributors includes people who:

- created or edited the pages
- commented on the pages
- added labels to the pages, or
- are simply watching the pages.

This macro is great for:

- document control and compliance
- providing a list of people to contact about a page.

On this page:

- Add this macro to your page
- Change the macro parameters
- Other ways to add this macro

Screenshot: Page showing contributors to the printer how-to pages.

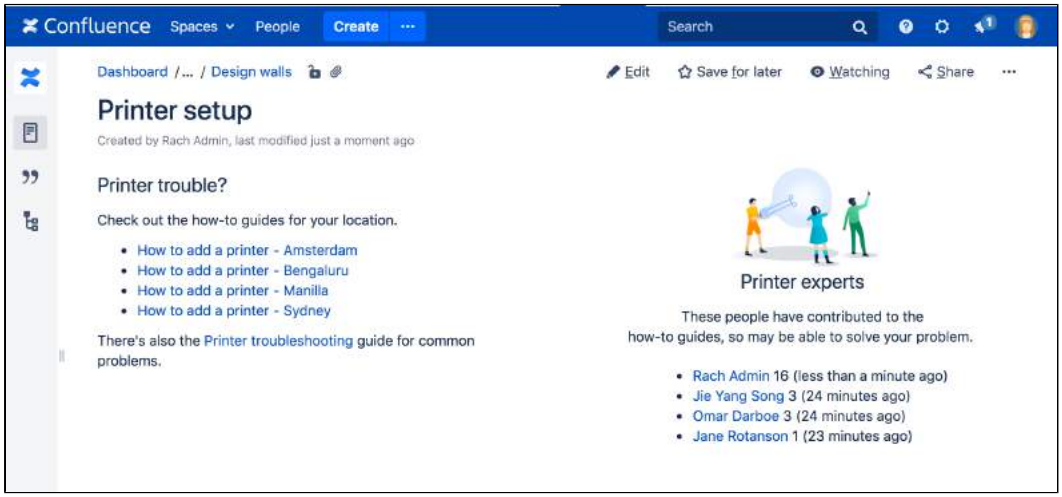

In this example, the**Display Format**parameter has been set to**list**.

Add this macro to your page

To add the Contributors macro to a page:

1. From the editor toolbar, choose **Insert** + > **Other Macros**.
2. Choose**Contributors**from the**Confluence content**category.
3. Enter the type of contributor you want to display, and set any parameters.
4. Choose**Insert**.

You can then publish your page to see the macro in action.

Screenshot: Setting parameters in the Contributors macro

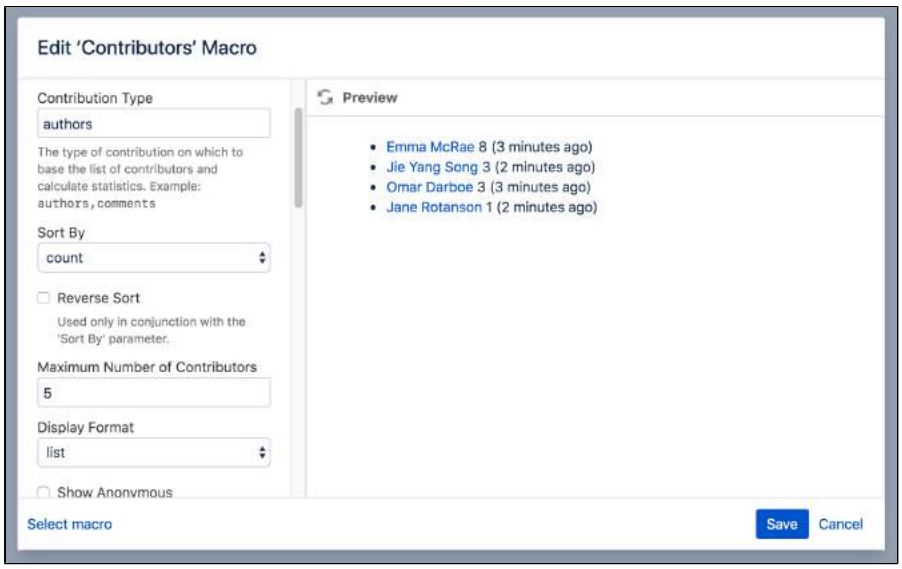

Change the macro parameters

Macro parameters are used to change the behaviour of a macro.

To change the macro parameters:

1. In the editor, click the macro placeholder and choose**Edit**.

2. Update the parameters as required then choose**Insert**.

Here's a list of the parameters available in this macro.

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.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

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.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

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.

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`).

Other ways to add this macro

Add this macro as you type

Type**{**followed by the start ofthe macro name, to see a list of macros.

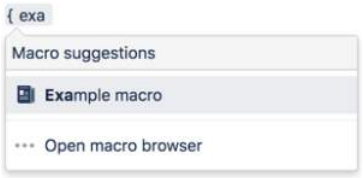

Add this macro using wiki markup

This is useful when you want to add a macro outside the editor, for example as custom content in the sidebar, header or footer of a space.

Macro name:`contributors`

Macro body: None.

This example specifies a content type of blog posts:

```
{contributors:limit=10|spaces=ds,@personal|reverse=true|labels=chocolate,
cake|showPages=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=chocolate,
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}
```


Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Contributors Summary Macro

Add the Contributors macro to a page to display statistics about people contributing to this page, a specific page, or set of pages. Contributors includes people who:

- created or edited the pages
- commented on the pages
- added labels to the pages, or
- are simply watching the pages.

This macro is great for:

- collaboration leaderboards
- recognising frequent documentation contributors
- tracking contributions to high value pages.

On this page:

- Add this macro to your page
- To add the Contributors Summary macro to a page:
- Change the macro parameters
- Other ways to add this macro

Screenshot: Page with two Contributors Summary macros, one displaying contributions by user, the other by page.

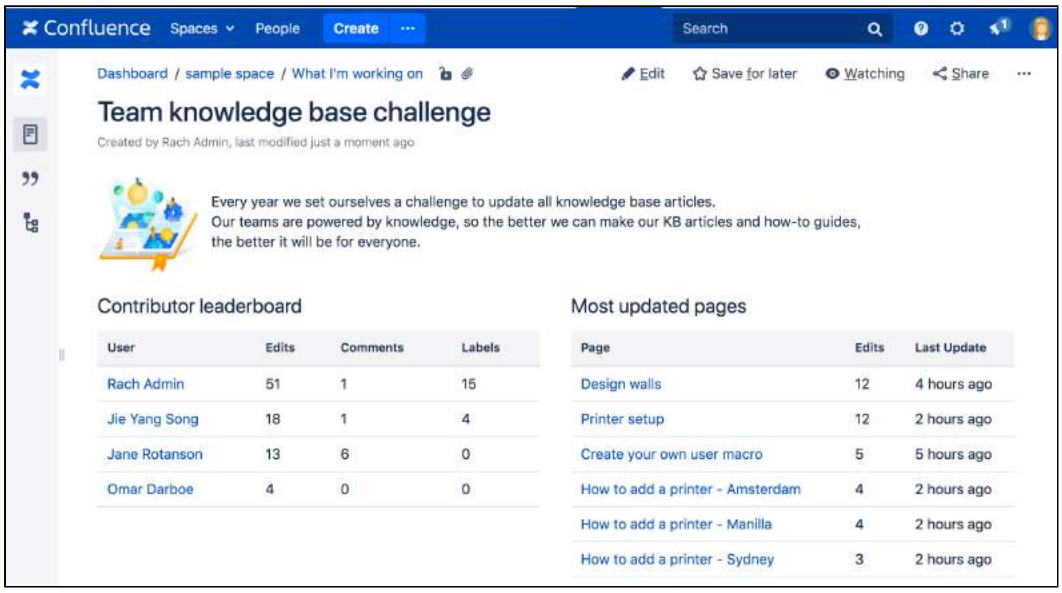

Add this macro to your page

To add the Contributors Summary macro to a page:

1. From the editor toolbar, choose **Insert** + > **Other Macros**.
2. Choose**Contributors Summary**from the**Confluence content**category.
3. Choose whether to show contributions by user or page
4. Set any parameters to refine your query.
5. Choose**Insert**.

You can then publish your page to see the macro in action.

Screenshot: Setting parameters in the Contributors Summary macro

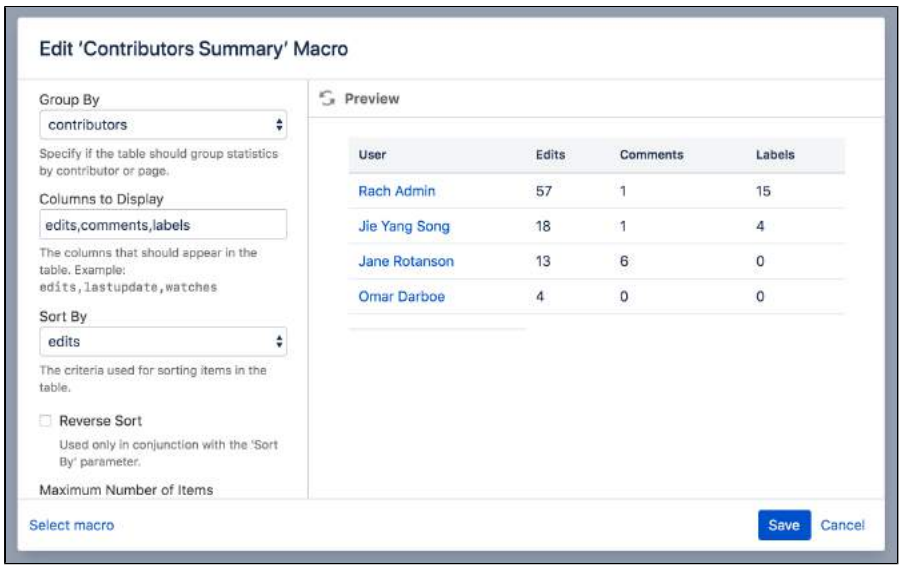

Change the macro parameters

Macro parameters are used to change the behaviour of a macro.

To change the macro parameters:

1. In the editor, click the macro placeholder and choose**Edit**.

2. Update the parameters as required then choose**Insert**.

Here's a list of the parameters available in this macro.

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.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

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.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

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.

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`).

Other ways to add this macro

Add this macro as you type

Type**{**followed by the start ofthe macro name, to see a list of macros.

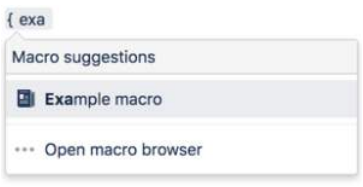

Add this macro using wiki markup

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

This is useful when you want to add a macro outside the editor, for example as custom content in the sidebar, header or footer of a space.

Macro name:`contributors-summary`

Macro body: None.

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}
```


Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Create from Template Macro

Add the Create from Template macro to a page to provide a way to create pages in a specific location, using a particular template.

When someone clicks the button, the macro opens the editor, loads the template or blueprint, and is ready to use. When saved, the page will be a child of the page containing the macro.

This macro is great for:

- project pages
- documentation
- recurring processes.

On this page:

- Add this macro to your page
- Change the macro parameters
- Other ways to add this macro

Screenshot: Page showing the Create from template button, used to create DACI decision pages.

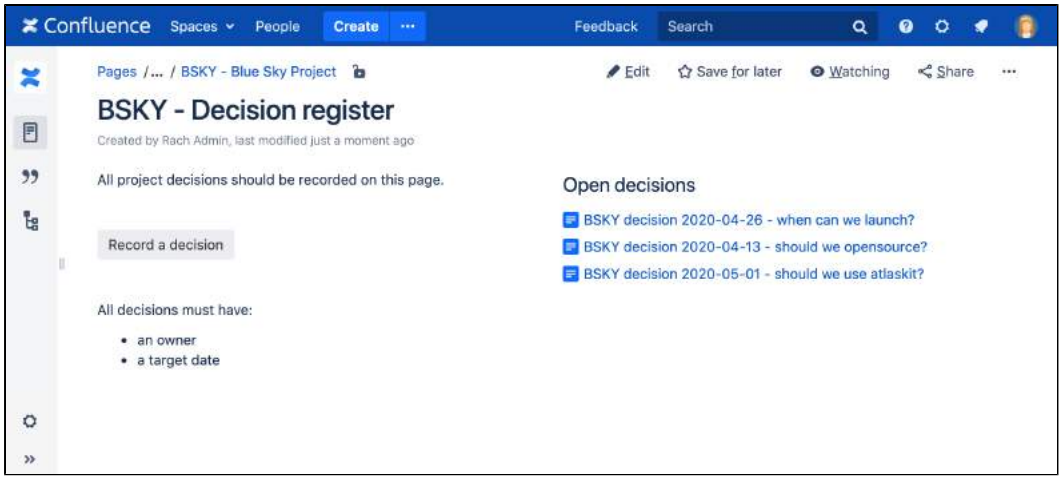

Add this macro to your page

To add the Create from Template macro to a page:

1. From the editor toolbar, choose **Insert** > **Other Macros**.
2. Choose**Create from Template**from the**Confluence content**category.
3. Select a template.
4. Choose**Insert**.

You can then publish your page to see the macro in action.

Screenshot: Configuring the macro to have custom button text, and a default page title which includes the current date.

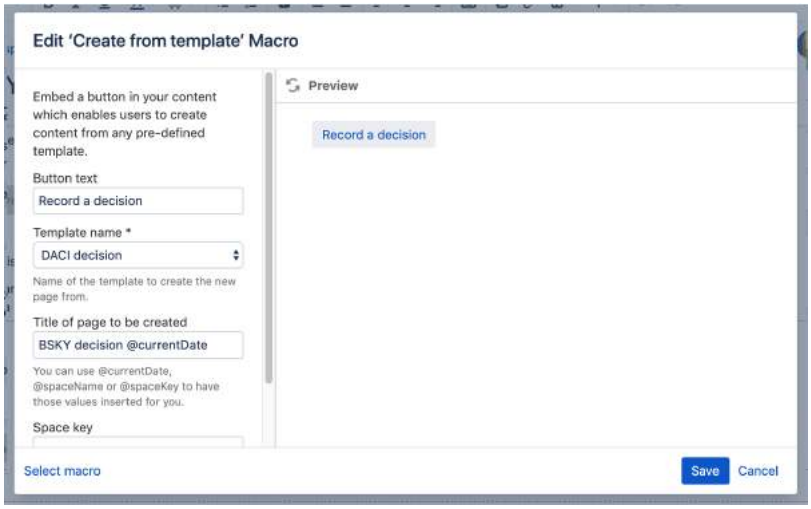

Change the macro parameters

Macro parameters are used to change the behaviour of a macro.

To change the macro parameters:

1. In the editor, click the macro placeholder and choose**Edit**.
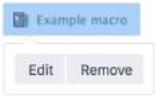

2. Update the parameters as required then choose**Insert**.

Here's a list of the parameters available in this macro.

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. Some blueprints, such as the Decision blueprint, will override your custom 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.

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`).

Other ways to add this macro

Add this macro as you type

Type**{**followed by the start ofthe macro name, to see a list of macros.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

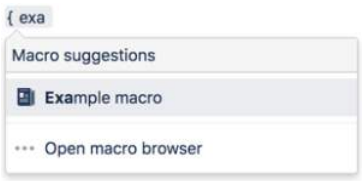

Add this macro using wiki markup

You can't use wiki markup to add this macro.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Create Space Button Macro

Add the Create Space Button macro to a page to provide a quick, visual prompt to create a space in Confluence.The Create Space Button provided by the macro is only be visible to people who have thewho have the Create Spaceglobal permission.

This is a**legacy macro**and was primarily used before you could create a space directly from the Spaces dropdown in the header. However, it might be useful when documenting a procedure that requires creating a space, for example a set-up guide for new teams.

On this page:

- Add this macro to your page
- Change the macro parameters
- Other ways to add this macro

Screenshot: page showing the Create Space Button macro.

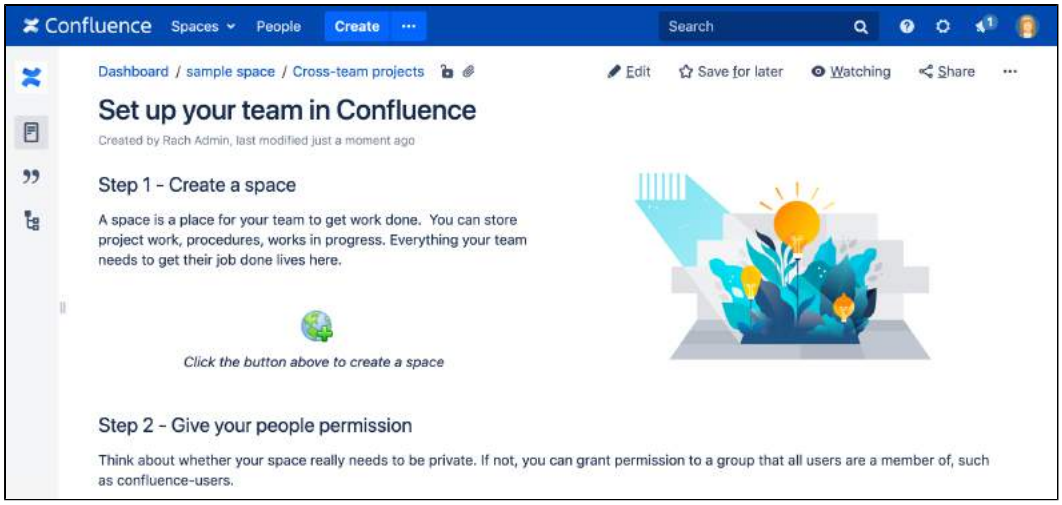

Add this macro to your page

To add the Create Space Button macro to a page:

1. From the editor toolbar, choose **Insert** > **Other Macros**.
2. Choose**Create Space Button**from the**Confluence content**category.
3. Choose**Insert**.

You can then publish your page to see the macro in action.

Change the macro parameters

Macro parameters are used to change the behaviour of a macro.

To change the macro parameters:

1. In the editor, click the macro placeholder and choose**Edit**.

2. Update the parameters as required then choose**Insert**.

Here's a list of the parameters available in this macro.

Parameter	Default	Description

Icon Size `size`	large	Specify whether to use large or small icon. Available values: • **large** • **small**

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`).

Other ways to add this macro

Add this macro as you type

Type**{**followed by the start ofthe macro name, to see a list of macros.

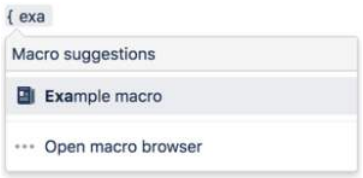

Add this macro using wiki markup

This is useful when you want to add a macro outside the editor, for example as custom content in the sidebar, header or footer of a space.

Macro name:`create-space-button`

Macro body: None.

```
{create-space-button:size=small}
{create-space-button:height=50px|width=50px}
```

The following additional parametersare available in wiki markup.

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.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Excerpt Include Macro

Add the Excerpt Include macro to display the contents of anExcerpt macroo n another page.

This is great single sourcing important information. For example you could provide the contact details and dates from each project page on a summary page. When the info is updated in the excerpt, it will flow though to all the other places it is used.

This is a two-step macro. You need to add theExcerpt macroto another page, before you can use the Excerpt Include macro.

On this page:

- Add this macro to your page
- Change the macro parameters
- Other ways to add this macro

The example below shows an example of an Excerpt Include macro as it appears in the editor, and as it would appear to someone viewing the page. We have set the options to show both the title of the page and the panel surrounding the content.

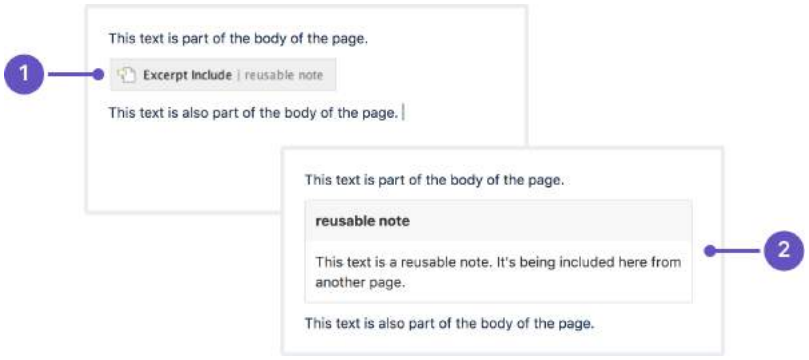

1. The Excerpt Include macro as it appears in the **editor**.
2. The Excerpt Include macro as it appears when **viewing the page**.

The content is being pulled from a Excerpt macro on a page called 'reusable note'.

Add this macro to your page

To add the Excerpt Include macro to a page:

1. From the editor toolbar, choose **Insert** + > **Other Macros**.
2. Choose**Excerpt Include**from the**Confluence content**category.
3. Enter the title of the page containing the Excerpt macro you want to include.
4. Choose**Insert**.

You can then publish your page to see the macro in action.

If the page contains more than one Excerpt macro, the Excerpt Include macro will display the contents of the first one on the page. It can't display content from multiple Excerpt macros on the same page.

> Want to show more than just an excerpt? The Include Page Macro allows you to include the entire contents of another page.

Change the macro parameters

Macro parameters are used to change the behaviour of a macro.

To change the macro parameters:

1. In the editor, click the macro placeholder and choose**Edit**.

2. Update the parameters as required then choose**Insert**.

Here's a list of the parameters available in this macro.

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.

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`).

Other ways to add this macro

Add this macro as you type

Type**{**followed by the start ofthe macro name, to see a list of macros.

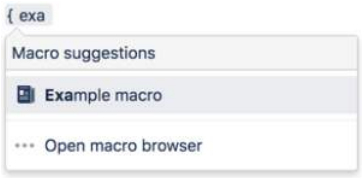

Add this macro using wiki markup

This is useful when you want to add a macro outside the editor, for example as custom content in the sidebar, header or footer of a space.

Macro name:`excerpt-include`

Macro body: None.

```
{excerpt-include:My page name|nopanel=true}
```


Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Excerpt Macro

Add the Excerpt macro to a page to define a snippet of content to be re-used on another page.

This is great single sourcing important information. For example you could provide the contact details and dates from each project page on a summary page. When the info is updated in the excerpt, it will flow though to all the other places it is used.

This is a two-step macro. Once you've added the Excerpt macro to one page, use theExcerpt Include,Children Display,orBlog Postsmacro to display the contents of that excerpt on another page.

On this page:

- Add this macro to your page
- Using this macro
- Change the macro parameters
- Other ways to add this macro

Screenshot: A project summary page containing a Children Display macro, configured to show the contents of the Excerpt macro on each child page.

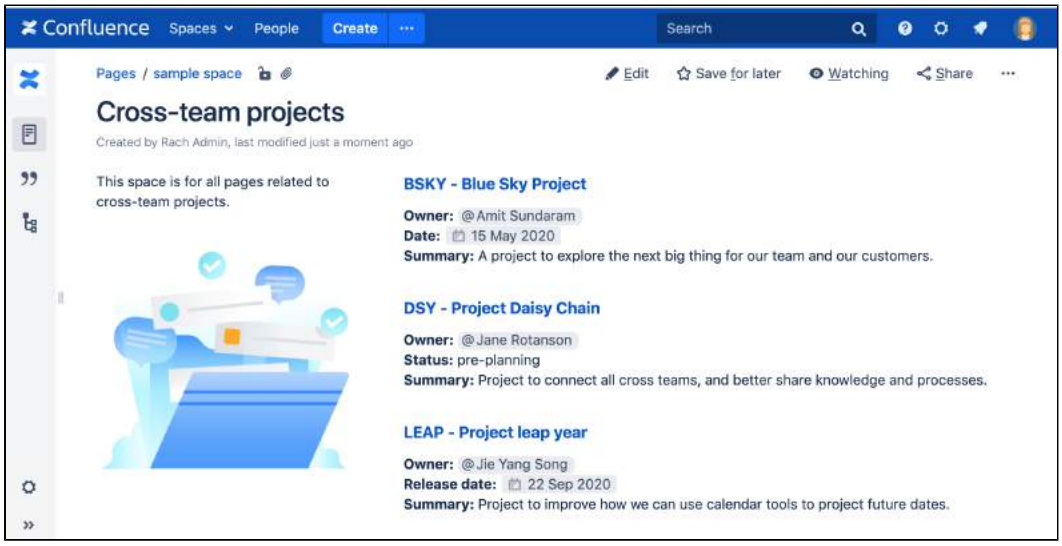

Add this macro to your page

To add the Excerpt macro to a page:

1. From the editor toolbar, choose **Insert** > **Other Macros**.
2. Choose**Excerpt**from the**Confluence content**category.
3. Choose**Insert**.

Enter your text into the macro body. You can then publish your page to see the macro in action.

Screenshot: the Excerpt macro in the editor, containing a short project summary.

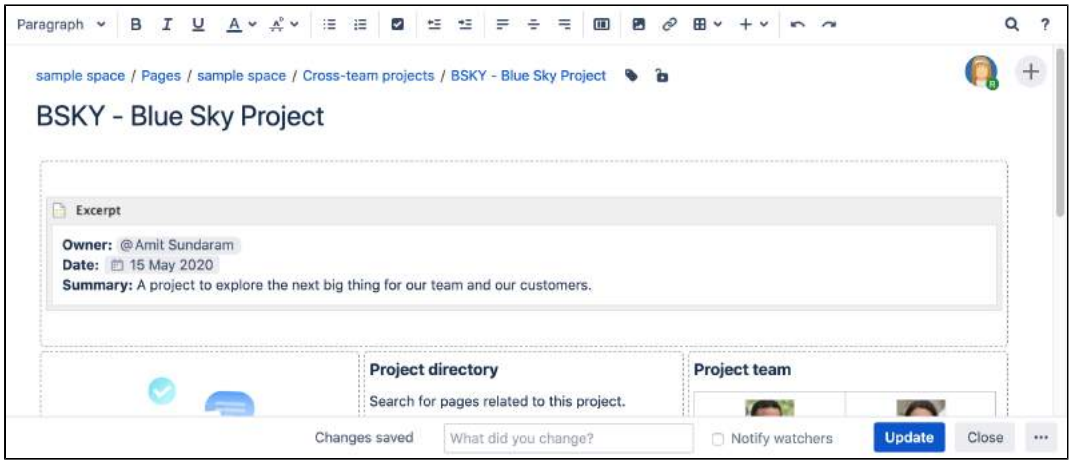

Using this macro

This macro must be used in conjunction with either theExcerpt Include,Children Display,orBlog Postsmacros.

You can only have one excerpt per page. If you add more than one excerpt macro to a page, only the first one will be used by these macros.

You can choose to position the except on its own line, or inline with the surrounding text.This option affects only the page that contains the Excerpt macro, it doesn't affect how the content displays when it is reused on other pages.

Screenshot: The Excerpt macro placeholder and options panel

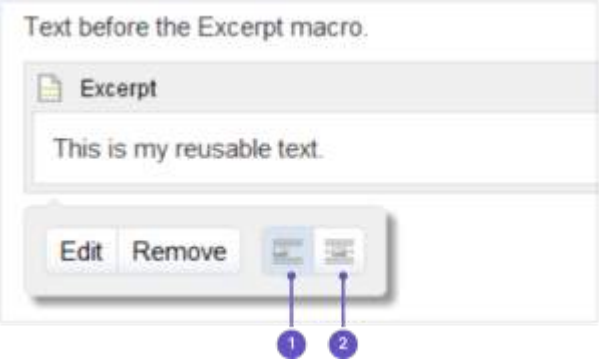

1. **New line**
2. **Inline**

Change the macro parameters

Macro parameters are used to change the behaviour of a macro.

To change the macro parameters:

1. In the editor, click the macro placeholder and choose**Edit**.

2. Update the parameters as required then choose**Insert**.

Here's a list of the parameters available in this macro.

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.

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`).

Other ways to add this macro

Add this macro as you type

Type**{**followed by the start ofthe macro name, to see a list of macros.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

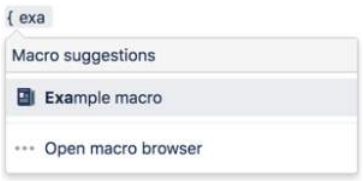

Add this macro using wiki markup

This is useful when you want to add a macro outside the editor, for example as custom content in the sidebar, header or footer of a space.

Macro name:`excerpt`

Macro body: Accepts rich text.

```
{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}
```


Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Expand Macro

This macro is available in **Confluence Server and Data Center**. Learn about the macros available in Confluence Cloud.

On this page:

- Add this macro to your page
- To add the Expand macro to a page:
- Change the macro parameters
- Other ways to add this macro
- Notes

Add the Expand macro to your page to provide content in an expandable / collapsible section.

This is one of Confluence's most popular macros. It's great for:

- visually reducing the amount of information on a page
- breaking process information down intoclickable steps
- hiding background or obsolete information, while still keeping it on the page for future reference.

The macro is collapsed by default, people need to click each one to expand it.There's no way to expand all macros on a page at once, however all Expand macros are automatically expanded when you print or export the page to PDF.

Screenshot: page showing four expand macros. Two are collapsed, and two are expanded.

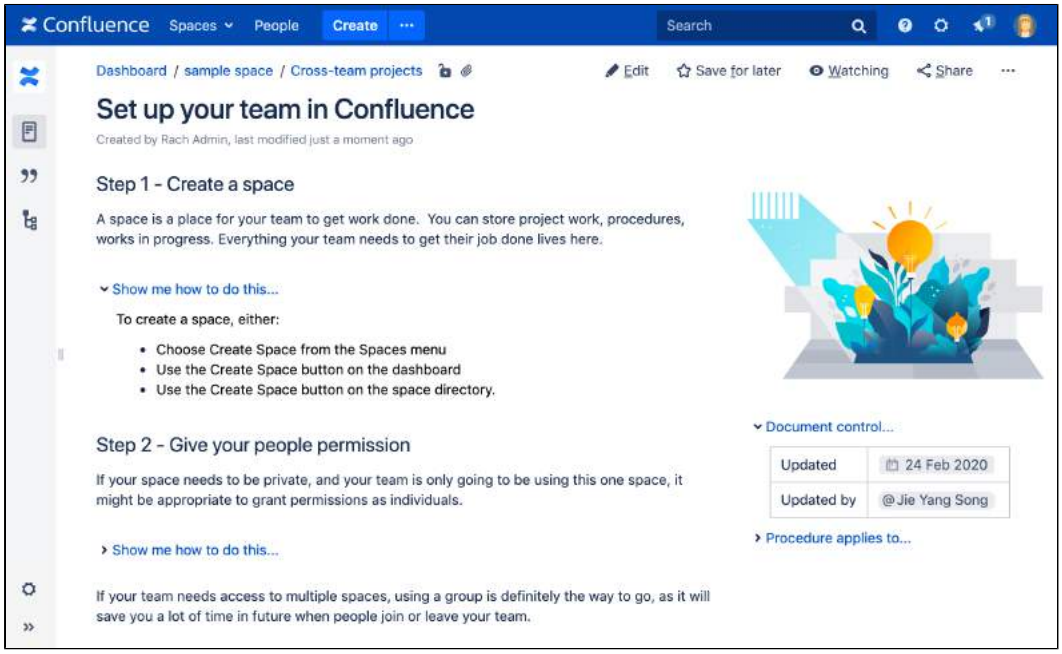

Add this macro to your page

To add the Expand macro to a page:

1. From the editor toolbar, choose **Insert** > **Other Macros**.
2. Choose**Expand**from the**Formatting**category.
3. Enter a **Title**. This is the text a user will click on to show the hidden content.
4. Type or paste your text into the body of the macro. This content will be visible when someone clicks the macro title.
5. Choose**Insert**.

You can then publish your page to see the macro in action.

Screenshot: The expand macro in the editor

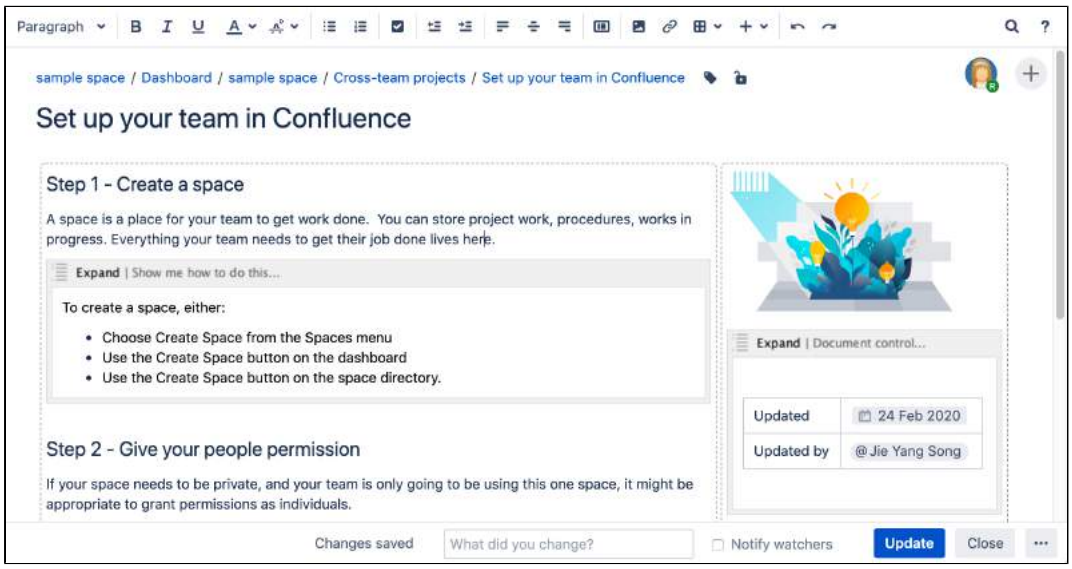

Change the macro parameters

Macro parameters are used to change the behaviour of a macro.

To change the macro parameters:

1. In the editor, click the macro placeholder and choose**Edit**.
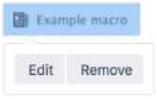

2. Update the parameters as required then choose**Insert**.

Here's a list of the parameters available in this macro.

Parameter	Default	Description
Title	Click here to expand...	Defines the text that appears next to the expand/collapse icon.

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`).

Other ways to add this macro

Add this macro as you type

Type**{**followed by the start ofthe macro name, to see a list of macros.

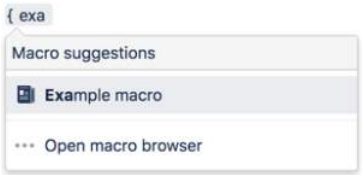

Add this macro using wiki markup

This is useful when you want to add a macro outside the editor, for example as custom content in the sidebar, header or footer of a space.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Macro name:expand

Macro body: Accepts rich text.

```
{expand}
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.
- **Using the Confluence Cloud editor?**This macro may not be available in the new editor. SeeExpand macro for more information.

Do more with Confluence

Extend Confluence with one of the hundreds of other macros in theAtlassian Marketplace, such as:

- Composition Tabs & Page Layouts: Toggle or expand the visible of portions of your pages with the Toggle and Cloak macro
- Content Formatting for Confluence: Over 30 easy-to-use Confluence macros gives you the ability to create better, more engaging content

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Favorite Pages Macro

Add the Favorite Pages macro to a page to display a list of all the pages you'vesaved for later.

This is a **legacy macro** and was primarily used in personal spaces, before you could access starred pages from the dashboard or profile menu. However, it can be a useful way to share pages you've saved for later with other people.

On this page:

- Add this macro to your page
- Change the macro parameters
- Other ways to add this macro

Screenshot: A page with the Favorite pages macro to help new people in a team.

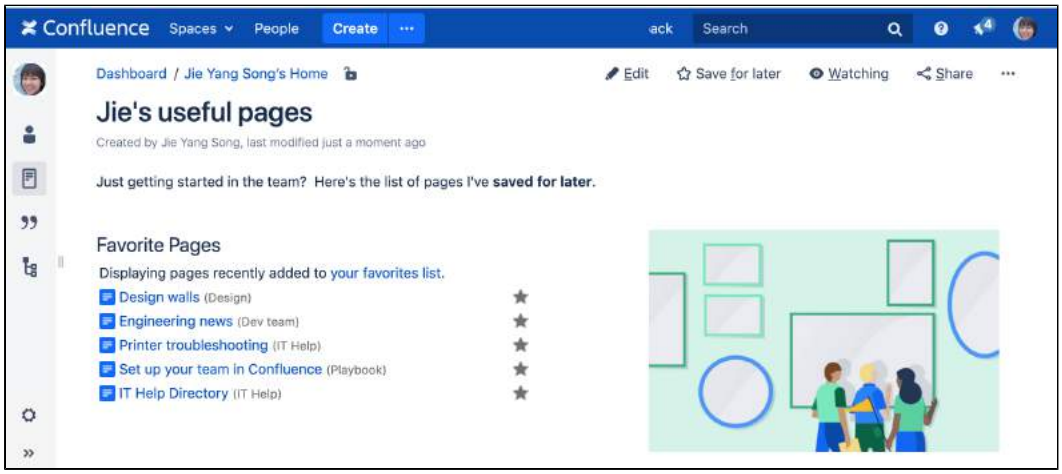

Add this macro to your page

To add the Favorite Pages macro to a page:

1. From the editor toolbar, choose **Insert** + > **Other Macros**.
2. Choose **Favorite Pages**from the**Confluence content**category.
3. Choose**Insert**.

You can then publish your page to see the macro in action.

Change the macro parameters

This macro has no parameters.

Other ways to add this macro

Add this macro as you type

Type**{**followed by the start ofthe macro name, to see a list of macros.

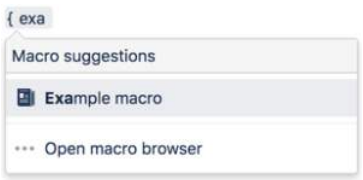

Add this macro using wiki markup

This is useful when you want to add a macro outside the editor, for example as custom content in the sidebar, header or footer of a space.

Macro name: `favpages`

Macro body: None.

```
{favpages}
```


Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Gallery Macro

Add the Gallery macro to a page to display a collection of images attached to this page, or other Confluence pages.

This macro is great for situations where you need to:

- curate a collection of images that are attached to several different pages
- display a set of images in a particular order, such as by date.

One benefit of the Gallery macro over simply inserting images into a page is that it will automatically update when files are added or removed from the page.

On this page:

- Add this macro to your page
- Change the macro parameters
- Image file formats
- Other ways to add this macro

Screenshot: Gallery macro on a page showing images from this page, with captions.

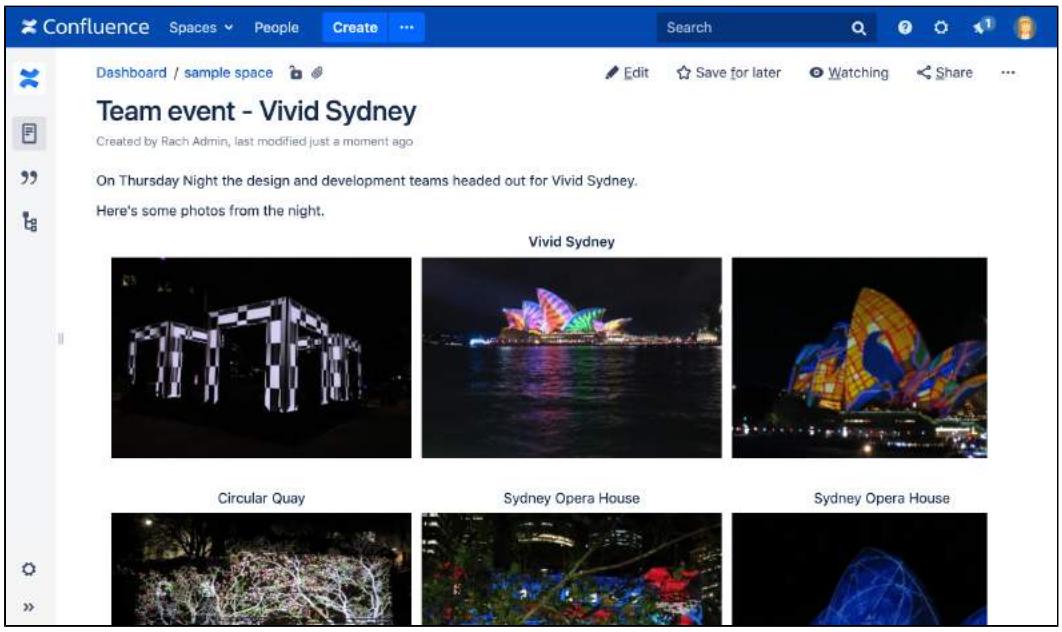

The captions below the images are drawn from the comments on the attachments. For information about adding comments to attachments, seeUpload Files.

For general information about working with files in Confluence, seeFiles

Add this macro to your page

To add the Gallery macro to a page:

1. From the editor toolbar, choose **Insert** > **Other Macros**.
2. Choose**Gallery**from the**Media** or **Visuals & Imagery**categories.
3. Use the macro parameters below to specify the files you want to display.
4. Choose**Insert**.

You can then publish your page to see the macro in action.

Screenshot: Entering parameters for the Gallery macro

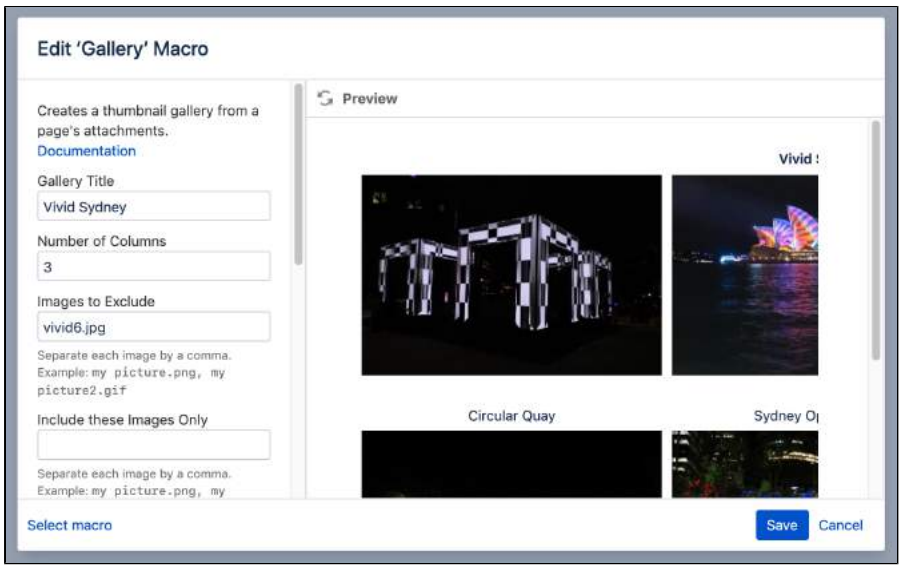

Change the macro parameters

Macro parameters are used to change the behaviour of a macro.

To change the macro parameters:

1. In the editor, click the macro placeholder and choose**Edit**.
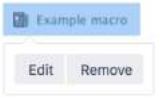

2. Update the parameters as required then choose**Insert**.

Here's a list of the parameters available in this macro.

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 recognized 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 recognized as 'my picture.png'.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Exclude Images with these Labels (`exclude Label`)	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 labeling the attachments, see Add, Remove and Search for Labels.
Include Images with these Labels Only (`include Label`)	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 labeling 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.

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`).

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`

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)

Other ways to add this macro

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Add this macro as you type

Type **{** followed by the start ofthe macro name, to see a list of macros.

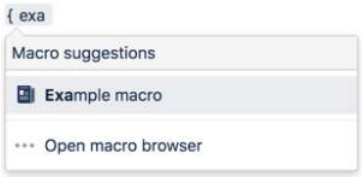

Add this macro using wiki markup

This is useful when you want to add a macro outside the editor, for example as custom content in the sidebar, header or footer of a space.

Macro name: `gallery`

Macro body: None.

```
{gallery:title=My holiday pictures|reverse=true|sort=size|page=My page1, ds:Welcome to
Confluence|excludeLabel=badlabel1, badlabel2|columns=3|exclude=badpicture.png}
```


Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Global Reports Macro

Add the Global Reports macro to a page for access to several feeds that can help you stay on top ofnew content in your site.

This is a**legacy macro**and is no longer fully functional. However, the link to new or updated pages since your last login can be quite useful.

On this page:

- Add this macro to your page
- Change the macro parameters
- Other ways to add this macro

Screenshot: The Global Reports macro

Add this macro to your page

To add the Global Reports macro to a page:

1. From the editor toolbar, choose **Insert** > **Other Macros**.
2. Choose**Global Reports**from the **Reporting**category.
3. Choose**Insert**.

You can then publish your page to see the macro in action.

Change the macro parameters

Macro parameters are used to change the behaviour of a macro.

To change the macro parameters:

1. In the editor, click the macro placeholder and choose**Edit**.
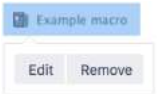

2. Update the parameters as required then choose**Insert**.

Here's a list of the parameters available in this macro.

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.

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`).

Other ways to add this macro

Add this macro as you type

Type **{** followed by the start ofthe macro name, to see a list of macros.

Add this macro using wiki markup

This is useful when you want to add a macro outside the editor, for example as custom content in the sidebar, header or footer of a space.

Macro name: `global-reports`

Macro body: None.

```
{global-reports:width=50%}
```


Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

HTML Include Macro

Add the HTML Include macro to a page toinclude the contents of specific URL in a Confluence page. This allows you to embed a webpage in your Confluence page.

This is a **legacy** macro, and is often disabled by Confluence administrators for security reasons.

On this page:

- Security considerations
- Add this macro to your page
- Change the macro parameters
- Enabling the HTML Include Macro
- Troubleshooting
- Other ways to add this macro

Security considerations

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.

Add this macro to your page

To add the HTML Include macro to a page:

1. From the editor toolbar, choose **Insert** > **Other Macros**.
2. Choose**HTML Include**from the**External content**category.
3. Enter the URL you want to display.
4. Choose**Insert**.

You can then publish your page to see the macro in action.

Change the macro parameters

Macro parameters are used to change the behaviour of a macro.

To change the macro parameters:

1. In the editor, click the macro placeholder and choose**Edit**.

2. Update the parameters as required then choose**Insert**.

Here's a list of the parameters available in this macro.

Parameter	Required	Default	Description
HTML Page's URL (`url`)	Yes	None	The URL of the page to include.

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`).

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 ⚙ > **Manage apps**
2. Select **System** from the drop down and search for the **Confluence HTML Macros** system app.
3. Expand the listing and enable the **html-include (html-include-xhtml)**module.

Administrators can also choose to use the allowlist to restrict URLs that can be displayed in the HTML Include macro.

Troubleshooting

- Administrators can define anallowlist of trusted URLs. If a URL is not in the allowlist, 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
  ~~CONFSERVER-6567~~ - HTML Include macro should rewrite relative links to point to remote site CLOSED
  .

Other ways to add this macro

Add this macro as you type

Type**{**followed by the start ofthe macro name, to see a list of macros.

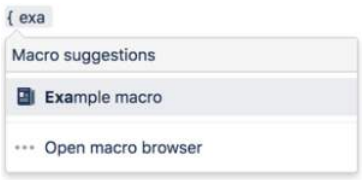

Add this macro using wiki markup

This is useful when you want to add a macro outside the editor, for example as custom content in the sidebar, header or footer of a space.

Macro name:`html-include`

Macro body: None.

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


Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

HTML Macro

This macro is available in **Confluence Server and Data Center**. Learn about the macros available in Confluence Cloud.

Add the HTML macro to a page to embed content from an external website. For example you can use this macro to embed content from your company website, or a web-based tool.

This is a **legacy** macro, and is often disabled by Confluence administrators for security reasons.

On this page:

- Security considerations
- Add this macro to your page
- Change the macro parameters
- Enabling the HTML Macro
- Other ways to add this macro

Security considerations

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.

Add this macro to your page

To add the HTML macro to a page:

1. From the editor toolbar, choose **Insert** > **Other Macros**.
2. Choose **HTML**from the**Development**category.
3. Choose**Insert**
4. Paste the HTML embed code from the website you want to display into the body of the macro.

You can then publish your page to see the macro in action.

Change the macro parameters

This macro has no parameters.

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 > **Manage apps**.
2. Select **System** from the drop down and search for the **Confluence HTML Macros** system app.
3. Expand the listing and enable the **html (html-xhtml)** module.

Other ways to add this macro

Add this macro as you type

Type**{**followed by the start ofthe macro name, to see a list of macros.

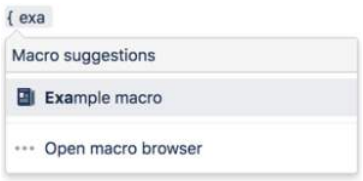

Add this macro using wiki markup

This is useful when you want to add a macro outside the editor, for example as custom content in the sidebar, header or footer of a space.

Macro name:`html`

Macro body: Text, consisting of HTML code.

```
{html}<a href="http://www.atlassian.com">Click here</a> to see the <b>Atlassian</b> website.{html}
```


Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

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

> **We ended support for this macro in Confluence 7.0**
> The macro no longer appears in the macro browser and can't be added to a page.
> Any macro already on a page will still work.

Parameters

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.

Wiki markup example

This is useful when you want to add a macro outside the editor, for example as custom content in the sidebar, header or footer of a space.

Macro name:`im`

Macro body: None.

```
{im:MySkypeName|service=skype|showid=false}
```

Include Page Macro

Add the Include Page macro to a page to display the contents of another page or blog post in this page.

This macro is great for:

- single-sourcing instructions and procedures
- sharing useful information in multiple spaces
- all types of content reuse.

You can add multiple Include Page macros to your page, and combine them with text, images, tables and other macros.

Because you're simply including the content of the other page, rather than duplicating it, any changes to the original automatically flow through to wherever the page is used.

The Include Page macro respects space permissions and page restrictions, so be sure tocheck who can viewthe page you're including.

On this page:

- Add this macro to your page
- Change the macro parameters
- Limitations
- Other ways to add this macro

Screenshot: Meeting notes page with a reusable warning at the top.

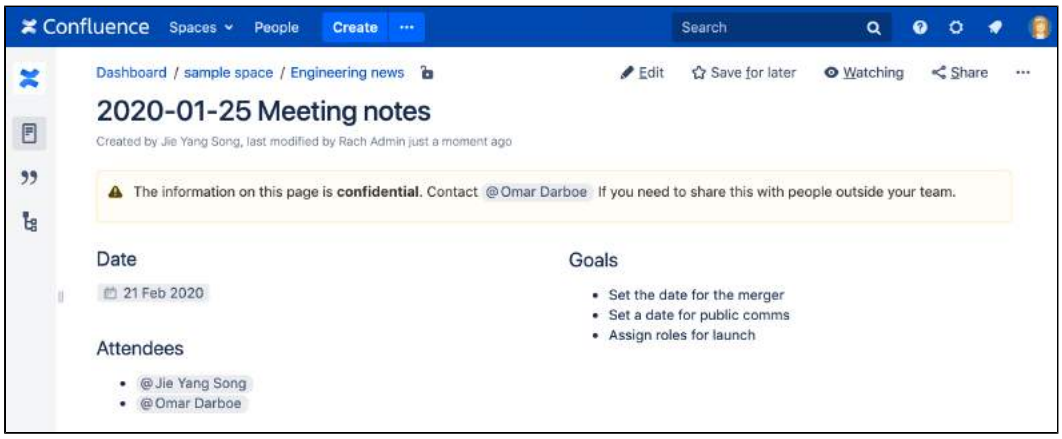

To learn how to create a reusable content library, check out Develop Technical Documentation in Confluence.

Add this macro to your page

To add the Include Page macro to a page:

1. From the editor toolbar, choose **Insert** + > **Other Macros**.
2. Choose**Include Page**from the **Confluence content**categories.
3. Enter the title of the page you want to include. It can be in this space, or another space.
4. Choose**Insert**.

You can then publish your page to see the macro in action.

Screenshot: The Include Page macro nested within a Note macro in the editor.

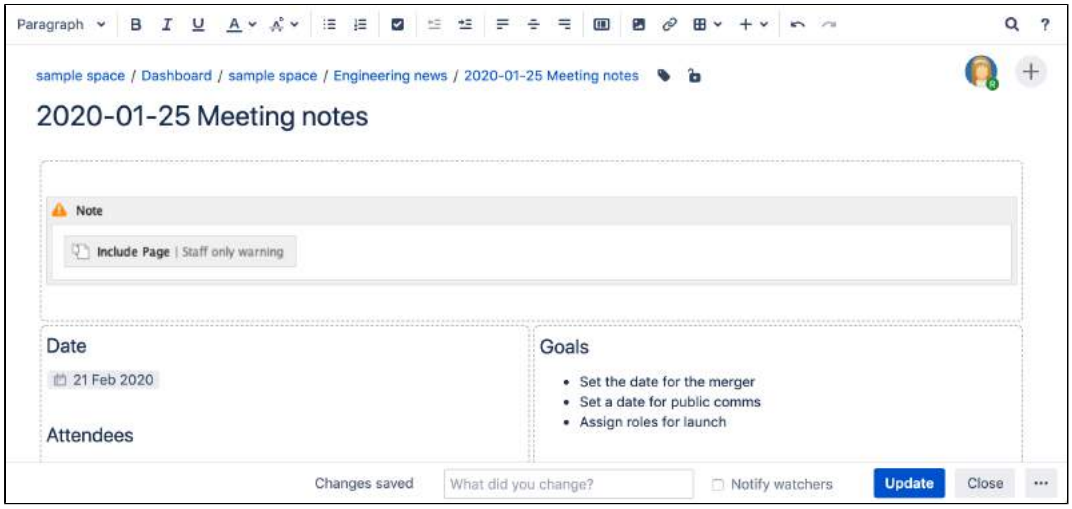

Change the macro parameters

Macro parameters are used to change the behaviour of a macro.

To change the macro parameters:

1. In the editor, click the macro placeholder and choose**Edit**.
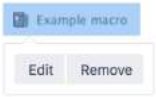

2. Update the parameters as required then choose**Insert**.

Here's a list of the parameters available in this macro.

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`.

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`).

Limitations

This macro has a few limitations you need to be aware of:

- The macro respects space permissions and page restrictions. It won't display the contents of the included page to anyone who doesn't have adequate permissions to see the included page.
- The macro will include the entire page content. If you only want to display part of a page, use theExcerpt Include Macroinstead.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

- The macro can only display pages that exist in your current site. You can't use the Include Page macro to display the contents of pages in other Confluence sites.

Other ways to add this macro

Add this macro as you type

Type**{**followed by the start ofthe macro name, to see a list of macros.

Add this macro using wiki markup

This is useful when you want to add a macro outside the editor, for example as custom content in the sidebar, header or footer of a space.

Macro name:`include`

Macro body: None.

```
{include:DOC:My chocolate page}
```


Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Info, Tip, Note, and Warning Macros

Add an Info, Tip, Note, or Warning macro to your page to make important information stand out. The macro is used to format your text in a colored panel.

As the names suggest, these macros are great for:

- Tips and tricks
- Warnings
- Important notes and other info.

You can configure this macro with or without the icon and title.

On this page:

- Add this macro to your page
- Change the macro parameters
- Other ways to add this macro

Screenshot: examples of the Info, Note, Tip, and Warning macros on a page

Want more control over the color of the boxes? Try the Panel Macro.

Add this macro to your page

To add the Info, Tip, Note, or Warning macro to a page:

1. From the editor toolbar, choose **Insert** > **Other Macros**.
2. Choose either **Info**, **Tip**, **Note**, or**Warning**from the**Formatting**categories.
3. Enter a title for the panel.
4. Choose**Insert**.

You can then publish your page to see the macro in action.

Screenshot: The Info, Tip, Note, and Warning macros in the editor.

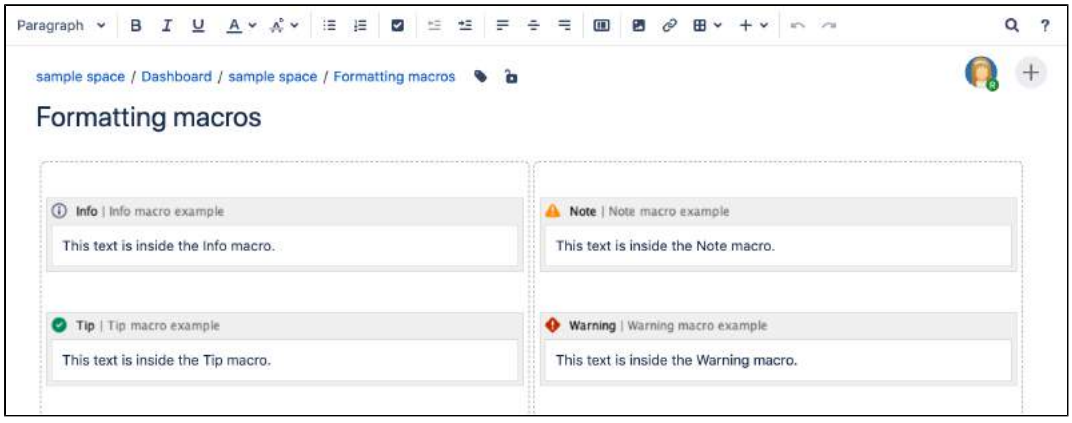

Change the macro parameters

Macro parameters are used to change the behaviour of a macro.

To change the macro parameters:

1. In the editor, click the macro placeholder and choose**Edit**.

2. Update the parameters as required then choose**Insert**.

Here's a list of the parameters available in this macro.

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.

Other ways to add this macro

Add this macro as you type

Type**{**followed by the start ofthe macro name, to see a list of macros.

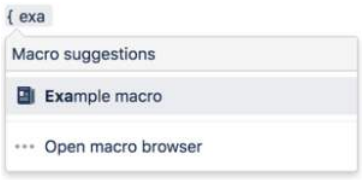

Add this macro using wiki markup

This is useful when you want to add a macro outside the editor, for example as custom content in the sidebar, header or footer of a space.

Macro name:`info/tip/note/warning`

Macro body: Accepts rich text.

```
{info:title=This is my title|icon=false}
This is _important_ information.
{info}
```

Using Confluence Cloud? Head to Add formatting to your page if your info panel looks more like the one below.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Jira Chart Macro

Add the Jira Charts macro to a page to build visual graphs and charts from information in Jira.

This is great for:

- team meetings and retrospectives
- project status updates
- sharing updates with people in your organization who don't use Jira regularly.

The macro can display issues from any connected Jira Server, Data Center, or Cloud application, including Jira Software and Jira Service Management.

On this page:

- Connect Confluence and Jira
- Add this macro to your page
- Pie chart
- Created vs Resolved chart
- Two Dimensional Chart
- Disable the Jira Chart macro
- Notes
- Other ways to add this macro

Screenshot: Meeting notes page with a Jira Chart macro showing a breakdown of all issues for a software version, by priority.

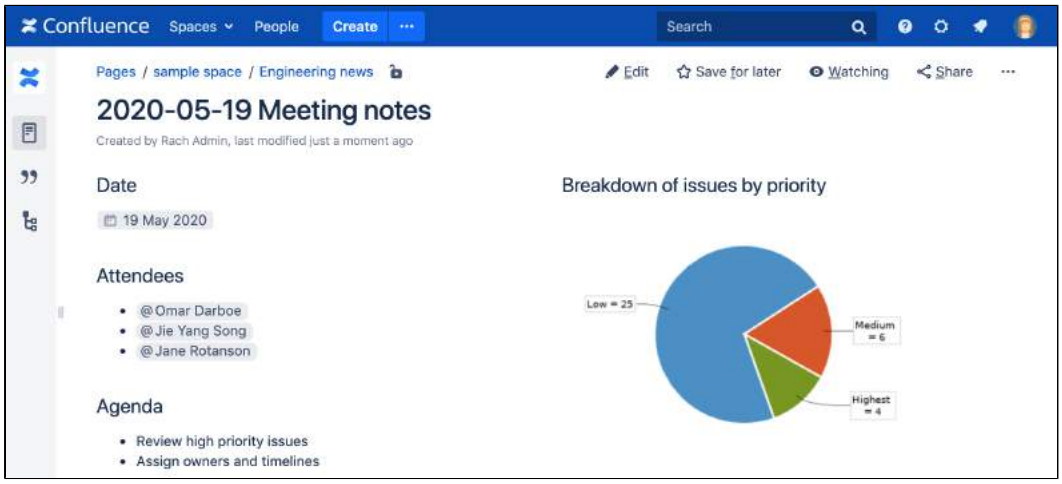

Connect Confluence and Jira

Before you can use this macro, your Confluence and Jira applications must be connected viaApplication 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.

SeeUse Jira applications and Confluence together for more information.

Add this macro to your page

To add the Jira Chart macro to a page:

1. From the editor toolbar, choose **Insert** > **Other Macros**.
2. Choose**Jira Chart**from the**Development**category.
3. Choose the type of report you want to create (for example Pie, Created vs Resolved)
4. Select your **Jira server**.
   If you have multiple Jira servers linked to Confluence the drop down will default to the primary application link.
5. Search for issues - you can enter the query in JQL or paste a Jira URL directly into the search field.
6. Choose**Preview**to generate the chart.
7. Choose**Display Options**to further control how your chart appears.
8. Choose**Insert**.

You can then publish your page to see the macro in action.

To find out more about searching for issues seeDisplaying issues using JIRA Query Language (JQL).

Screenshot: Configuring the Jira Chart Macro in the macro browser.

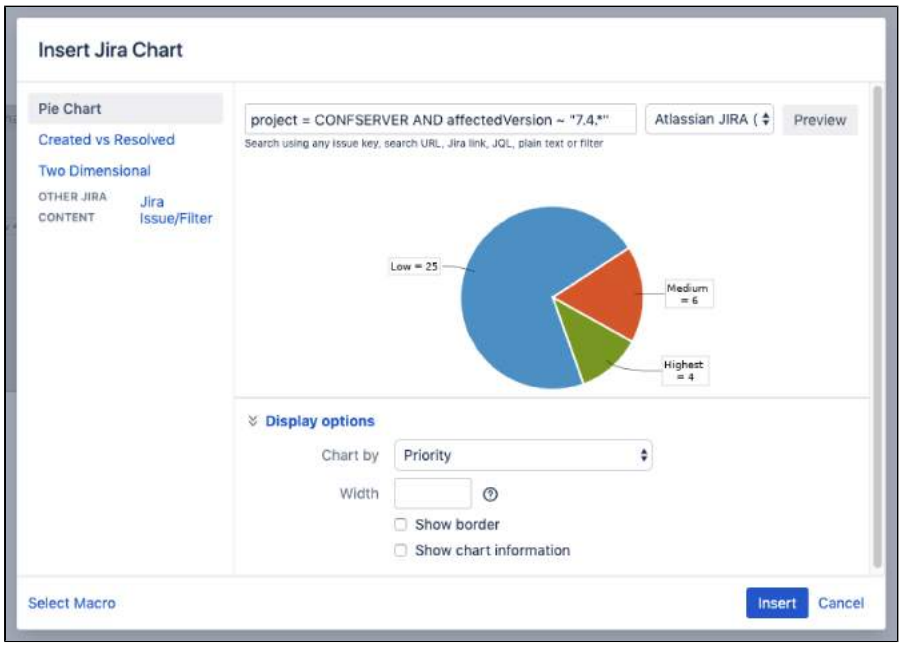

Pie chart

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.

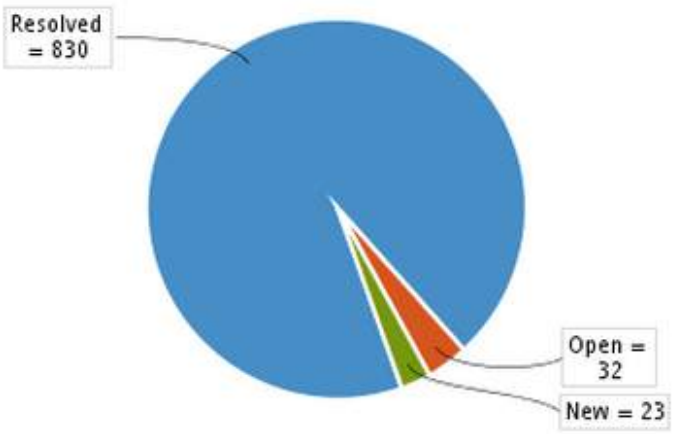

Created vs Resolved chart

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

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.

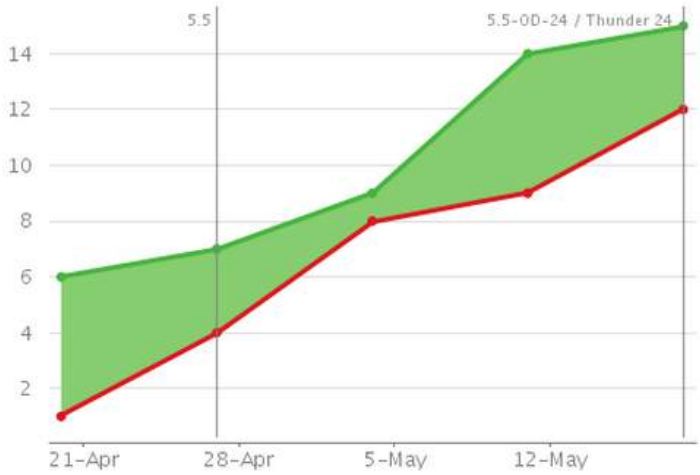

Two Dimensional Chart

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.

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.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Status	Priority			
	Highest	Medium	Low	T:
IN PROGRESS	1	1	0	2
CLOSED	1	0	4	5
NEEDS TRIAGE	0	0	5	5
AWAITING MERGE	1	0	0	1
Total Unique Issues:	3	1	9	13

Showing 4 of 4 statistics.
View in Jira

Disable the Jira Chart macro

The functionality is provided by a system app (plugin) called 'Jira Macros'. This is also used for the Jira Issues macro. To make the macro unavailable on your site, you can disable the app. SeeDisabling and enabling apps.

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. SeeConnecting 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 will need to authenticate.

Other ways to add this macro

Add this macro as you type

Type**{**followed by the start ofthe macro name, to see a list of macros.

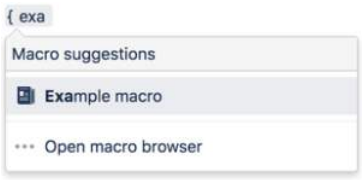

Add this macro using wiki markup

You can't use wiki markup to add this macro.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Jira Issues Macro

Add the Jira Issues macro to a page to display information from Jira. You can display a single issue, a list of issues, or a count, based on aJIRA Query Language (JQL)search, filter, or URL.

This is great for:

- team meetings and retrospectives
- project status updates
- release notes and customer communications
- sharing updates with people in your organization who don't use Jira regularly.

The macro can display issues from any connected Jira Server, Data Center, or Cloud application, including Jira Software and Jira Service Management.

On this page:

- Connect Confluence and Jira
- Add this macro to your page
- Displaying issues via a Jira Query Language (JQL) search
- Displaying issues via a Jira URL
- Displaying a single issue, or selected issues
- Displaying a count of issues
- Creating a new issue
- Configuring application links to display restricted issues
- Rendering HTML from Jira applications
- Disabling the Jira Issues macro
- Notes
- Other ways to add this macro

Screenshot: Project status page with a Jira issues macro showing issues that must be resolved before release.

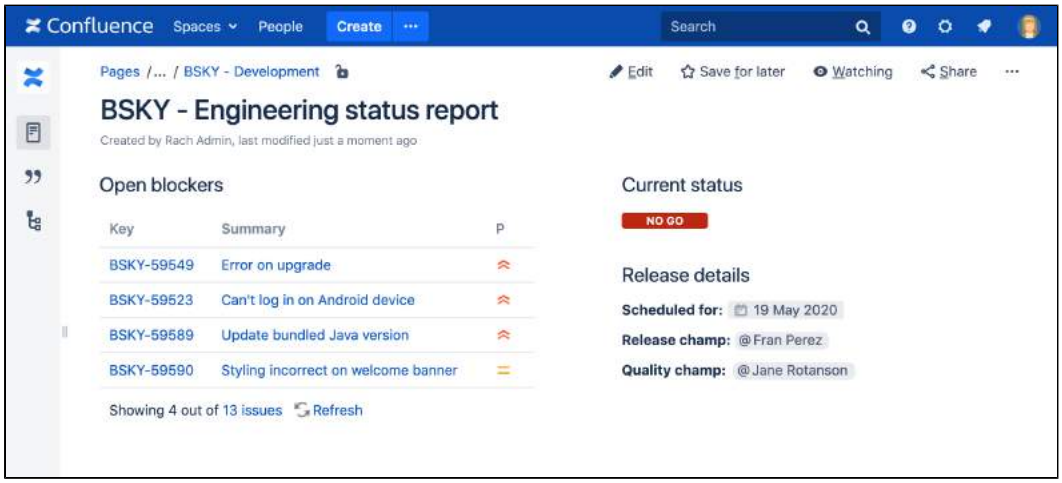

Connect Confluence and Jira

Before you can use this macro, your Confluence and Jira application must be connected viaApplication 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.

SeeUse Jira applications and Confluence togetherfor more information.

Add this macro to your page

To add the Jira macro to a page:

1. From the editor toolbar, choose **Insert** + > **Other Macros**.
2. Choose**Jira**from the**Development**category.
3. Enter a filter or search for a Jira issue.
4. Choose**Insert**.

You can then publish your page to see the macro in action.

Screenshot: Configuring the Jira Issues Macro to show a list of issues.

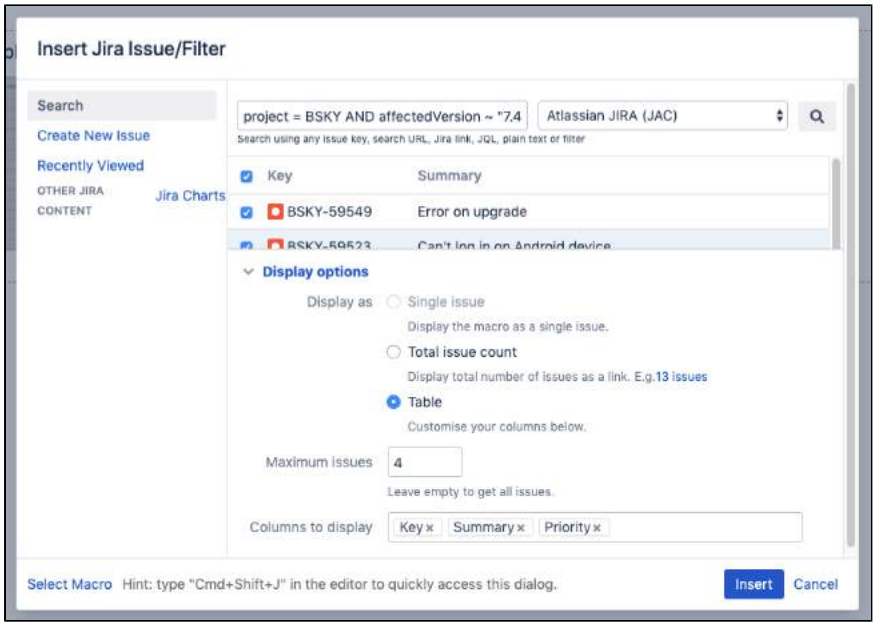

Displaying issues via a Jira Query Language (JQL) search

You can use the macro to display atableof issues on your page, based on the results of a search usingJIRA 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 *operator* (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:

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

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 customize 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.

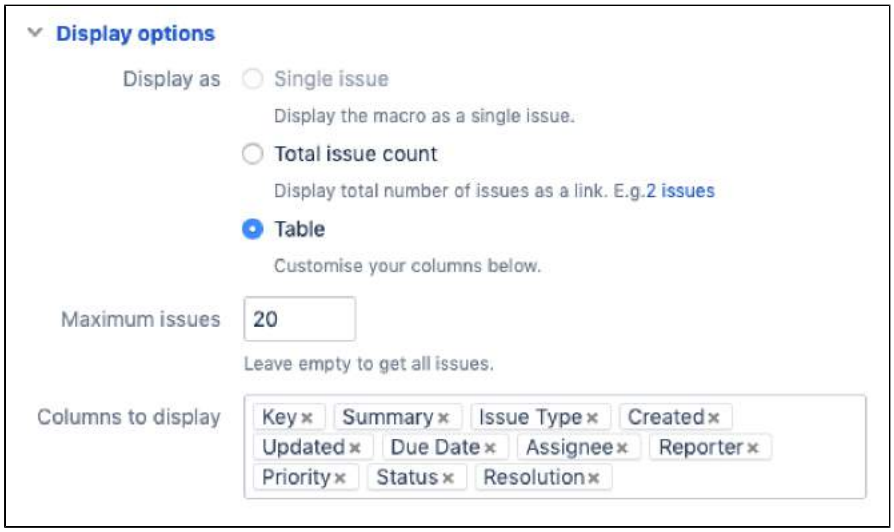

Displaying issues via a Jira URL

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.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Screenshot: Selecting a subset of issues to display

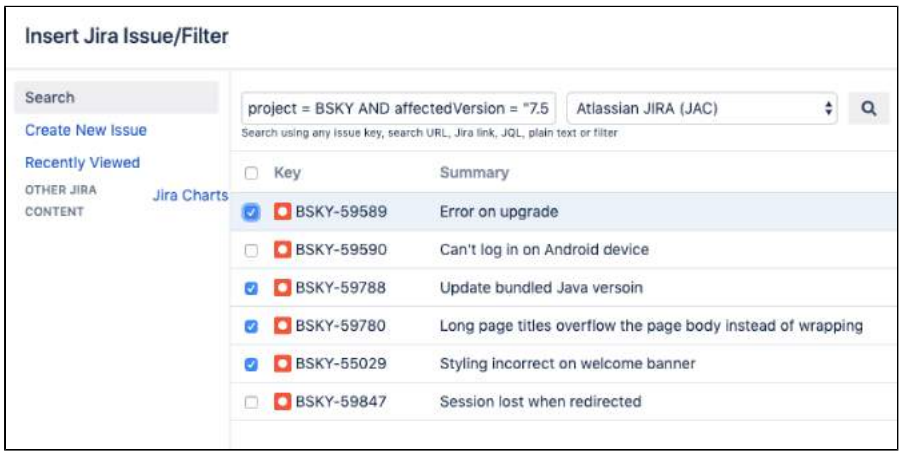

Displaying a count of issues

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.

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 anew 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 customizationthat is not supported by the Jira Issues macro. In this situation you will need to create the issue directly in your Jira application.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Configuring application links to display restricted issues

Before you can use this macro, your Confluence and Jira sites must be connected viaApplication 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 be prompted to **Log in & Approve** before seeing the restricted issues.

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 a system app called 'Jira Macros'. To make the macro unavailable on your site, you can disable the app. SeeDisabling and enabling apps.

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. SeeConnecting 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`.

Other ways to add this macro

Add this macro as you type

Type**{**followed by the start ofthe macro name, to see a list of macros.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

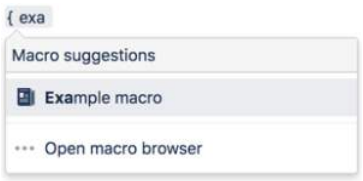

Add this macro using wiki markup

This is useful when you want to add a macro outside the editor, for example as custom content in the sidebar, header or footer of a space.

Macro name:`jiraissues`

Macro body: None.

This example uses a Jira filter.

```
{jiraissues:anonymous=true|url=https://jira.atlassian.com/issues/?filter=41225}
```

A number of additional parameters that are not available via the macro browser are available in storage format and wiki markup.

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 authorized 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' paramet er	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.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

`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 many columns recognized by your Jira application, including custom columns. Some columns, such as those that need to be calculated by Jira like 'work ratio' or 'time to resolution', can't be viewed in Confluence.
`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.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

`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 customize 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.

Do more with Confluence and Jira

Take displaying Jira issues to the next level, with these apps on the Atlassian Marketplace:

- Issue Macro from Jira to Confluence:customize the look of a single Jira Issue report or generate a well-formatted filter report
- Issues Forms for Confluence: Create and display Jira issues/tickets on Confluence pages

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Labels List Macro

Add the Labels List macro to a page to create a directory of all the labels used in a space. People can then click a label to see a list of all pages with that label.

This macro is great for providing an alternative way of navigating the content of a space, especially if you use Confluence for process, procedure, or other documentation.

On this page:

- Add this macro to your page
- Change the macro parameters
- Other ways to add this macro

Screenshot: Page with a Labels List macro to help people find help guides on a particular topic

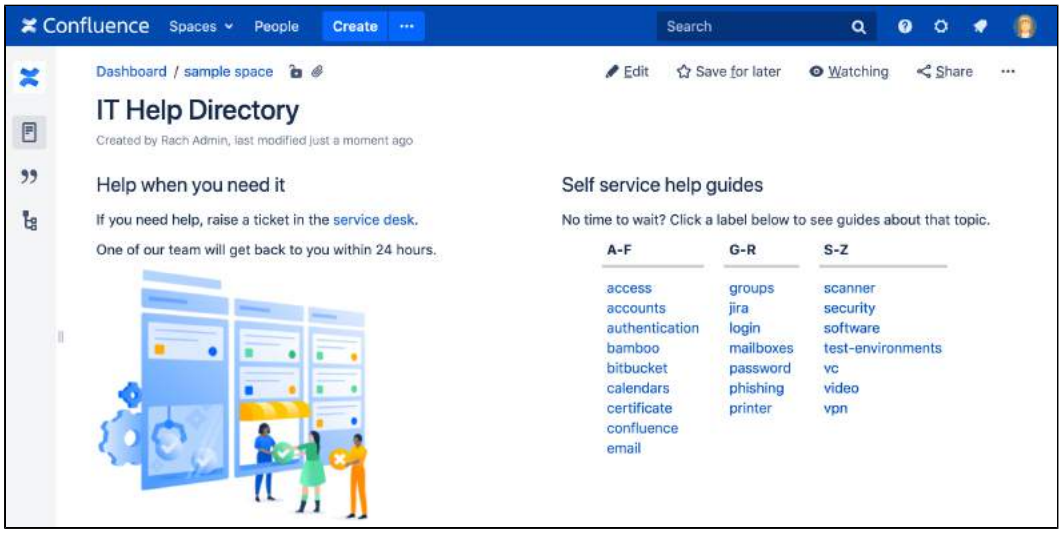

For general information about using labels in Confluence, see Add, Remove and Search for Labels.

Add this macro to your page

To add the Labels List macro to a page:

1. From the editor toolbar, choose **Insert** + > **Other Macros**.
2. Choose**Labels List**from the**Confluence content**categories.
3. Enter a space key, and any labels you might want to exclude.
4. Choose**Insert**.

You can then publish your page to see the macro in action.

Screenshot: The Labels List macro configured to show labels from the IT Help space, and exclude the labels 'test' and 'testing'.

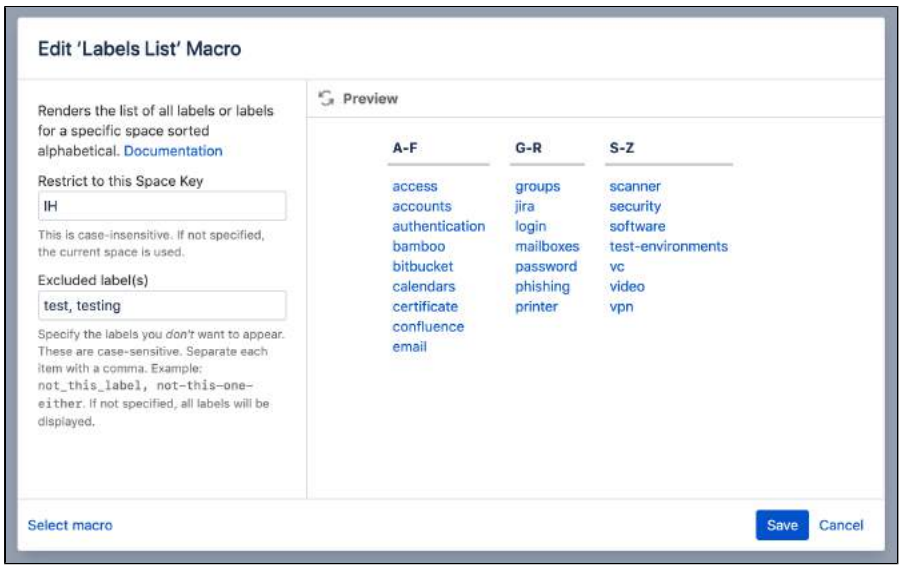

Change the macro parameters

Macro parameters are used to change the behaviour of a macro.

To change the macro parameters:

1. In the editor, click the macro placeholder and choose**Edit**.

2. Update the parameters as required then choose**Insert**.

Here's a list of the parameters available in this macro.

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.

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`).

Other ways to add this macro

Add this macro as you type

Type**{**followed by the start ofthe macro name, to see a list of macros.

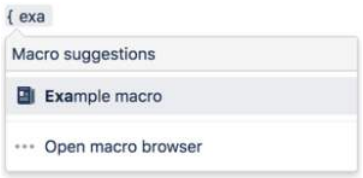

Add this macro using wiki markup

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

This is useful when you want to add a macro outside the editor, for example as custom content in the sidebar, header or footer of a space.

Macro name: `listlabels`

Macro body: None.

```
{listlabels:spaceKey=DOC}
```


Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Livesearch Macro

Add the Livesearch macro to a page to provide a fully customizable search field. When someone starts typing into the search field, Confluence will suggest matching pages, blogs, or comments.

This macro is great when you want to:

- provide a way to search the current space, such as in a knowledge base
- encourage people to search for particular types of content, such as pages with a particular label.

Because you can limit the search by space, label, content type, you can provide a very targeted search experience for people viewing your space.

On this page:

- Add this macro to your page
- Change the macro parameters
- Other ways to add this macro

Screenshot: page with a Livesearch macro showing search results for the search term 'printer'

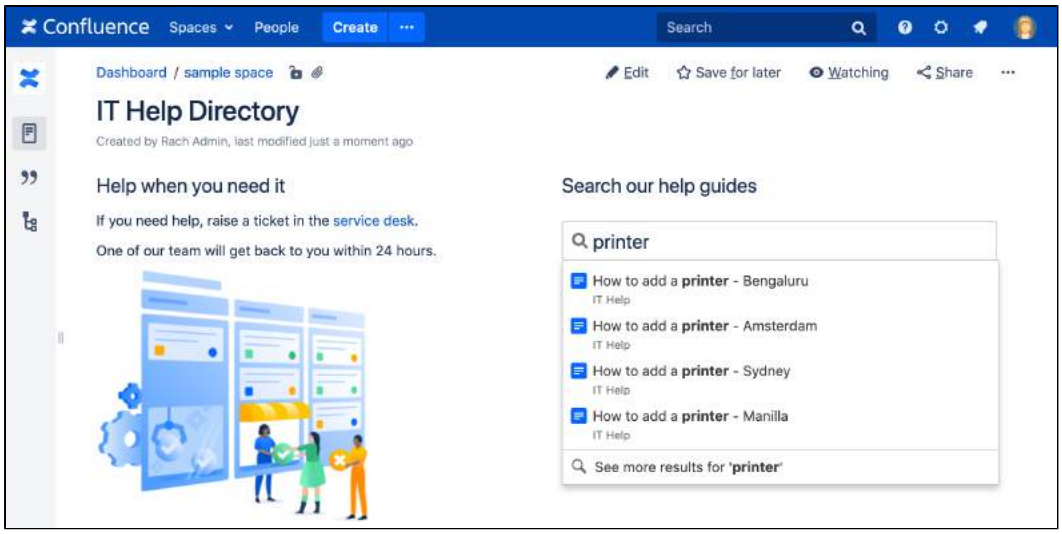

Add this macro to your page

To add the Livesearch macro to a page:

1. From the editor toolbar, choose **Insert** > **Other Macros**.
2. Choose**Livesearch**from the**Navigation**category.
3. Use the parameters below to narrow down the content to be searched.
4. Choose**Insert**.

You can then publish your page to see the macro in action.

Screenshot: configuring the Livesearch macro to search for pages with particular labels in a specific space.

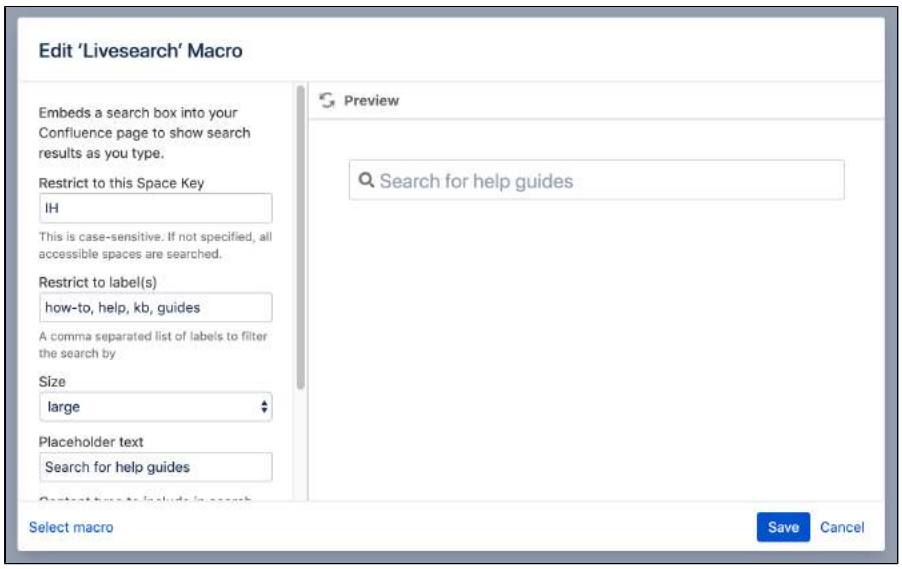

Change the macro parameters

Macro parameters are used to change the behaviour of a macro.

To change the macro parameters:

1. In the editor, click the macro placeholder and choose**Edit**.

2. Update the parameters as required then choose**Insert**.

Here's a list of the parameters available in this macro.

Parameter	Default	Description
Restrict to this Space Key (`spaceKey`)	all spaces	Specify a space key to limit the search to a single space. Case-sensitive. You can't specify multiple spaces. 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.

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`).

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Other ways to add this macro

Add this macro as you type

Type**{**followed by the start ofthe macro name, to see a list of macros.

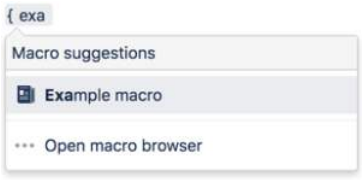

Add this macro using wiki markup

This is useful when you want to add a macro outside the editor, for example as custom content in the sidebar, header or footer of a space.

Macro name:`livesearch`

Macro body: None.

```
{livesearch:spaceKey=DOC|size=large|placeholder=Search this space}
```


Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Loremipsum Macro

Add the Loremipsum macro to a page to display a paragraph of Lorem ipsumplaceholder text.

This is a**legacy macro**, and is mostly used for demonstrating how a particular page layout or format might look with content.The text is deliberately non-meaningful so that it does not influence the viewer's perception of the page arrangement or design.

Screenshot: a page with headings and the Loremipsum macro in place of actual content.

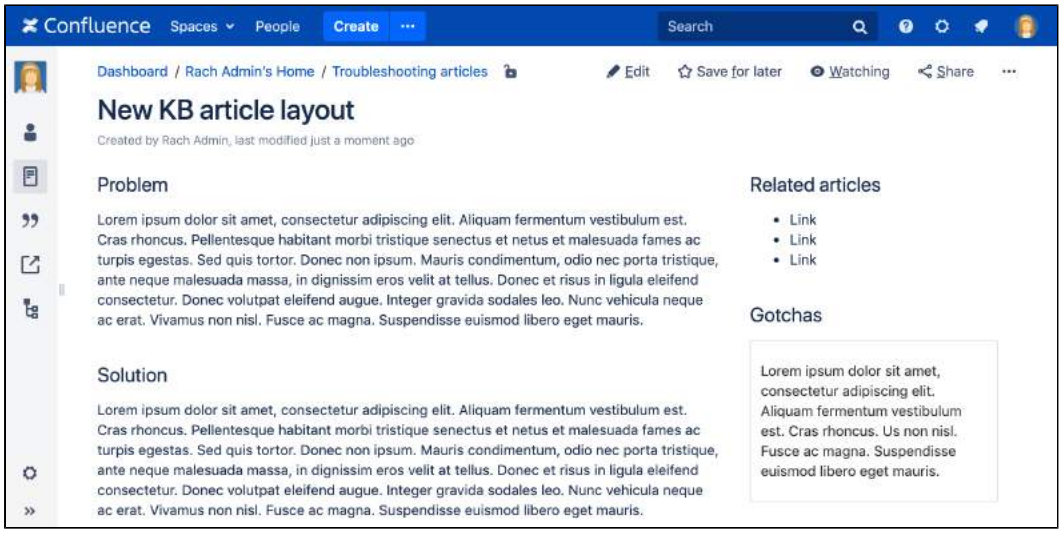

Add this macro to your page

To add the Loremipsum macro to a page:

1. From the editor toolbar, choose **Insert** > **Other Macros**.
2. Choose**Loremipsum**from the**Development**category.
3. Enter the number of paragraphs to include.
4. Choose**Insert**.

You can then publish your page to see the macro in action.

Change the macro parameters

Macro parameters are used to change the behaviour of a macro.

To change the macro parameters:

1. In the editor, click the macro placeholder and choose**Edit**.
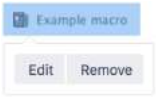

2. Update the parameters as required then choose**Insert**.

Here's a list of the parameters available in this macro.

Parameter	Default	Description
Number of Paragraphs	3	Determines the amount of placeholder text to display. The macro will display a maximum number of 30 paragraphs. Parameter is unnamed in storage format and wiki markup.

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`).

Other ways to add this macro

Add this macro as you type

Type**{**followed by the start ofthe macro name, to see a list of macros.

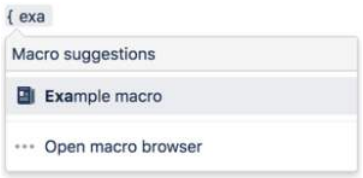

Add this macro using wiki markup

This is useful when you want to add a macro outside the editor, for example as custom content in the sidebar, header or footer of a space.

Macro name:`loremipsum`

Macro body: None.

```
{loremipsum:2}
```


Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Multimedia Macro

Add the Multimedia macro to a page to embed an attached video, animation, or other multimedia file.

The macro uses the HTML5<video>tag, so the type of video your page viewers can see depends on the video formats their browser supports with the HTML5<video>tag.

This macro is great for:

- training videos
- customer interviews and research recordings
- workshop or meeting recordings.

On this page:

- Add this macro to your page
- Change the macro parameters
- Other ways to add this macro

Screenshot: a meeting notes page containing a demo video.

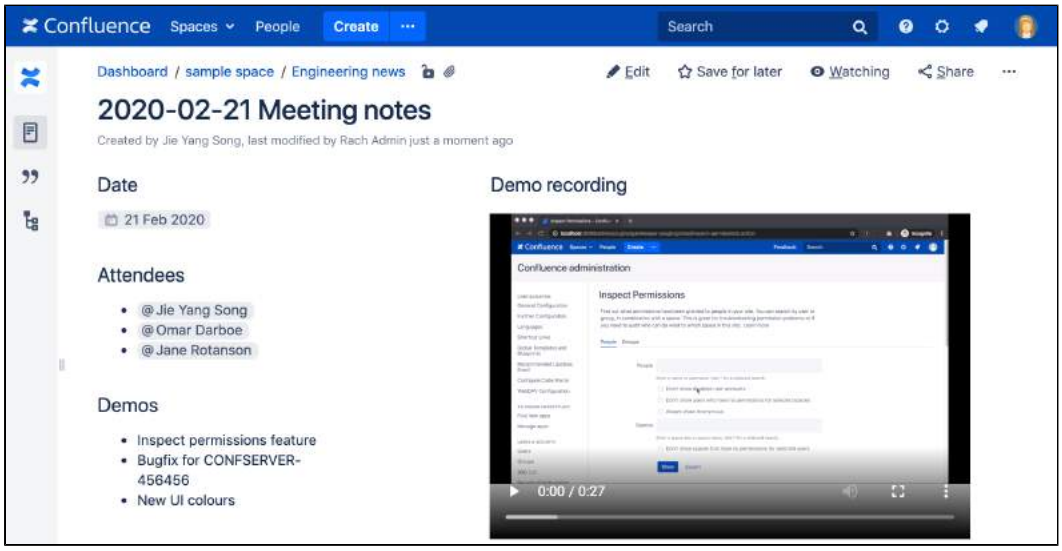

If you want to display online multimedia content, like YouTube and Vimeo videos, take a look at theWidget Connector Macro.

The file previewalso supports MP3 audio and MP4 video files. This is handy when you want to play a video in a larger format.

Add this macro to your page

To add the Multimedia macro to a page:

1. Upload the file to your page.
2. From the editor toolbar, choose **Insert** > **Other Macros**.
3. Choose**Multimedia**from the **Media**category.
4. Select the file you want to play.
5. Choose**Insert**.

You can then publish your page to see the macro in action.

Screenshot: configuring the Multimedia macro

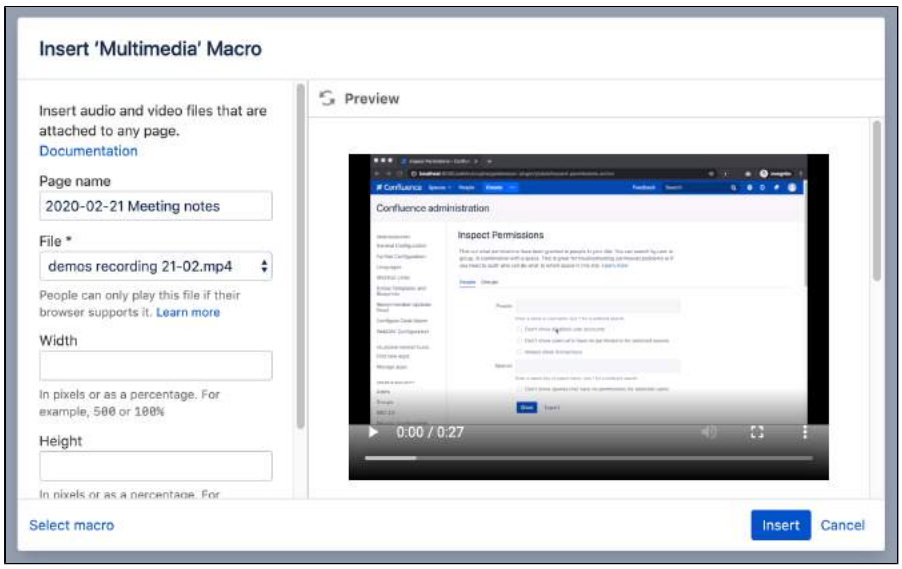

Change the macro parameters

Macro parameters are used to change the behaviour of a macro.

To change the macro parameters:

1. In the editor, click the macro placeholder and choose**Edit**.

2. Update the parameters as required then choose**Insert**.

Here's a list of the parameters available in this macro.

Parameter name	Required	Default	Parameter description and accepted values
Page name (`page`)	No	Current page	The 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`)
File* (`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.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

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`).

Other ways to add this macro

Add this macro as you type

Type**{**followed by the start ofthe macro name, to see a list of macros.

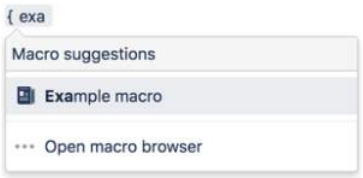

Add this macro using wiki markup

This is useful when you want to add a macro outside the editor, for example as custom content in the sidebar, header or footer of a space.

Macro name:`multimedia`

Macro body: None.

```
{multimedia:space=DOC|page=My macros|name=ninjas.swf|autostart=true}
```


Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Navigation Map Macro

Add the Navigation Map macro to a page to display pages tagged with a specific label in a grid layout.

This macro is great forfor visually representing a small set of pages with a particular label. You could:

- display of all pages that have the label 'needs-review' for highlighting pages that need work
- display all pages with the label 'how-to' in your knowledge base.

If you want to get really fancy, you can style the macro by creating your own Velocity theme. This does require writing some code though.

On this page:

- Add this macro to your page
- Change the macro parameters
- Create your own navmap theme
- Other ways to add this macro

Screenshot: page with a Navigation Map macro displaying pages with the label 'printer-how-to'.

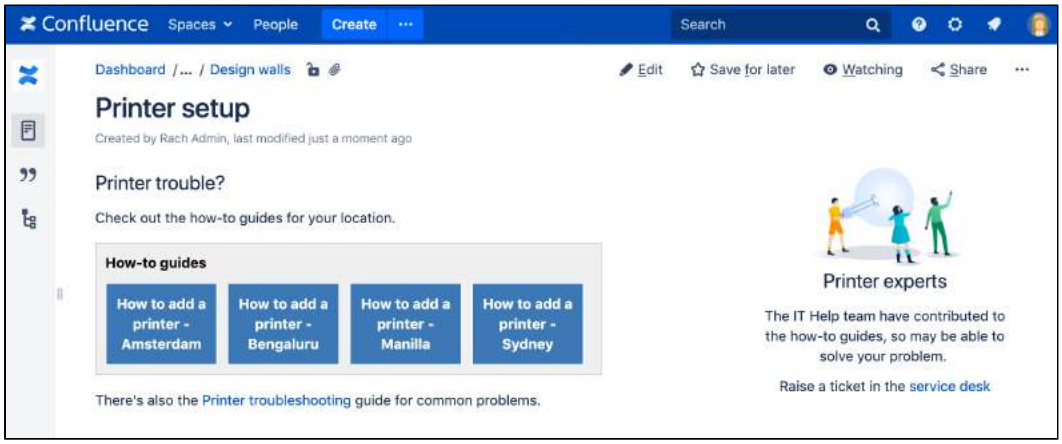

Want more flexibility? Check out the Content by Label Macro for a more modern way to display a list of pages with specific labels and more.

Add this macro to your page

To add the Navigation Map macro to a page:

1. From the editor toolbar, choose **Insert** > **Other Macros**.
2. Choose**Navigation Map**from the**Navigation**category.
3. Enter a label.
4. Choose**Insert**.

You can then publish your page to see the macro in action.

Screenshot: specifying a label and title in the Navigation Map macro

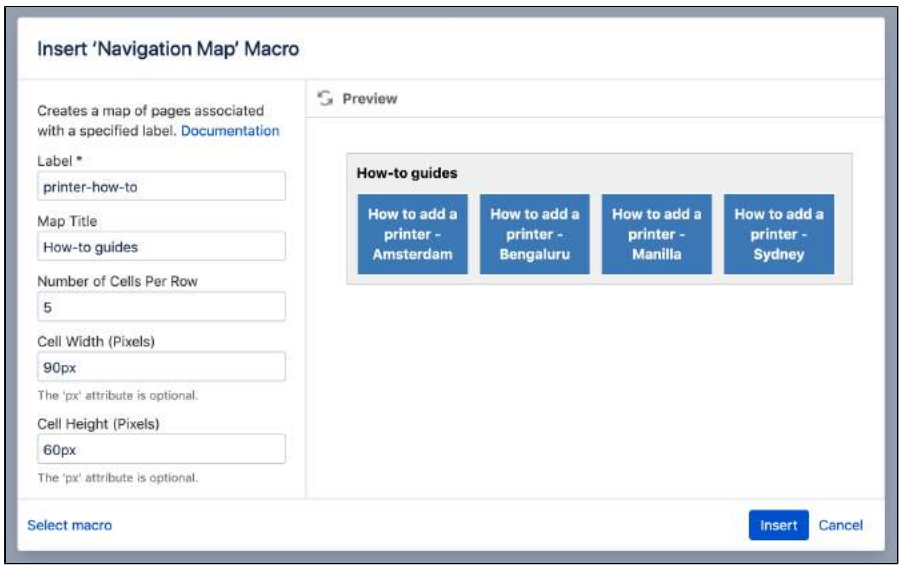

Change the macro parameters

Macro parameters are used to change the behaviour of a macro.

To change the macro parameters:

1. In the editor, click the macro placeholder and choose**Edit**.

2. Update the parameters as required then choose**Insert**.

Here's a list of the parameters available in this macro.

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. See further info below.

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`).

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Create your own navmap theme

If you want to create your own navmap 'look and feel' (for example, one with rounded corners), you need to add a customized navmap macro theme file to the `WEB-INF/classes/com/atlassian/confluence/plugins/macros/advanced` 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.

The theme must be written in Velocity. SeeVelocity User Guide for more information.

Other ways to add this macro

Add this macro as you type

Type**{**followed by the start ofthe macro name, to see a list of macros.

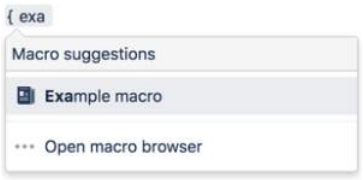

Add this macro using wiki markup

This is useful when you want to add a macro outside the editor, for example as custom content in the sidebar, header or footer of a space.

Macro name:`navmap`

Macro body: None.

```
{navmap:mylabel|wrapAfter=4|title=My map name|cellHeight=50px|theme=navmap-mytheme.vm|cellWidth=80px}
```


Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

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

We ended support for this macro in Confluence 7.0
The macro no longer appears in the macro browser and can't be added to a page.
Any macro already on a page will still work.

Parameters

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.

Wiki markup example

This is useful when you want to add a macro outside the editor, for example as custom content in the sidebar, header or footer of a space.

Macro name:`network`

Macro body: None.

```
{network:followers|username=admin|max=10|theme=full}
```

Disabling the Network macro

The Network macro is provided by the 'network' module in the 'Profile Macros' system app (plugin). To remove the macro from your site, you can disable this module. SeeDisabling and enabling apps.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Noformat Macro

Add the Noformat macro to a page to display text in monospace font with no other formatting.

This is a **legacy macro**, and is similar to theCode Block Macro, but doesn't provide any additional functionality such as line numbering or syntax highlighting.

On this page:

- Add this macro to your page
- Change the macro parameters
- Other ways to add this macro

Add this macro to your page

To add the Noformat macro to a page:

1. From the editor toolbar, choose **Insert** > **Other Macros**.
2. Choose**Noformat**from the**Formatting**category.
3. Choose**Insert**.
4. Paste or type your text into the macro body.

You can then publish your page to see the macro in action.

Screenshot: the Noformat macro in the editor

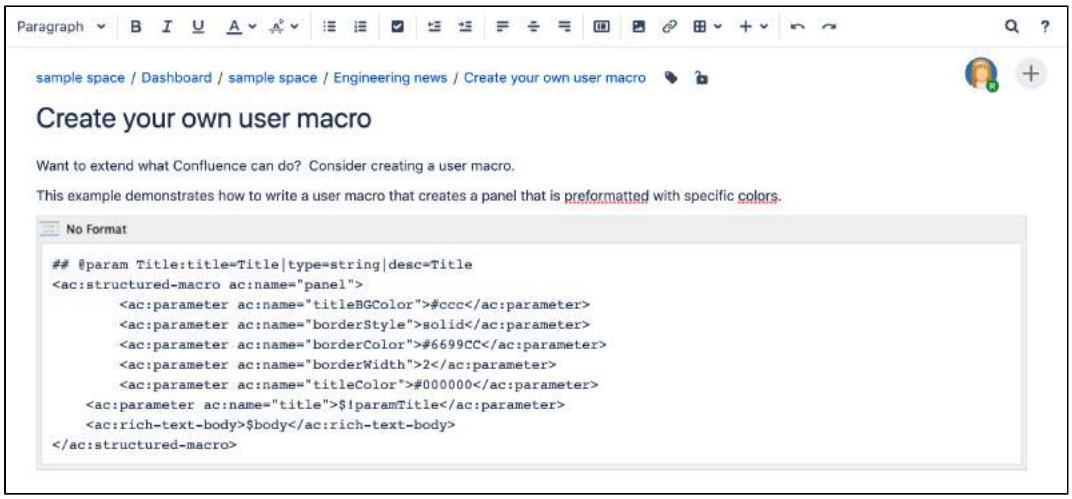

Change the macro parameters

Macro parameters are used to change the behaviour of a macro.

To change the macro parameters:

1. In the editor, click the macro placeholder and choose**Edit**.

2. Update the parameters as required then choose**Insert**.

Here's a list of the parameters available in this macro.

Parameter	Default	Description
No Panel (nopanel)	False	Removes the panel around the content.

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`).

Other ways to add this macro

Add this macro as you type

Type**{**followed by the start ofthe macro name, to see a list of macros.

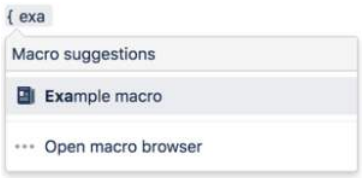

Add this macro using wiki markup

This is useful when you want to add a macro outside the editor, for example as custom content in the sidebar, header or footer of a space.

Macro name:`noformat`

Macro body: Accepts plain text.

```
{noformat:nopanel=true}http://www.example.com{noformat}
```


Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Office Excel Macro

Add the Office Excel macro to a page to displaythe contents of an Excel spreadsheet.

This is great for situations where you need more than a basic Confluence table can provide, such as for financial information or planning data.

This macro embeds your spreadsheet in the page, rather than showing a simple preview. People viewing the page don't need Excel installed to be able to see the spreadsheet.

On this page:

- Add this macro to your page
- Edit the attached file
- Change the macro parameters
- Limitations
- Other ways to add this macro

Screenshot: a page with an Office Excel macro displaying an Excel spreadsheet.

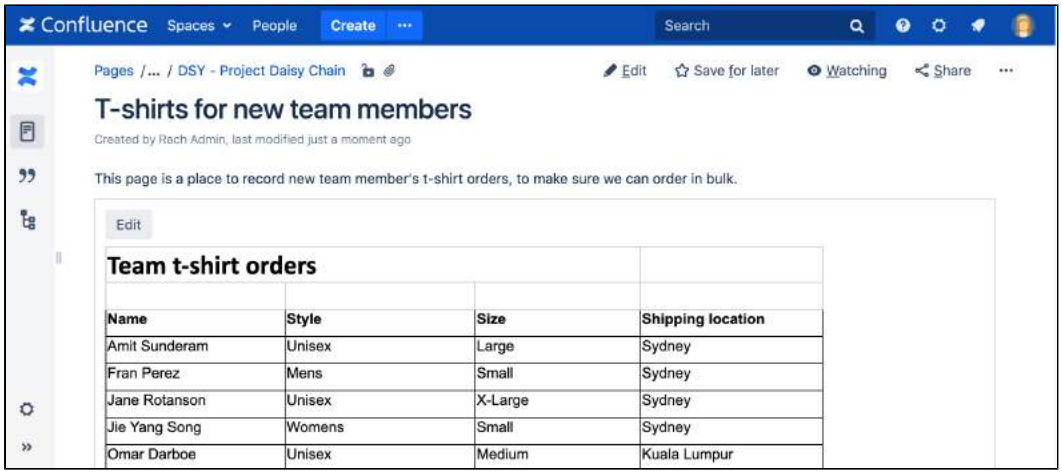

There are other ways to add a spreadsheet to your page:

- Insert the file directly into the page. We'll display a PDF thumbnail of the sheet. This is okay for simple spreadsheets but may not be suitable for complex or multi-sheet files.
- Use the Widget Connector Macro to embed a Google Sheet.

Add this macro to your page

To add the Office Excel macro to a page:

1. Upload the Excel file to your page, then publish the page. SeeUpload Filesto learn how to do this.
2. From the editor toolbar, choose **Insert** + > **Other Macros**.
3. Choose**Office Excel**from the**Confluence content**category.
4. Select the attached file you want to display.
5. Choose**Insert**.

You can then publish your page to see the macro in action.

Screenshot: configuring the Office Excel macro in the macro browser.

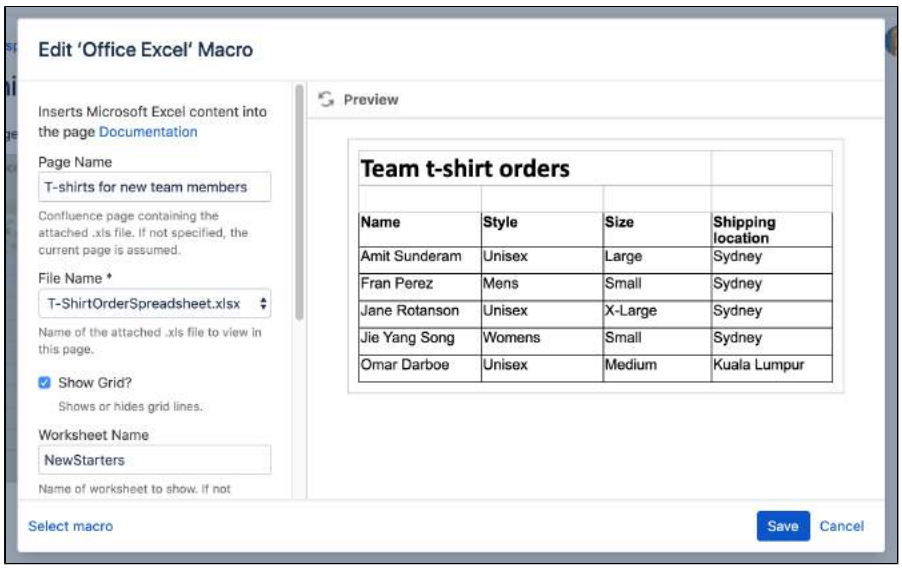

Edit the attached file

If you have Excel installed, you can edit the attached file, and automatically re-upload the file back to Confluence.

SeeEdit Files for more information on the ways to do this.

Change the macro parameters

Macro parameters are used to change the behaviour of a macro.

To change the macro parameters:

1. In the editor, click the macro placeholder and choose**Edit**.
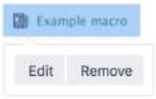

2. Update the parameters as required then choose**Insert**.

Here's a list of the parameters available in this macro.

Parameter	Default	Description
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. If the file does not appear, publish the page, then head back into the editor and try again.
Show Grid?	true	Select to show grid lines around each cell of the Excel spreadsheet. Clear to hide these grid lines.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

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.

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`).

Limitations

The Office file must be attached to the current page, or another Confluence page. The macro can't display live Office 365 files. If you use Office 365, you'll need to download the file, and then upload it to Confluence to display it with this macro. Alternatively, you could just link to the Office 365 file.

If your uploaded file does not appear in the **File Name** menu in the macro browser, you'll need to publish the page, and then hit**Edit**to return to the editor.

Rendering very large or complex files can put a lot of load on Confluence. For this reason, in Confluence Data Center we'll prompt you to download the file if we can't display with a set time limit. This limit varies depending on system properties set by your administrator, but is generally about 30 seconds. You can continue to view other content on the page while we attempt to display the file contents.

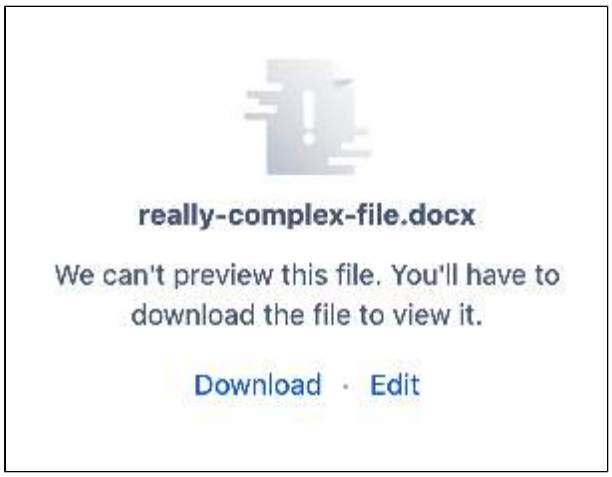

Other ways to add this macro

Add this macro as you type

Type**{**followed by the start ofthe macro name, to see a list of macros.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

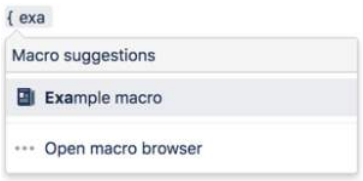

Add this macro using wiki markup

This is useful when you want to add a macro outside the editor, for example as custom content in the sidebar, header or footer of a space.

Macro name:`viewxls`

Macro body: None.

```
{viewxls:col=5|page=Docs|name=My document.xls|grid=false|sheet=mysheet|row=5}
```


Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Office PowerPoint Macro

Add the Office PowerPoint macro to a page to display the content of a PowerPoint presentation.

This is great for sharing presentations, training sessions, and other visual data.

This macro displays your presentation in a viewer with next and back buttons, rather than showing a simple preview. People viewing the page don't need PowerPoint installed to be able to see the presentation.

On this page:

- Add this macro to your page
- Edit the attached file
- Change the macro parameters
- Limitations
- Other ways to add this macro

Screenshot: Project page with an Office PowerPoint macro.

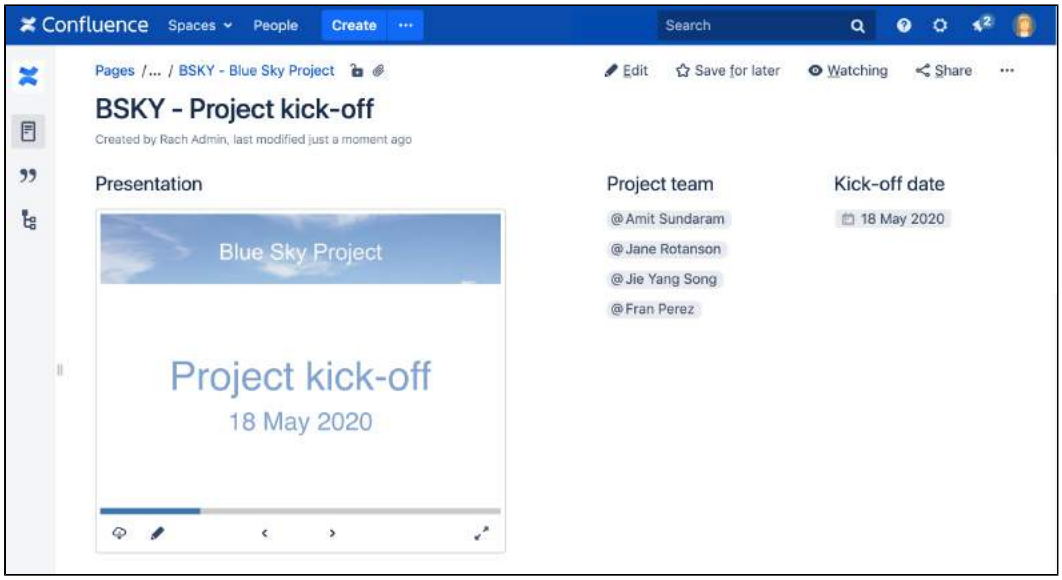

There are multiple ways to show a file on a page. See Display Files and Images for more.

Add this macro to your page

To add the Office PowerPoint macro to a page:

1. Upload a PowerPoint file to your page, then publish the page. SeeUpload Filesto learn how to do this.
2. From the editor toolbar, choose **Insert** + > **Other Macros**.
3. Choose**Office PowerPoint**from the**Confluence content**category.
4. Select the attached PowerPoint file you want to display.
5. Choose**Insert**.

You can then publish your page to see the macro in action.

Screenshot: Configuring the Office PowerPoint macro.

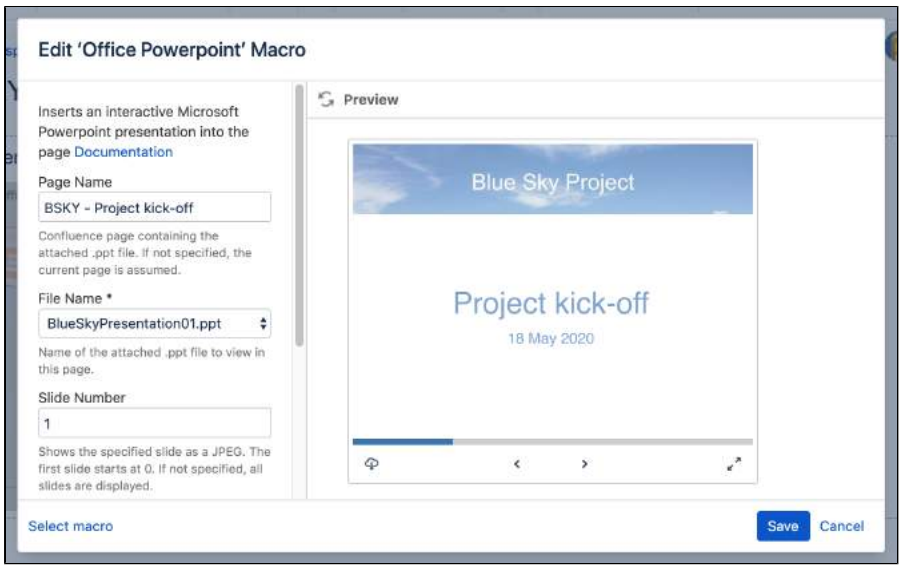

Edit the attached file

If you have PowerPoint installed, hit the Edit icon on the macroto edit the attached file, and automatically re-upload it back to Confluence.

SeeEdit Filesfor more information on the ways to do this.

Change the macro parameters

Macro parameters are used to change the behaviour of a macro.

To change the macro parameters:

1. In the editor, click the macro placeholder and choose**Edit**.
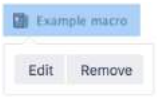

2. Update the parameters as required then choose**Insert**.

Here's a list of the parameters available in this macro.

Parameter	Default	Description
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 PowerPoint file to be displayed. The document must be attached to a page on your Confluence site.
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.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Width		Specify the width of the display, in pixels (default) or as a percentage of the window's width.

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`).

Limitations

The Office file must be attached to the current page, or another Confluence page. The macro can't display live Office 365 files. If you use Office 365, you'll need to download the file, and then upload it to Confluence to display it with this macro. Alternatively, you could just link to the Office 365 file.

If your uploaded file does not appear in the **File Name** menu in the macro browser, you'll need to publish the page, and then hit**Edit**to return to the editor.

Rendering very large or complex files can put a lot of load on Confluence. For this reason, in Confluence Data Center we'll prompt you to download the file if we can't display with a set time limit. This limit varies depending on system properties set by your administrator, but is generally about 30 seconds. You can continue to view other content on the page while we attempt to display the file contents.

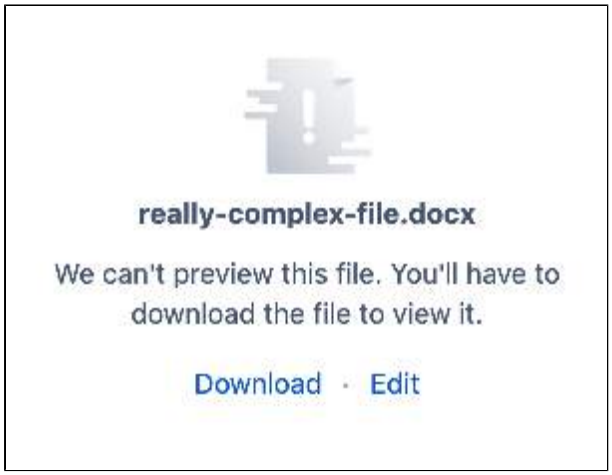

Other ways to add this macro

Add this macro as you type

Type**{**followed by the start ofthe macro name, to see a list of macros.

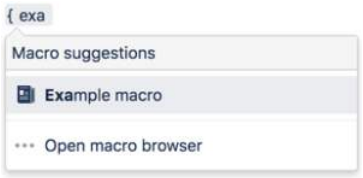

Add this macro using wiki markup

This is useful when you want to add a macro outside the editor, for example as custom content in the sidebar, header or footer of a space.

Macro name:`viewppt`

Macro body: None.

```
{viewppt:height=20%|page=Docs|width=20%|name=My document.ppt|slide=4}
```


Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Office Word Macro

Add the Office Word macro to a page to displaythe contents of a Word document.

This is great for situations where you can't simply copy the content into the Confluence page, because you want to preserve formatting or other Word functionality.

This macro embeds your document in the page, rather than showing a simple preview. People viewing the page don't need Word installed to be able to see the document.

On this page:

- Add this macro to your page
- Edit the attached file
- Change the macro parameters
- Limitations
- Other ways to add this macro

Screenshot: A page with an Office Word macro displaying a Word document with images and tracked changes.

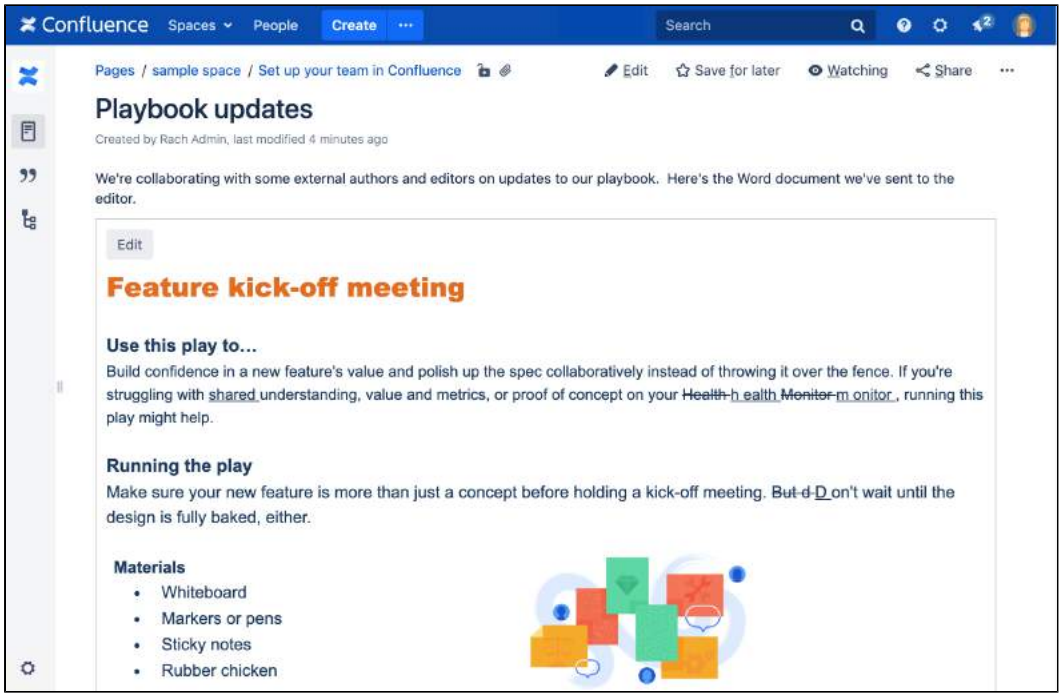

There are other ways to add a document to your page:

- Insert the file directly into the page. We'll display a PDF thumbnail of the document.
- Use the Widget Connector Macro to embed a Google Doc.

Add this macro to your page

To add the Office Word macro to a page:

1. Upload the Word file to your page, then publish the page. SeeUpload Filesto learn how to do this.
2. From the editor toolbar, choose **Insert** + > **Other Macros**.
3. Choose**Office Word**from the**Confluence content**category.
4. Select the attached file you want to display.
5. Choose**Insert**.

You can then publish your page to see the macro in action.

Screenshot: Configuring the Office Word macro in the macro browser.

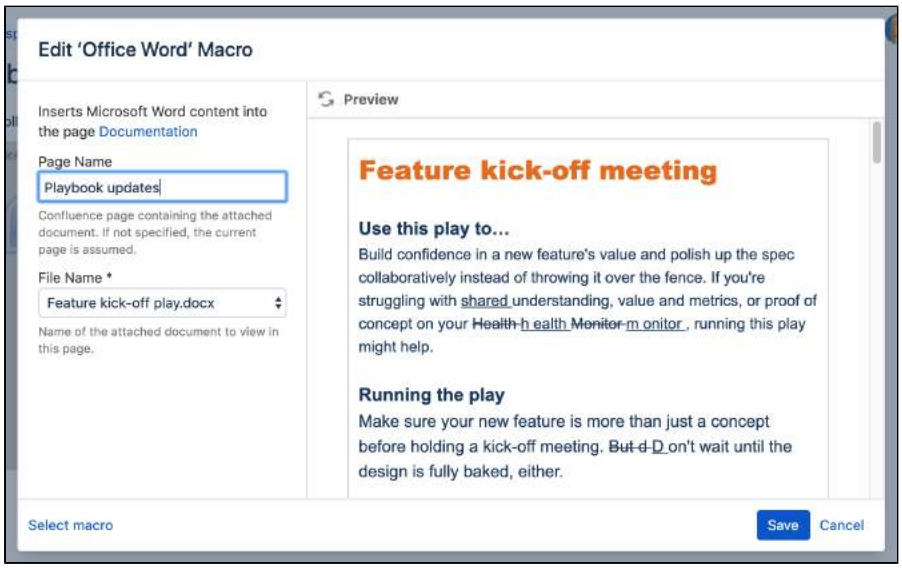

Edit the attached file

If you have Word installed, you can edit the attached file, and automatically re-upload the file back to Confluence.

SeeEdit Filesfor more information on the ways to do this.

Change the macro parameters

Macro parameters are used to change the behaviour of a macro.

To change the macro parameters:

1. In the editor, click the macro placeholder and choose**Edit**.

2. Update the parameters as required then choose**Insert**.

Here's a list of the parameters available in this macro.

Parameter	Default	Description
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. If the file does not appear, publish the page, then head back into the editor and try again.

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`).

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Limitations

The Office file must be attached to the current page, or another Confluence page. The macro can't display live Office 365 files. If you use Office 365, you'll need to download the file, and then upload it to Confluence to display it with this macro. Alternatively, you could just link to the Office 365 file.

If your uploaded file does not appear in the **File Name** menu in the macro browser, you'll need to publish the page, and then hit**Edit**to return to the editor.

Rendering very large or complex files can put a lot of load on Confluence. For this reason, in Confluence Data Center we'll prompt you to download the file if we can't display with a set time limit. This limit varies depending on system properties set by your administrator, but is generally about 30 seconds. You can continue to view other content on the page while we attempt to display the file contents.

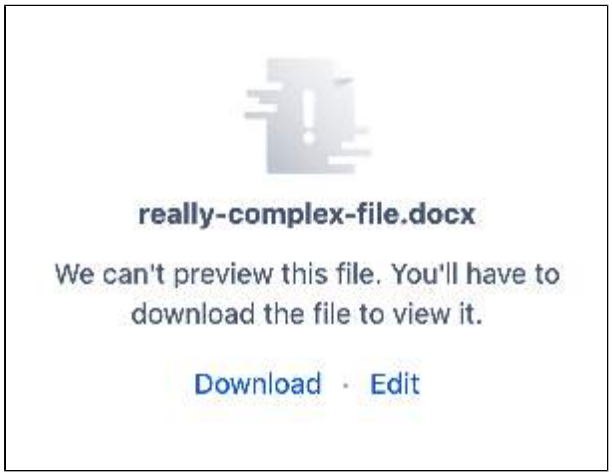

Other ways to add this macro

Add this macro as you type

Type**{**followed by the start ofthe macro name, to see a list of macros.

Add this macro using wiki markup

This is useful when you want to add a macro outside the editor, for example as custom content in the sidebar, header or footer of a space.

Macro name:`viewdoc`

Macro body: None.

```
{viewdoc:page=Docs|name=My document.doc}
```


Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Page Index Macro

Add the Page Index macro to a page to create an alphabetical index of all pages in the current space.

This is a **legacy macro**, and does have some limitations, so it's not suitable for use in very large spaces.

In small spaces, this macro can be useful for providing a directory of pages, such as:

- project pages
- knowledge base articles
- process and procedure pages.

On this page:

- Add this macro to your page
- Limitations
- Change the macro parameters
- Other ways to add this macro

Add this macro to your page

To add the Page Index macro to a page:

1. From the editor toolbar, choose **Insert** + > **Other Macros**.
2. Choose**Page Index**from the **Navigation**category.
3. Choose**Insert**.

You can then publish your page to see the macro in action.

Screenshot: Page Index macro on a page.

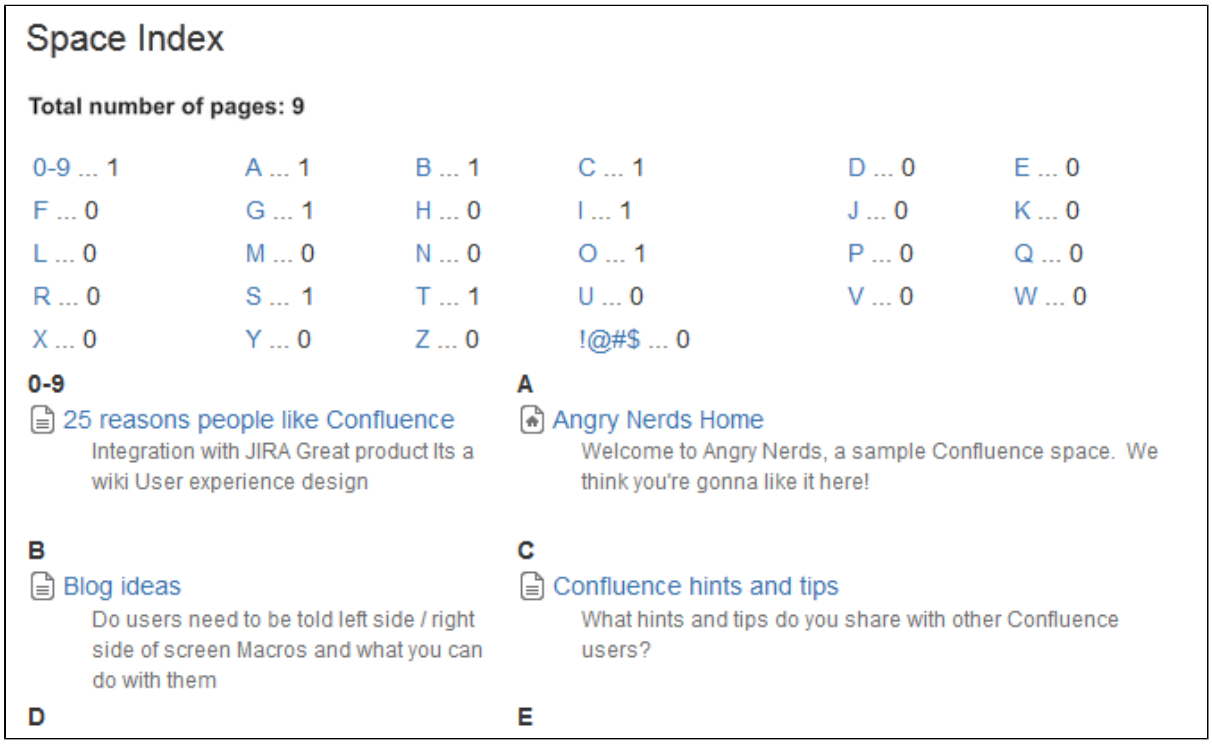

Limitations

The Page Index macro can be quite memory hungry in large spaces. To prevent it causing out of memory errors in your site, we:

- don't show page excerpts when there are more than 200 pages in the space,
- don't list any pages if there are more than 1000 in the space. This limit is configurable. System Administrators can use the `page.index.macro.max.pages` system property to reduce the number of pages displayed.

Change the macro parameters

This macro has no parameters.

Other ways to add this macro

Add this macro as you type

Type**{**followed by the start ofthe macro name, to see a list of macros.

Add this macro using wiki markup

This is useful when you want to add a macro outside the editor, for example as custom content in the sidebar, header or footer of a space.

Macro name:`index`

Macro body: None.

```
{index}
```


Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Page Properties Macro

When used together, the Page Properties and Page Properties Report macro can be used to create a table of information drawn from multiple pages.

These macros are great for:

- Decision registers
- Status reports
- Policy and procedure documentation

In short anywhere you have several distinct pieces of information you want to be able to roll-up, and cross-reference in a table on another page.

To use this macro, you need to add a **Page Properties** macro on one or more pages, and then you can add a **Page Properties Report** macro on another page, as shown below.

On this page:

- Add this macro to your page
- Page properties layout options
- Using multiple Page Properties macros on one page
- Change the macro parameters
- Limitations
- Other ways to add this macro

Related pages:

- Page Properties Report Macro
- Decisions Blueprint
- Product Requirements Blueprint

Screenshot: A project page with status information presented in a Page Properties macro.

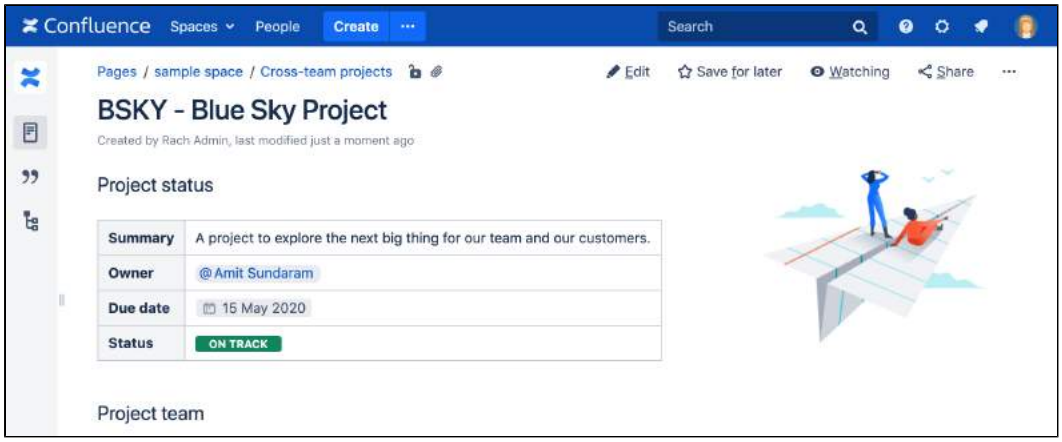

Screenshot: A page with a Page Properties Report macro showing status information from the Page Properties macro on several different project pages.

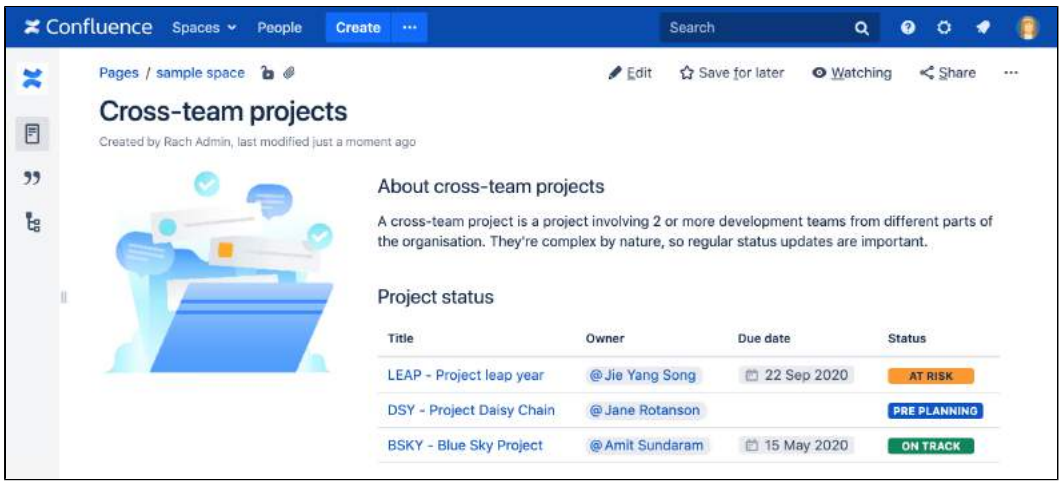

You can see examples of these two macros in action in the Decisions and Product Requirements blueprints.

Add this macro to your page

There are quite a few steps involved to set up this macro correctly, but once it's done, it's very easy to copy the macro to other pages, or add it to a template.

To add the Page Properties macro to a page:

1. From the editor toolbar, choose **Insert** + > **Other Macros**.
2. Choose **Page Properties** from the **Confluence content** category.
3. Choose Insert.
4. In the macro body create a two column table
5. In the left column list your 'properties' these will be the column headings in your report table.
6. Apply the **Heading column** style to the left column - this will indicate to the Page Properties Report macro that these are your properties.
7. In the right column list the value for each property these will populate the rows in your report table
8. Save your page.
9. Add a label to your page -you'll need to specify this label in the page properties report macro

Next you need to add the Page Properties Report macro to a different page.

Screenshot: The Page Properties macro on a page in the editor, with a vertical layout.

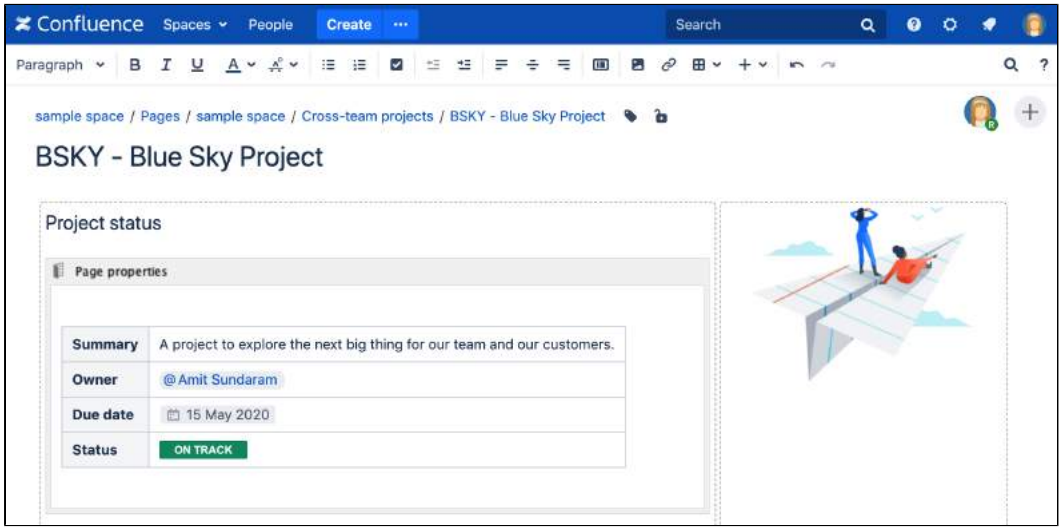

Page properties layout options

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Horizontal or vertical layout

You can choose to arrange the properties in your table vertically or horizontally. Just make sure you apply the **Heading row** or **Heading column** style to your properties, to tell the Page Properties Report macro, where to find them.

Here's an example of a horizontal layout.

Screenshot: The Page Properties macro on a page in the editor with a horizontal layout.

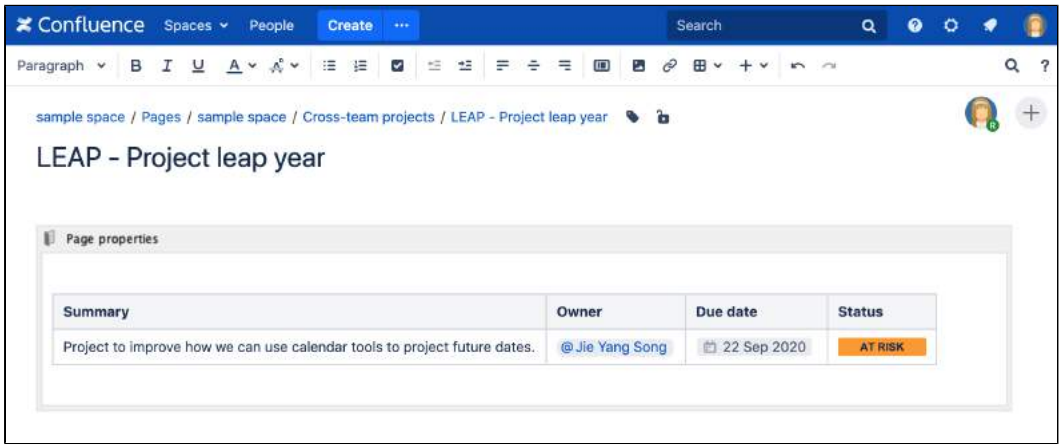

Hidden or visible

You can choose whether the contents of the Page Properties macro should be visible when someone views the page.

This is useful when the information isn't relevant to everyone. For example if you're using this macro to track when a policy was last reviewed and approved, you may only want that info to be visible on the page containing the Page Properties Report macro, not the page itself.

Screenshot: Configuring the Page Properties macro to be hidden.

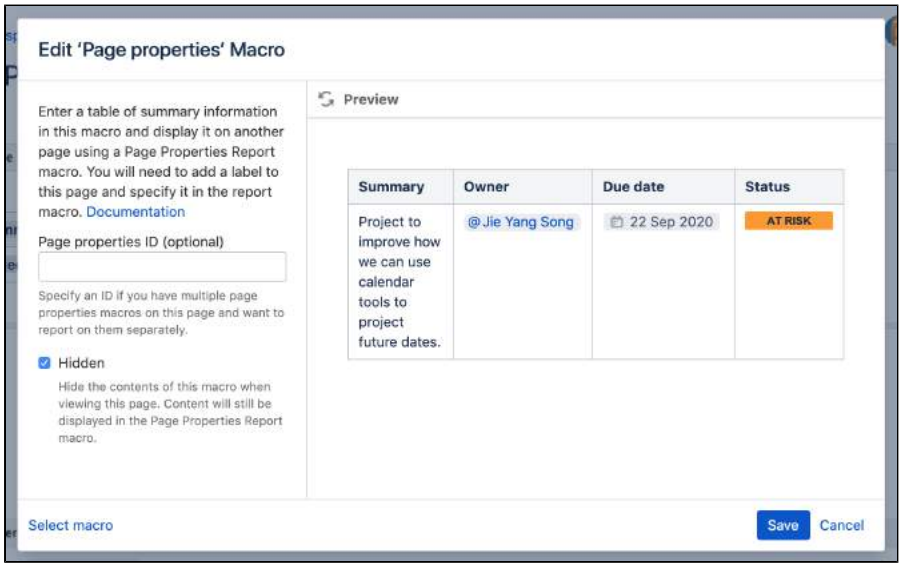

Using multiple Page Properties macros on one page

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.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

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:

1. Add a label to the page containing the Page Properties macros
2. Specify this label in the Page Properties Report macro

To show the contents of**selected**Page Properties macros in the report:

1. Add a label to the page containing the Page Properties macros
2. Specify an ID in the Page Properties macro that you want to report on
3. 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 thePage Properties macroin action in How to document product requirements in Confluence. This powerful macro lets you create a summary page thatpulls in information from multiple pages.

Change the macro parameters

Macro parameters are used to change the behaviour of a macro.

To change the macro parameters:

1. In the editor, click the macro placeholder and choose**Edit**.

2. Update the parameters as required then choose**Insert**.

Here's a list of the parameters available in this macro.

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.

Limitations

- You can't use macros in the left column as the data in this column is used to populate the column headings in your Page Properties Report macro.
- It is not possible to reference the metadatausing the metadata key from within the page, or anywhere else on a Confluence page.
- There's a known issue where the macro does not work correctly when placed inside an expand macro, which is inside a panel macro. See CONFSERVER-59594 GATHERING IMPACT .

Other ways to add this macro

Add this macro as you type

Type**{**followed by the start ofthe macro name, to see a list of macros.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

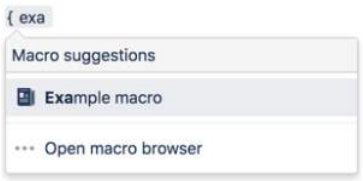

Add this macro using wiki markup

You can't use wiki markup to add this macro.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

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 another page. You can see examples of these two macros in action on theDecisionandProduct Requirementsblueprints.

This macro was previously known as the Details Summary macro.

On this page:

- Adding the Page Properties Report macro to a page
- Reporting on specific Page Properties macros
- CQL fields
- Macro display options
- Troubleshooting
- Limitations

Related pages:

- Page Properties Macro
- Decisions Blueprint
- Product Requirements Blueprint

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:

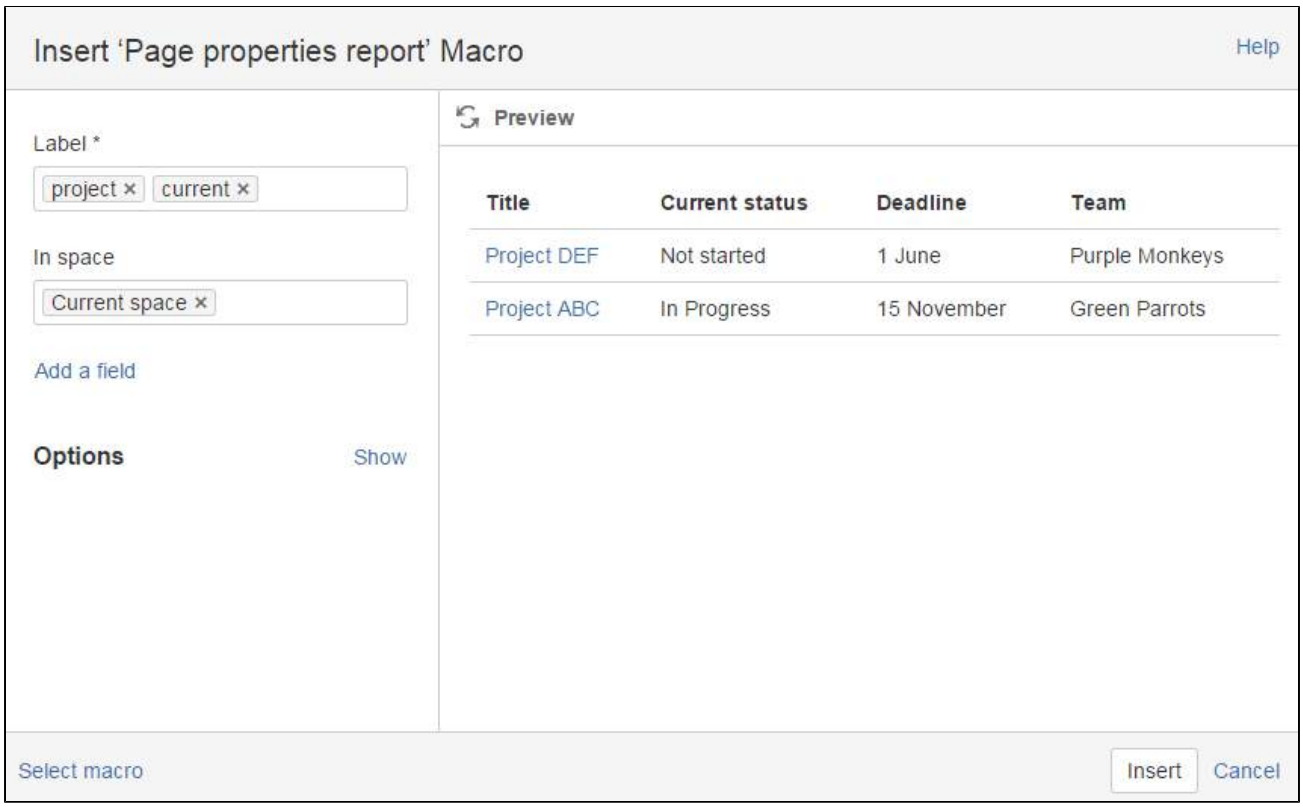

Reporting on specific Page Properties macros

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.

- 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:

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

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

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

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. The macro can display a maximum of 3000 pages. System administratorscan increase or decrease this limit. It is a good idea to use pagination, rather than listing all your pages in one go.
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.
- ThePage Properties macroson 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).

Limitations

You can enter a maximum of 60 labels in the macro browser.

The macro can display a maximum of 3000 pages. System administratorscan increase or decrease this limit using the`pagePropertiesReportContentRetrieverMaxResult`system property.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Page Tree Macro

Add the Page Tree macro to a page to display all or part of the hierarchy of pages in a space.

This macro is great for providing:

- quick and easy navigation to pages about a specific project, if the pages are organised together under one parent page
- a table of contents like experience to help people navigate a set of procedures
- an overview of pages in the current part of the hierarchy in documentation spaces.

On this page:

- Add this macro to your page
- Change the macro parameters
- Other ways to add this macro

Screenshot: The Page Tree macro in Confluence showing two levels of hierarchy.

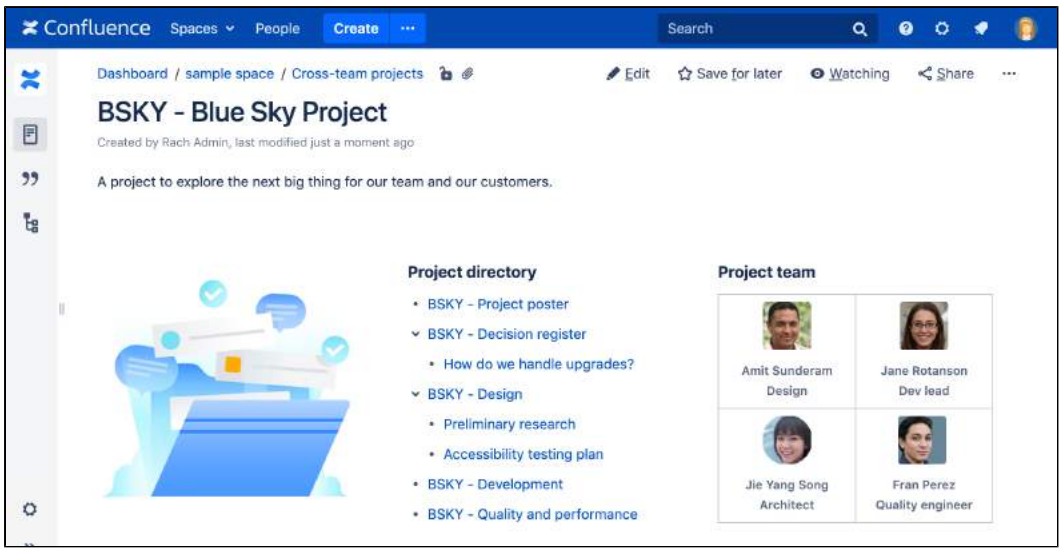

Add this macro to your page

To add the Page Tree macro to a page:

1. From the editor toolbar, choose **Insert** + > **Other Macros**.
2. Choose**Page Tree**from the**Confluence content**category.
3. Choose the number of page versions to display.
4. Choose**Insert**.

You can then publish your page to see the macro in action.

Screenshot: Configuring the Page Tree macro to display all pages below a specific page, in another space.

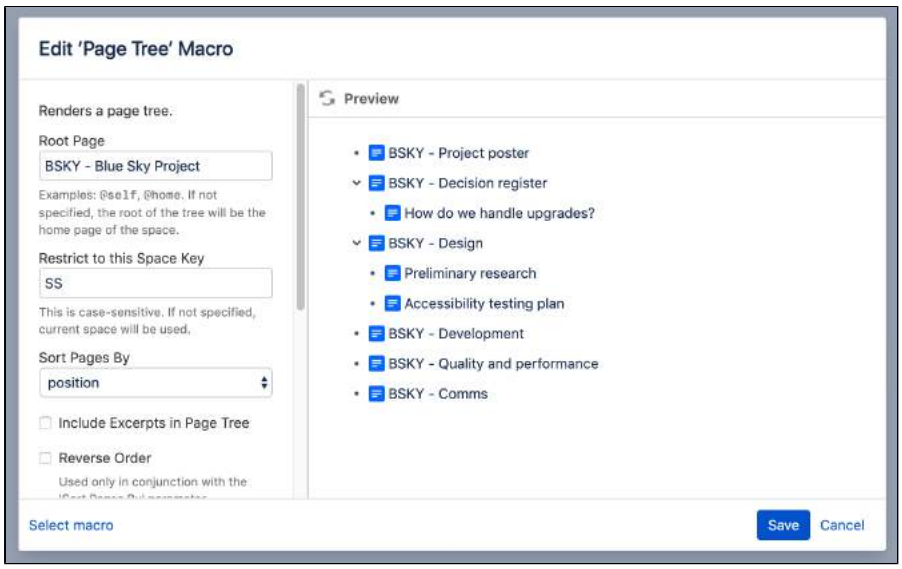

Change the macro parameters

Macro parameters are used to change the behaviour of a macro.

To change the macro parameters:

1. In the editor, click the macro placeholder and choose**Edit**.

2. Update the parameters as required then choose**Insert**.

Here's a list of the parameters available in this macro.

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.
Restrict to this space key	Current space	Enter a space name to show the page tree for a different space. Leave blank to use the current space. Set this parameter *before* the Root Page if you intend to specify a root page in another space.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

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.

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`).

Other ways to add this macro

Add this macro as you type

Type**{**followed by the start ofthe macro name, to see a list of macros.

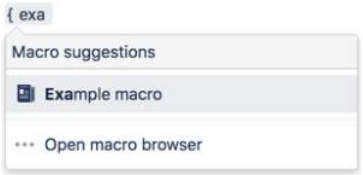

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Add this macro using wiki markup

This is useful when you want to add a macro outside the editor, for example as custom content in the sidebar, header or footer of a space.

Macro name: `pagetree`

Macro body: None.

```
{pagetree:root=Page
Name|sort=natural|excerpt=true|reverse=false|startDepth=3|expandCollapseAll=true|searchBox=true}
```


Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Page Tree Search Macro

Add the Page Tree Search macro to a page to provide users with a way to search for pages in a specific page hierarchy.

This macro is useful when you want to provide a way to search:

- one section of the current space, such as in a knowledge base
- a specific part of a page hierarchy, such as one project in a space that contains multiple projects.

After someone enters a keyword and clicks the Search button on this macro, the results are presentedon Confluence's advanced search screen.

On this page:

- Add this macro to your page
- Change the macro parameters
- Other ways to add this macro

Screenshot: Page Tree Search macro on a Confluence page.

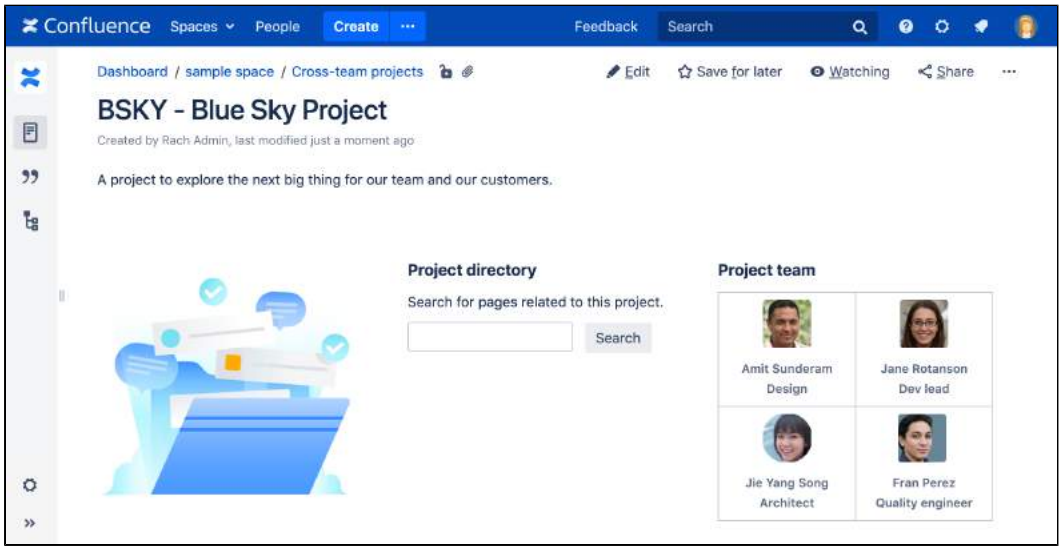

For a better search experience, check out the Livesearch Macro, or enable the integrated search field in the Page Tree Macro.

Add this macro to your page

To add the Page Tree Search macro to a page:

1. From the editor toolbar, choose **Insert** + > **Other Macros**.
2. Choose**Page Tree Search**from the**Confluence content**category.
3. Choose the number of page versions to display.
4. Choose**Insert**.

You can then publish your page to see the macro in action.

Change the macro parameters

Macro parameters are used to change the behaviour of a macro.

To change the macro parameters:

1. In the editor, click the macro placeholder and choose**Edit**.

2. Update the parameters as required then choose**Insert**.

Here's a list of the parameters available in this macro.

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.

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`).

Other ways to add this macro

Add this macro as you type

Type**{**followed by the start ofthe macro name, to see a list of macros.

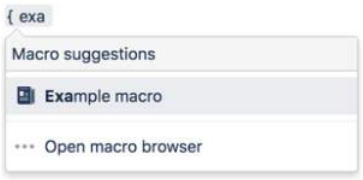

Add this macro using wiki markup

This is useful when you want to add a macro outside the editor, for example as custom content in the sidebar, header or footer of a space.

Macro name:`pagetreesearch`

Macro body: None.

```
{pagetreesearch:root=My page name}
```


Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Panel Macro

This macro is available in **Confluence Server and Data Center**. Learn about the macros available in Confluence Cloud.

On this page:

- Add this macro to your page
- Change the macro parameters
- Available colors
- Other ways to add this macro

Add the Panel macro to a page to format your text in a customizable coloured panel. It's similar to theInfo, Tip, Note, and Warning Macros except you have complete control over the border, background, title and text colours.

This is great for adding some visual interest to your pages. You can use panels in table cells and in page layouts, as in the example below.

Screenshot: page with a purple Panel macro containing a list of useful links.

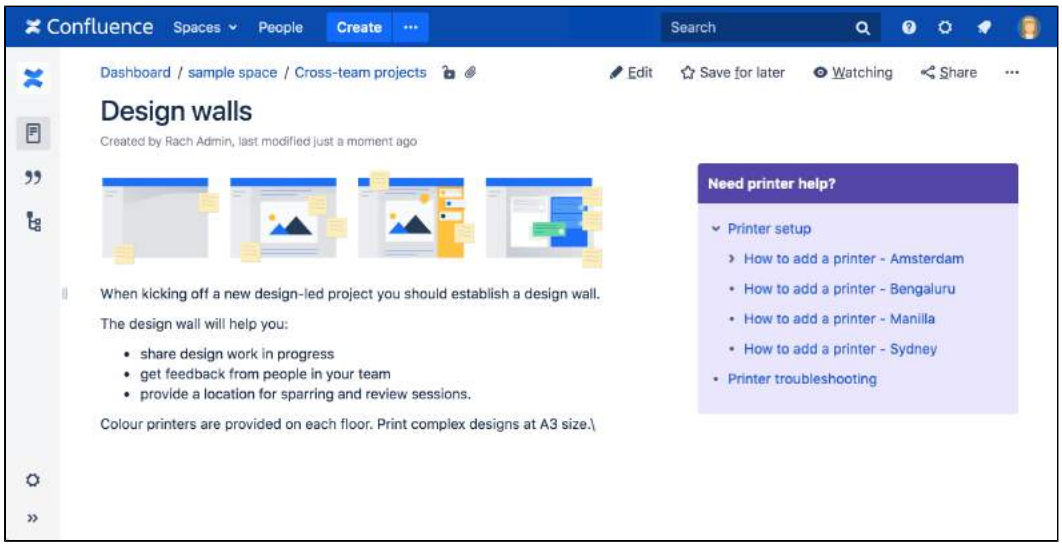

Add this macro to your page

To add the Panel Map macro to a page:

1. From the editor toolbar, choose **Insert** > **Other Macros**.
2. Choose**Panel**from the**Formatting**category.
3. Enter any parameters. Leave blank for a simple white panel with a grey border.
4. Choose**Insert**.

You can then publish your page to see the macro in action.

Screenshot: specifying a title, border, and background colour in the Panel macro.

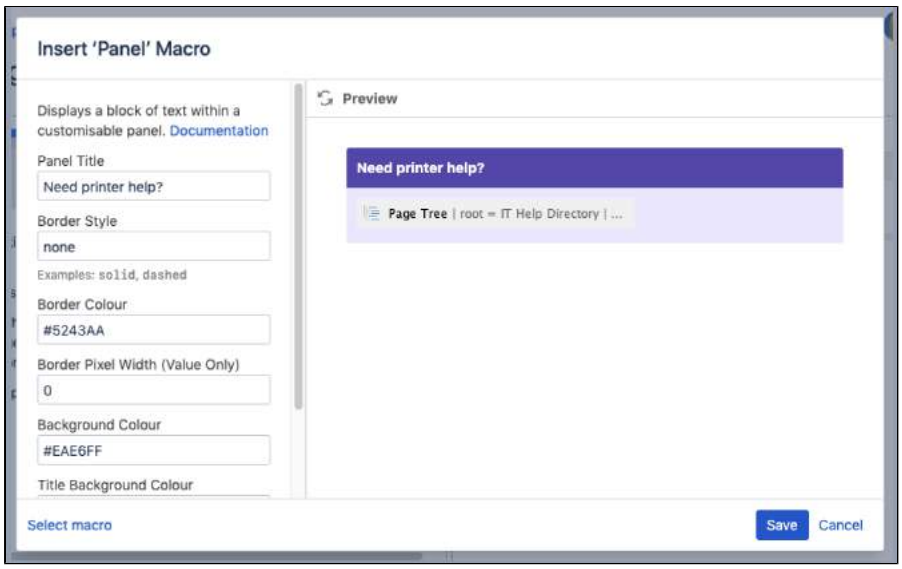

Change the macro parameters

Macro parameters are used to change the behaviour of a macro.

To change the macro parameters:

1. In the editor, click the macro placeholder and choose**Edit**.

2. Update the parameters as required then choose**Insert**.

Here's a list of the parameters available in this macro.

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 Color (`borderColor`)		The color of the panel's border. Colors can be specified as HTML color names or hexadecimal codes.
Border Pixel Width (Value Only) (`borderWidth`)		The width of the panel's border (in pixels).
Background Color (`bgColor`)		The background color of the panel. Colors can be specified as HTML color names or hexadecimal codes.
Title Background Color (`titleBGColor`)		The background color of the title row of the panel. Colors can be specified as HTML color names or hexadecimal codes.
Title Text Color (`titleColor`)		The color of the text in the title row of the panel. Colors can be specified as HTML color names or hexadecimal codes.

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`).

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Available colors

You can enter the HTML or X11 name for a color such as`Fuchsia`, `Teal`, or `MediumOrchid`, or you can enter hexadecimal values such as`#FF00FF`,`#008080`, and `#BA55D3`. You will need to include the # symbol when entering a hexadecimal value. SeeWeb colors for general information.

Other ways to add this macro

Add this macro as you type

Type**{**followed by the start ofthe macro name, to see a list of macros.

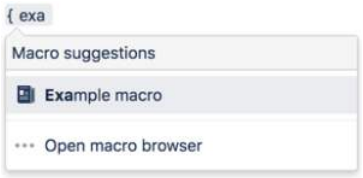

Add this macro using wiki markup

This is useful when you want to add a macro outside the editor, for example as custom content in the sidebar, header or footer of a space.

Macro name:`panel`

Macro body: Accepts rich text.

```
{panel:title=My
title|borderStyle=dashed|borderColor=blue|titleBGColor=#00a400|titleColor=white|bgColor=#72bc72}
A formatted panel
{panel}
```

Do more in Confluence

To further customize panels, check out these apps on the Atlassian Marketplace:

- Panelbox:Create a set of designed panelboxes to display identical topics in the same style, keeping your pages clear and easy to read
- Panels: Make your intranet more interactive with a customizable panels

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

PDF Macro

Add the PDF macro to a page to display the content of a PDF document. First attach the document to a Confluence page, then use the macro to display the document.

This is great for sharing presentations, design documents, whitepapers, and other visual data.

This macro displays your file in a viewer, with next and back buttons, rather than showing a simple preview.

On this page:

- Add this macro to your page
- Change the macro parameters
- Other ways to add this macro

Screenshot: A page with a PDF macro displaying an A4 PDF document.

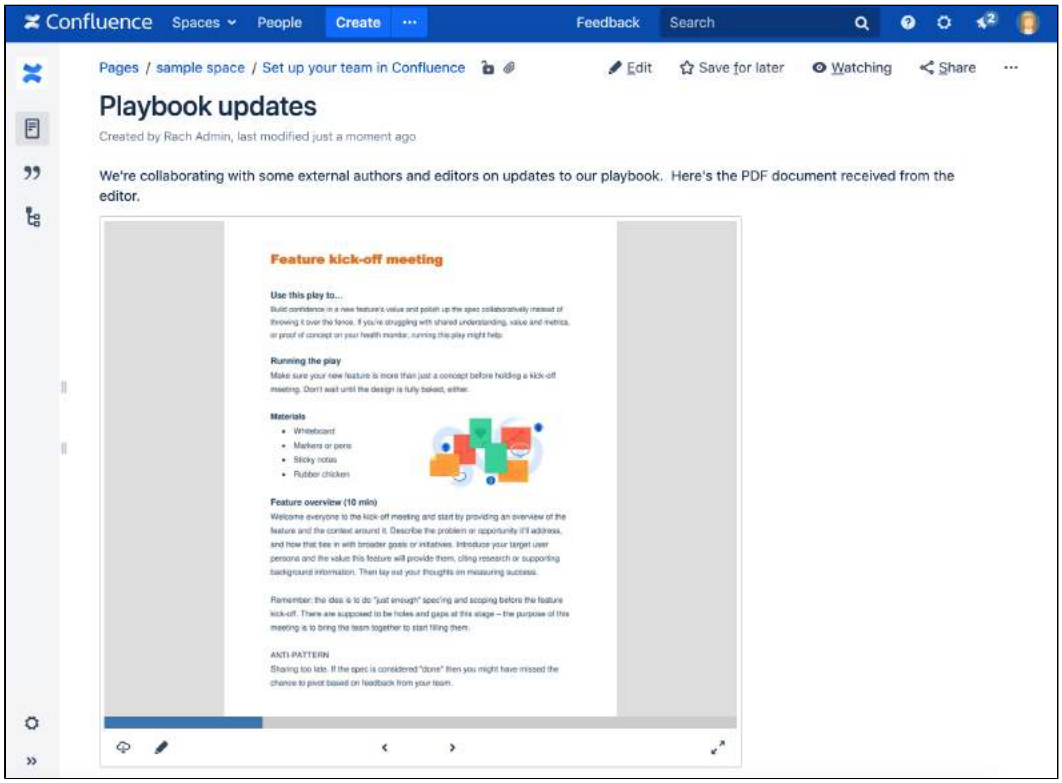

Add this macro to your page

To add the PDF macro to a page:

1. Upload the PDF file to your page, then publish the page. SeeUpload Filesto learn how to do this.
2. From the editor toolbar, choose **Insert** + > **Other Macros**.
3. Choose**PDF**from the**Confluence content**category.
4. Select the attached file you want to display.
5. Choose**Insert**.

You can then publish your page to see the macro in action.

Screenshot: Configuring the PDF macro in the macro browser.

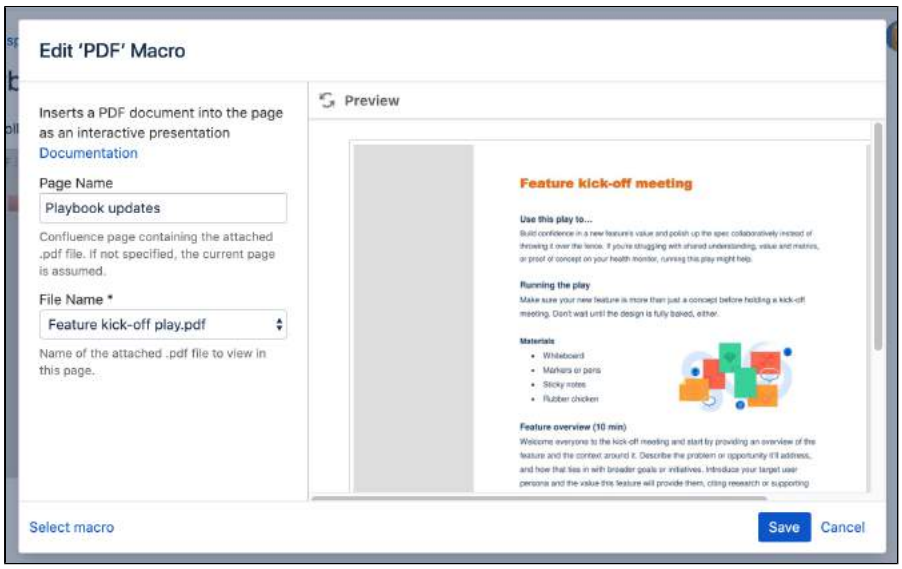

Change the macro parameters

Macro parameters are used to change the behaviour of a macro.

To change the macro parameters:

1. In the editor, click the macro placeholder and choose**Edit**.

2. Update the parameters as required then choose**Insert**.

Here's a list of the parameters available in this macro.

Parameter	Default	Description
Page Name (`page`)	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 (`name`)	none	The file name of the PDF document to be displayed. The document must be attached to a page on your Confluence site.

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`).

Other ways to add this macro

Add this macro as you type

Type**{**followed by the start ofthe macro name, to see a list of macros.

Add this macro using wiki markup

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

This is useful when you want to add a macro outside the editor, for example as custom content in the sidebar, header or footer of a space.

Macro name:`viewpdf`

Macro body: None.

```
{viewpdf:page=Docs|name=My document.pdf}
```


Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Popular Labels Macro

Add the Popular Labels macro to a page tohighlight 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.

This is great for surfacing labels that might be trending in your site, for example in a knowledge base.

On this page:

- Add this macro to your page
- Change the macro parameters
- Other ways to add this macro

Screenshot: The Popular Labels macro showing a heatmap of the 20 most popular labels.

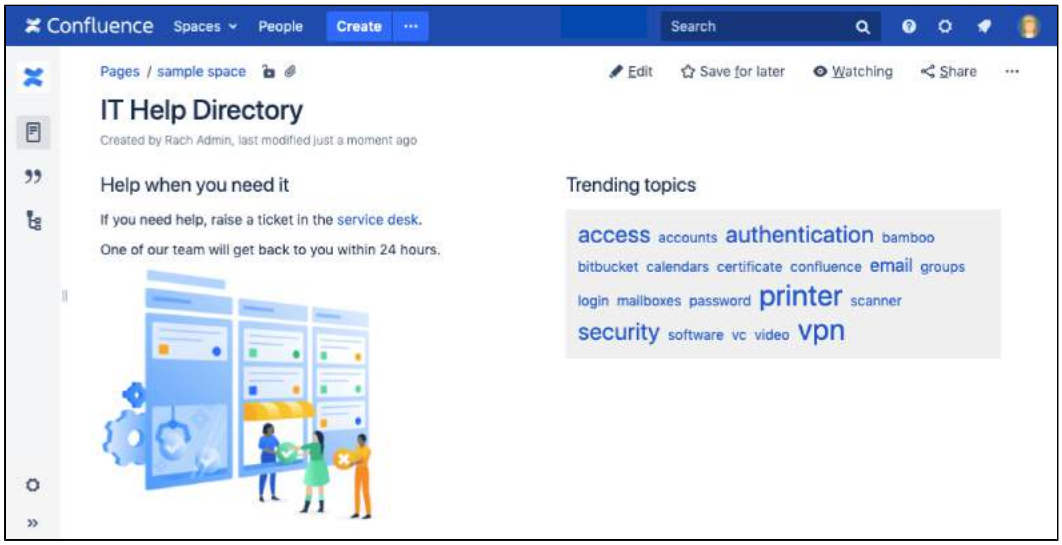

For general information about using labels in Confluence, see Add, Remove and Search for Labels.

Add this macro to your page

To add the Panel Map macro to a page:

1. From the editor toolbar, choose **Insert** + > **Other Macros**.
2. Choose**Popular Labels**from the**Reporting**category.
3. Enter any parameters and choose how you want the labels to display.
4. Choose**Insert**.

You can then publish your page to see the macro in action.

Screenshot: specifying a space key, number of labels, and heatmap display style.

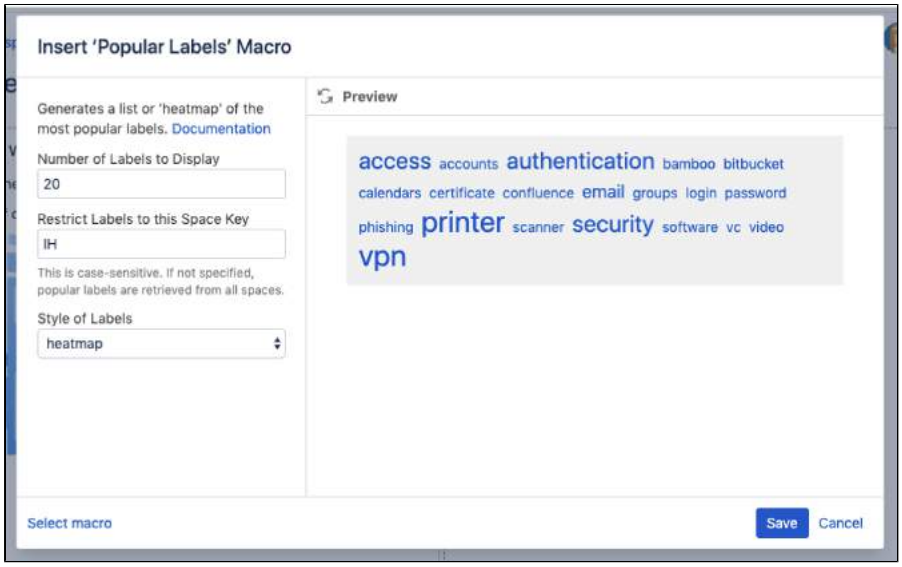

Change the macro parameters

Macro parameters are used to change the behaviour of a macro.

To change the macro parameters:

1. In the editor, click the macro placeholder and choose**Edit**.

2. Update the parameters as required then choose**Insert**.

Here's a list of the parameters available in this macro.

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.

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`).

Other ways to add this macro

Add this macro as you type

Type**{**followed by the start ofthe macro name, to see a list of macros.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

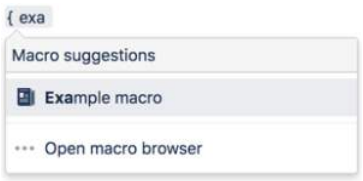

Add this macro using wiki markup

This is useful when you want to add a macro outside the editor, for example as custom content in the sidebar, header or footer of a space.

Macro name:`popular-labels`

Macro body: None.

```
{popular-labels:style=heatmap|count=20|spaceKey=ds}
```


Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Profile Picture Macro

Add the Profile Picture macro to a page to display a user's profile picture.

It's great for putting a face to the name on team and project pages.

On this page:

- Interacting with the macro on a page
- Change the macro parameters
- Other ways to add this macro

Screenshot: Page with several Profile Picture macros.

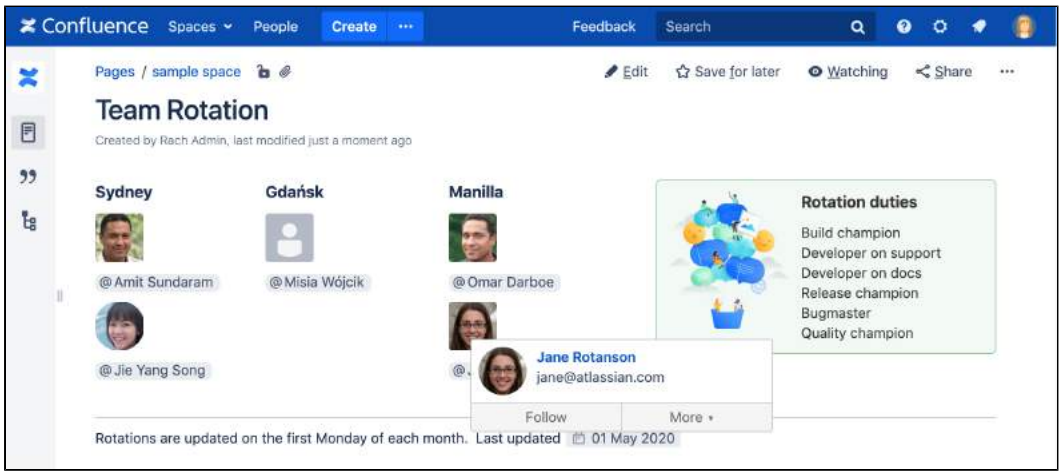

Add this macro to your page

To add the Profile Picture macro to a page:

1. From the editor toolbar, choose **Insert** > **Other Macros**.
2. Choose**Profile Picture**from the**Confluence content**category.
3. Enter a username.
4. Choose**Insert**.

You can then publish your page to see the macro in action.

Screenshot: specifying a user in the macro browser.

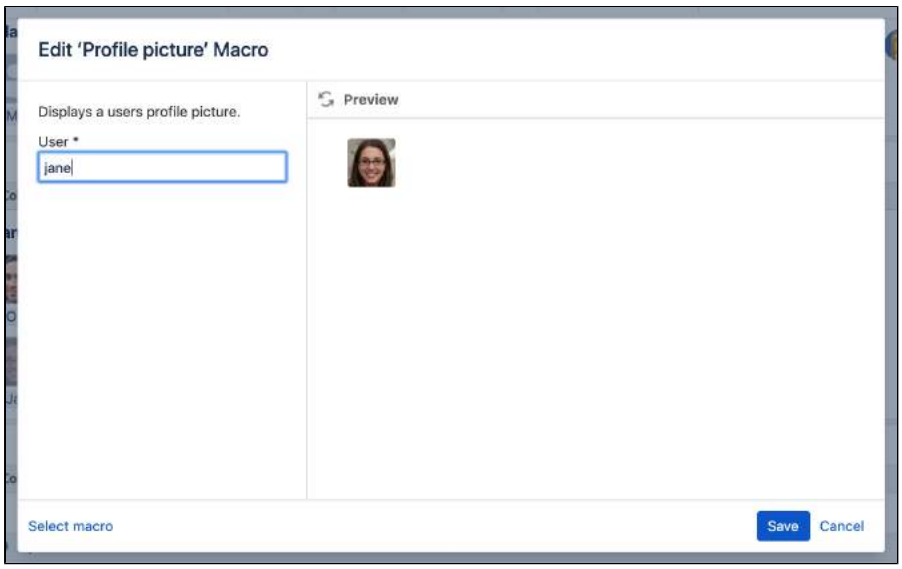

Interacting with the macro on a page

When viewing a page, 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**View User Profile**to see theprofilefor the user.

Screenshot: The user profile macro displaying the profile picture for a user.

Change the macro parameters

Macro parameters are used to change the behaviour of a macro.

To change the macro parameters:

1. In the editor, click the macro placeholder and choose**Edit**.

2. Update the parameters as required then choose**Insert**.

Here's a list of the parameters available in this macro.

Parameter	Default	Description

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

User	none	The username of the person you want to display a profile picture for.

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`).

Other ways to add this macro

Add this macro as you type

Type**{**followed by the start ofthe macro name, to see a list of macros.

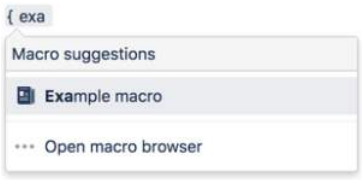

Add this macro using wiki markup

You can't use wiki markup to add this macro.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Recently Updated Dashboard Macro

Add the Recently Updated Dashboard macro to a page to show a list of pages, blogs, files, and comments that have been created or edited recently.

This is a **legacy** macro, and was previously used to display recently updated content on the dashboard. It is very similar to theRecently Updated Macro, except that it has a tabbed view, that lets you switch between all updates, updates from your favorite spaces, updates from your network (the people you follow), or particular space categories.

On this page:

- Add this macro to your page
- To add the Recently Updated Dashboard macro to a page:
- Change the macro parameters
- Other ways to add this macro

Screenshot: The Recently Updated Dashboard macro My Spaces tab showing a personalised view of recent updates in the site.

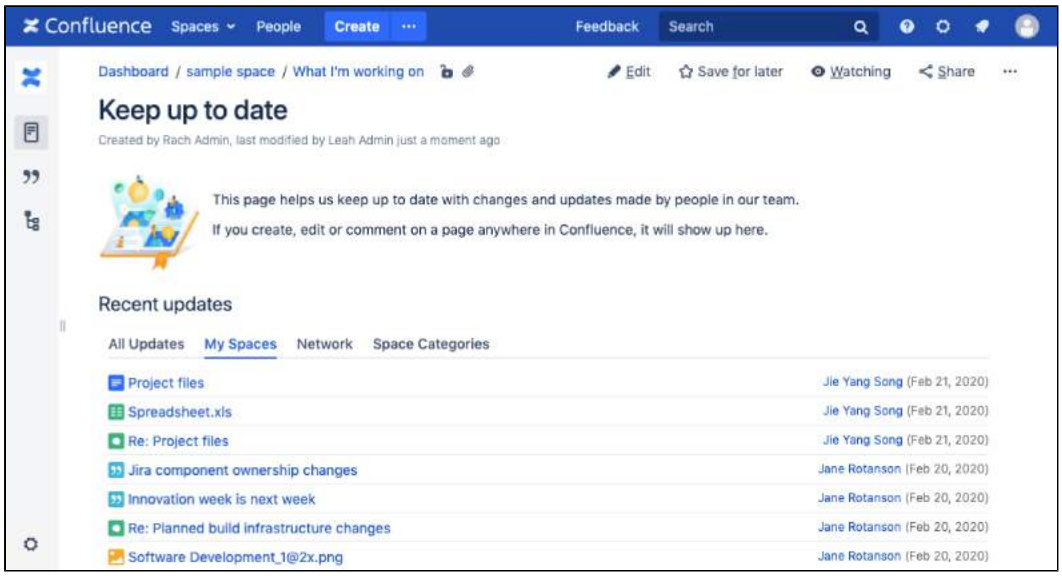

Add this macro to your page

To add the Recently Updated Dashboard macro to a page:

1. From the editor toolbar, choose **Insert** + > **Other Macros**.
2. Choose**Recently Updated Dashboard**from the**Confluence content**category.
3. Enter any parameters.
4. Choose**Insert**.

You can then publish your page to see the macro in action.

Screenshot: Configuring the Recently Updated Dashboard macro to show updates from specific people.

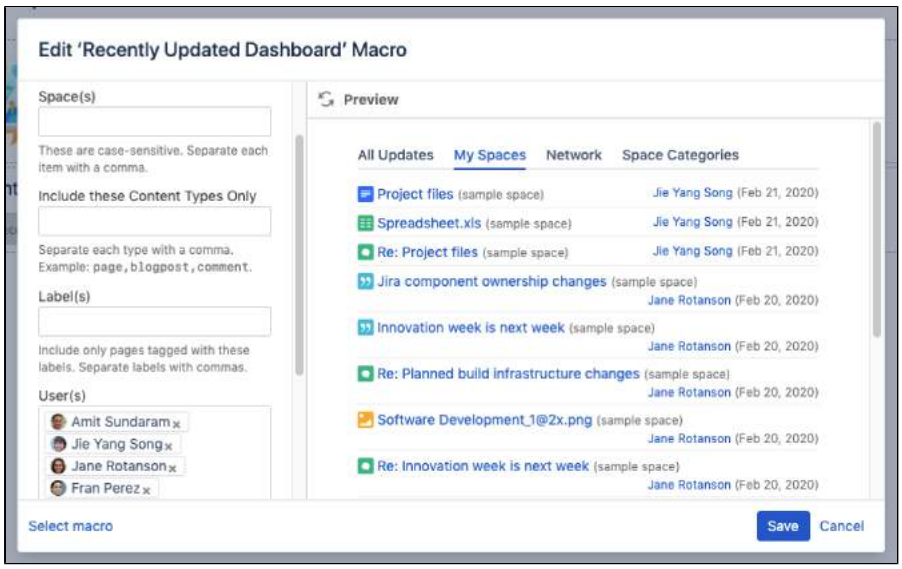

Change the macro parameters

Macro parameters are used to change the behaviour of a macro.

To change the macro parameters:

1. In the editor, click the macro placeholder and choose**Edit**.

2. Update the parameters as required then choose**Insert**.

Here's a list of the parameters available in this macro.

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`, `spaced esc`, `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.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

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.

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`).

Other ways to add this macro

Add this macro as you type

Type**{**followed by the start ofthe macro name, to see a list of macros.

Add this macro using wiki markup

This is useful when you want to add a macro outside the editor, for example as custom content in the sidebar, header or footer of a space.

Macro name:`recently-updated-dashboard`

Macro body: None.

```
{recently-updated-dashboard:spaces=ds|users=admin|width=50%|showProfilePic=true|labels=choc|types=page}
```


Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Recently Updated Macro

Add the Recently Updated macro to a page to show a list of pages, blogs, files, and comments that have been created or edited recently.

This is great for:

- project landing or information pages
- team space home pages

It's very flexible, you can limit the list to specific people, spaces, types of content, and more.

On this page:

- Add this macro to your page
- Change the macro parameters
- Other ways to add this macro

Screenshot: project landing page showing recently created and updated pages.

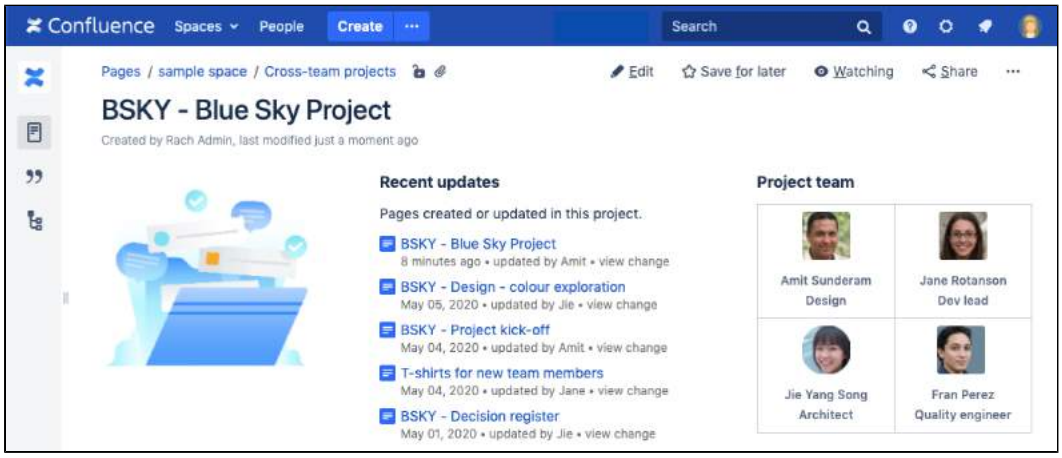

Add this macro to your page

To add the Recently Updated macro to a page:

1. From the editor toolbar, choose **Insert** > **Other Macros**.
2. Choose**Recently Updated**from the**Confluence content**category.
3. Enter any parameters.
4. Choose**Insert**.

You can then publish your page to see the macro in action.

Screenshot: configuring the Recently Updated macro to show updates from specific people in a particular space.

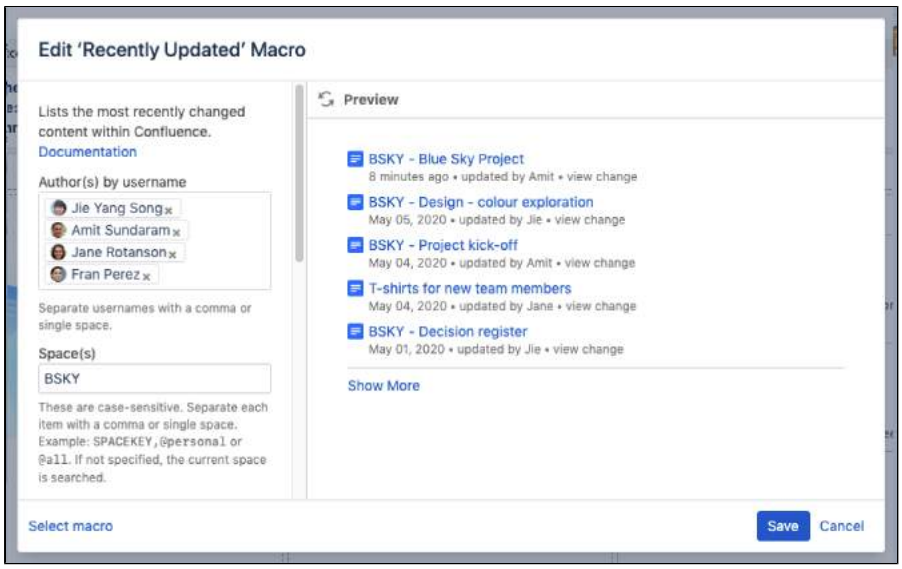

Change the macro parameters

Macro parameters are used to change the behaviour of a macro.

To change the macro parameters:

1. In the editor, click the macro placeholder and choose**Edit**.

2. Update the parameters as required then choose**Insert**.

Here's a list of the parameters available in this macro.

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.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

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 favorite. • `@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 labeled 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.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

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.

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`).

Other ways to add this macro

Add this macro as you type

Type**{**followed by the start ofthe macro name, to see a list of macros.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

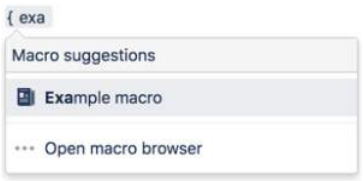

Add this macro using wiki markup

This is useful when you want to add a macro outside the editor, for example as custom content in the sidebar, header or footer of a space.

Macro name:`recently-updated`

Macro body: None.

```
{recently-updated:spaces=ds|author=admin|max=10|hideHeading=true|width=50%
|theme=sidebar|showProfilePic=true|labels=choc|types=page}
```


Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Recently Used Labels Macro

Add theRecently Used Labels macro to a page to display a list of labels that have recently been applied to a page, blog post or attached file.

This is great for keeping track of when new topics are added to things like:

- knowledge base articles
- process and procedure documentation
- project documentation.

You can confine the search to the current space, or the entire site.

On this page:

- Add this macro to your page
- Change the macro parameters
- Other ways to add this macro

Screenshot: a page using the Recently Used Labels macro to show the list of topics applied to knowledge base articles recently.

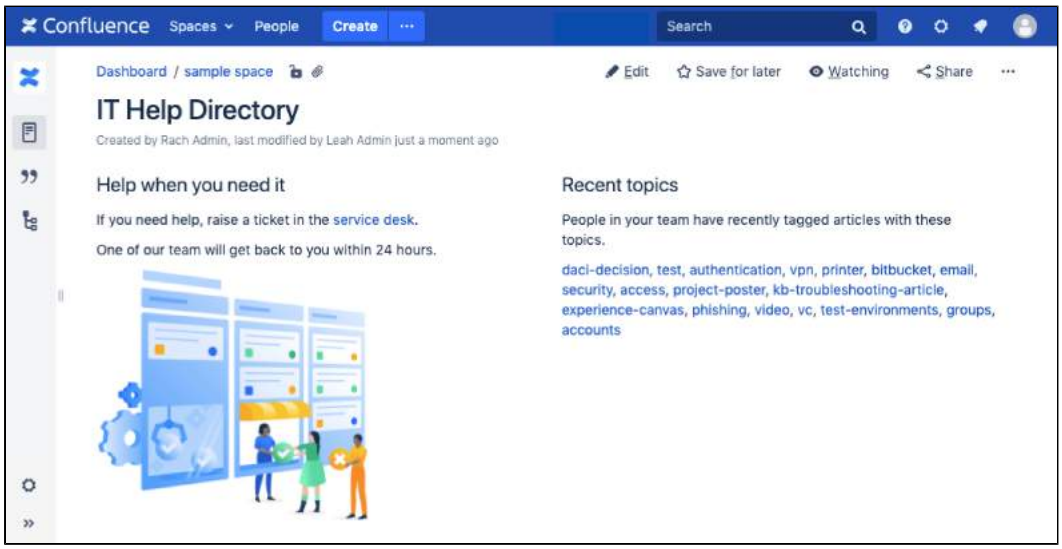

You can also choose to configure this macro to show more detail, including the page titles and information about the user who added the label.

Add this macro to your page

To add the Recently Used Labels macro to a page:

1. From the editor toolbar, choose **Insert** + > **Other Macros**.
2. Choose**Recently Used Labels**from the**Confluence content**category.
3. Enter any parameters.
4. Choose**Insert**.

You can then publish your page to see the macro in action.

Screenshot: configuring the Recently Used Labels macro to show 35 recently applied labels.

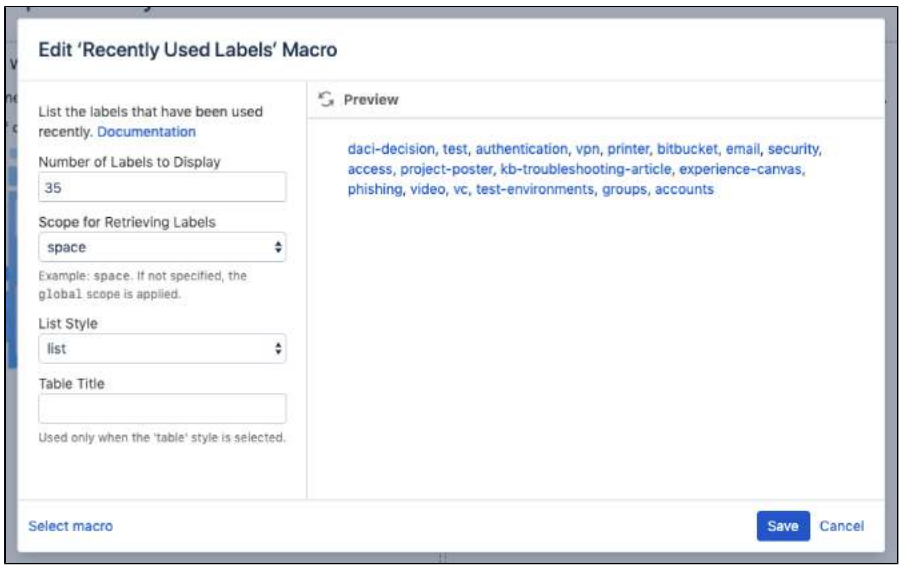

Change the macro parameters

Macro parameters are used to change the behaviour of a macro.

To change the macro parameters:

1. In the editor, click the macro placeholder and choose**Edit**.

2. Update the parameters as required then choose**Insert**.

Here's a list of the parameters available in this macro.

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**.

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`).

Other ways to add this macro

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Add this macro as you type

Type **{** followed by the start ofthe macro name, to see a list of macros.

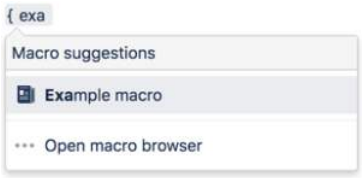

Add this macro using wiki markup

This is useful when you want to add a macro outside the editor, for example as custom content in the sidebar, header or footer of a space.

Macro name: `recently-used-labels`

Macro body: None.

```
{recently-used-labels:title=My title|scope=space|style=table|count=20}
```


Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Related Labels Macro

Add the Related Labels macro to give people a way to navigate to related content in your site.

It works by finding labels that pages have in common. It takes the label used on the current page, looks for other pages that use this label, and then lists any other labels that are found on those pages.

This macro is great for:

- Knowledge base articles
- Process and procedure documentation.

Basically any situation where you want a map of related pages in your site. The list updates automatically, as labels are added and removed from pages over time.

On this page:

- Add this macro to your page
- Change the macro parameters
- Other ways to add this macro

Screenshot: the Related Labels macro guiding people to related articles in a knowledge base.

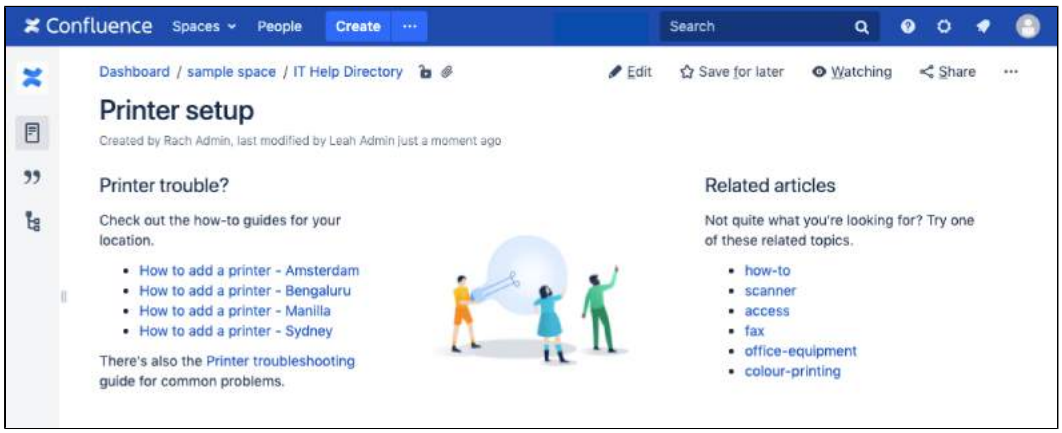

Add this macro to your page

To add the Related Labels macro to a page:

1. From the editor toolbar, choose **Insert** + > **Other Macros**.
2. Choose**Related Labels**from the**Confluence content**category.
3. Enter specific labels the macros should look for, or leave blank to use any labels on the page.
4. Choose**Insert**.

You can then publish your page to see the macro in action.

Screenshot: configuring the Related Labels macro in the macro browser.

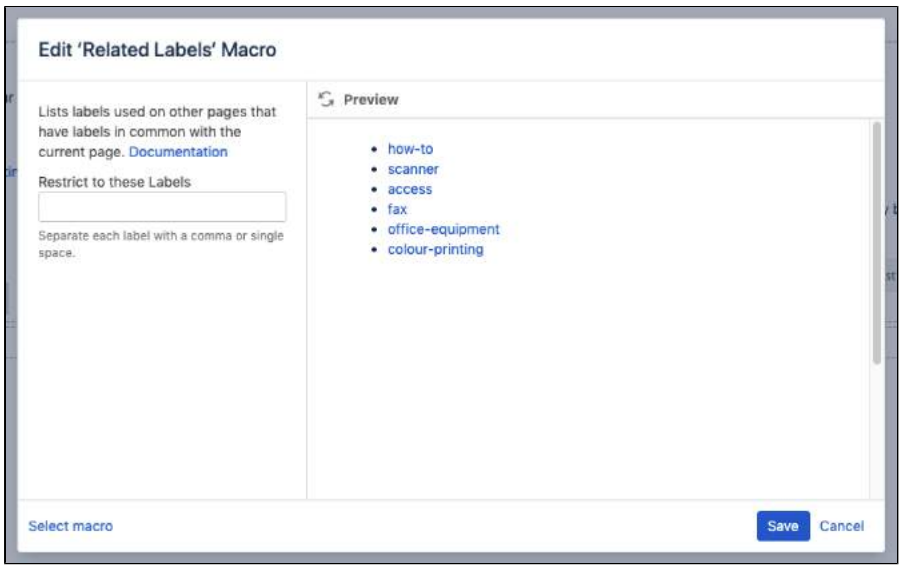

Change the macro parameters

Macro parameters are used to change the behaviour of a macro.

To change the macro parameters:

1. In the editor, click the macro placeholder and choose**Edit**.

2. Update the parameters as required then choose**Insert**.

Here's a list of the parameters available in this macro.

Parameter	Default	Description
Restrict to these Labels (`labels`)	none	Specify the labels for which you want to view related labels. For example, `documentation,how-to`.

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`).

Other ways to add this macro

Add this macro as you type

Type**{**followed by the start ofthe macro name, to see a list of macros.

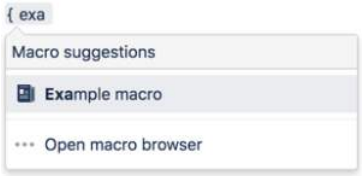

Add this macro using wiki markup

This is useful when you want to add a macro outside the editor, for example as custom content in the sidebar, header or footer of a space.

Macro name:`related-labels`

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Macro body: None.

```
{related-labels:labels=choc,cake}
```


Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Roadmap Planner Macro

Add the Roadmap Planner macro to a page to create a simple, visual timeline that's useful for planning projects, software releases and much more.

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.

On this page:

- Add this macro to your page
- Editing your Roadmap
- Change the macro parameters
- Other ways to add this macro
- Notes

Screenshot: Page with a Roadmap Planner macro showing the various stages of a project.

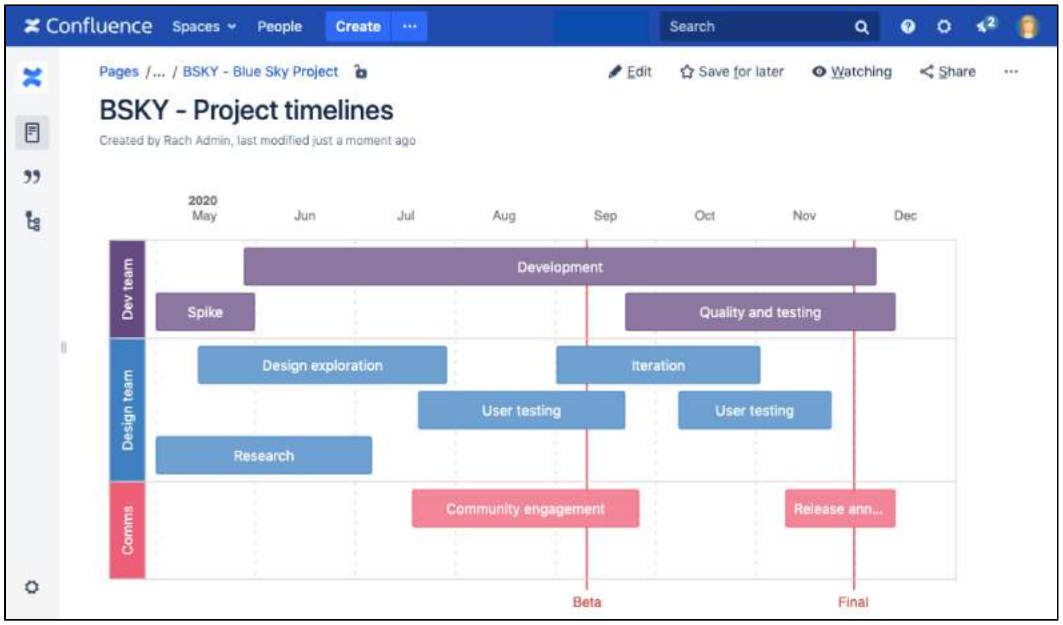

Add this macro to your page

To add the Roadmap Planner macro to a page:

1. From the editor toolbar, choose **Insert** + > **Other Macros**.
2. Choose**Roadmap Planner**from the**Visuals and images**category.
3. **Add**lanes, bars and markers as described below.
4. Choose**Insert**.

You can then publish your page to see the macro in action.

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 colors 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: Adding bars and markets to the Roadmap Planner macro in the editor.

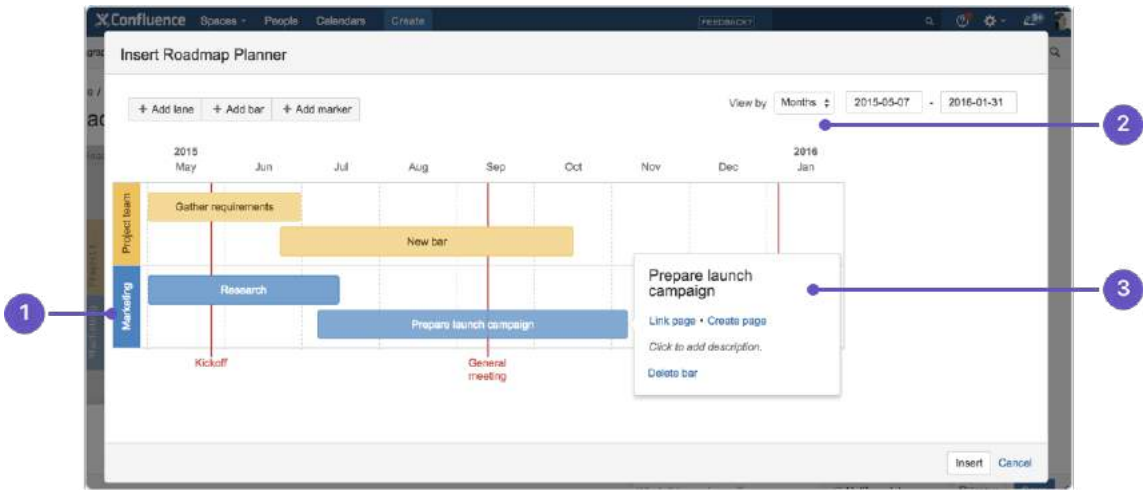

1. **Add lanes:** to differentiate your teams or streams of work.
2. **Date range:** display the plan by weeks or months.
3. **Add bar info:** to provide more info or link a bar to a page.

Change the macro 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.

Other ways to add this macro

Add this macro as you type

Type**{**followed by the start ofthe macro name, to see a list of macros.

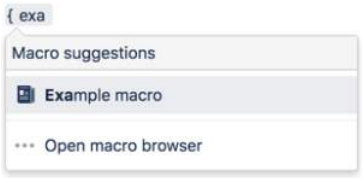

Add this macro using wiki markup

You can't use wiki markup to add this macro.

Notes

The Roadmap macro was previously available fromThe 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.

Do more with Confluence

To extend Confluence's roadmap capability, check out these apps on the Atlassian Marketplace:

- Live Roadmap: Build your Roadmap and keep it alive connecting it to your Jira
- ProductPlan for Confluence Server: Embed your roadmap in Confluence to keep your team aligned around high-level goals

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

RSS Feed Macro

Add the RSS Feed macro to a page to display the contents of an external or internal RSS feed. For example, to display blog posts or to list recently updated pages in a space, you can create an internal feed in theFeed Builder, then render it using this macro.

This is a **legacy** macro, and is often disabled by Confluence administrators for security reasons.

On this page:

- Security considerations
- Add this macro to your page
- Change the macro parameters
- How up to date is the feed?
- What happens to a page containing a disallowed URL?
- Authentication
- Enable or disable the RSS Feed macro
- Other ways to add this macro

Security considerations

The RSS Feed macro may be disabled by your Confluence administrator. Also, your Confluence administrator can define alist of trusted URLs. You will see an error message on the Confluence page, if the included URL is not in the allowlist.

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.

Add this macro to your page

To add the RSS Feed macro to a page:

1. From the editor toolbar, choose **Insert** > **Other Macros**.
2. Choose**RSS Feed**from the**External content**category.
3. Enter the RSS feed URL.
4. Choose**Insert**.

You can then publish your page to see the macro in action.

Change the macro parameters

Macro parameters are used to change the behaviour of a macro.

To change the macro parameters:

1. In the editor, click the macro placeholder and choose**Edit**.

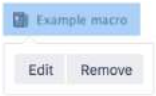

2. Update the parameters as required then choose**Insert**.

Here's a list of the parameters available in this macro.

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.

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`).

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 theCache pluginonto your Confluence site.

What happens to a page containing a disallowed URL?

Your Confluence Administrator can set up an allowlist 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 allowlist. 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 allowlist.

Here is an example of the error message, including the link shown only to Confluence Administrators:

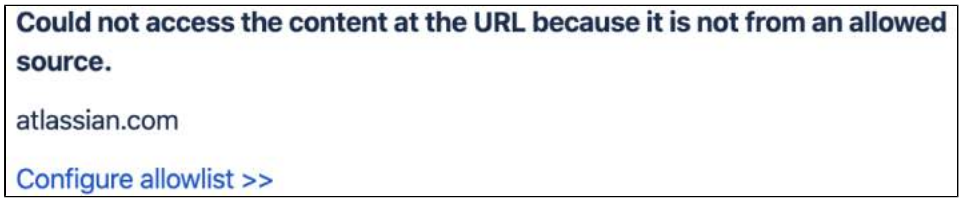

Here is an example of the error message, but without the link.

Authentication

Private feeds from external sites

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

RSS feeds which require authentication cannot be accessed using the RSS Feed macro.

Accessing internal HTTPS feeds

This applies only to Confluence instances which haveenabled 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 theSSL guideto 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 theIE 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.

Enable or disable the RSS Feed macro

To enable or disable the RSS Feed macro:

1. Go to > **Manage apps**.
2. Select**System**from the drop down and search for the**Confluence HTML Macros** system app.
3. Expand the listing and enable or disable the**rss (rss-xhtml)**module.

Other ways to add this macro

Add this macro as you type

Type**{**followed by the start ofthe macro name, to see a list of macros.

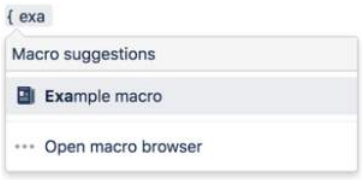

Add this macro using wiki markup

This is useful when you want to add a macro outside the editor, for example as custom content in the sidebar, header or footer of a space.

Macro name: rss

Macro body: None.

```
{rss:max=10|showTitlesOnly=true|url=http://myblog.com/feed|titleBar=false}
```


Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Search Results Macro

Add the Search Results macro to a page to display the results of a pre-defined search.

Add this macro to your page

> **We ended support for this macro in Confluence 7.0**
> The macro no longer appears in the macro browser and can't be added to a page.
> Any macro already on a page will still work.

Change the macro parameters

Macro parameters are used to change the behaviour of a macro.

To change the macro parameters:

1. In the editor, click the macro placeholder and choose**Edit**.
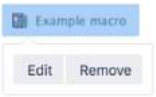

2. Update the parameters as required then choose**Insert**.

Here's a list of the parameters available in this macro.

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.

Other ways to add this macro

Add this macro using wiki markup

This is useful when you want to add a macro outside the editor, for example as custom content in the sidebar, header or footer of a space.

Macro name:`search`

Macro body: None.

```
{search:lastModified=3w|query=choc|contributor=admin|maxLimit=10|type=page|spacekey=ds}
```


Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Section Macro

Add the Section macro to a page to organise your content in sections and columns. This macro is used in conjunction with the Column macro, and provides more flexibility than page layouts.

This macro is great for situations where:

- you need more than three columns, or
- you need your columns to be a specific width.

On this page:

- Add this macro to your page
- Change the macro parameters
- Other ways to add this macro

Screenshot: page with a four column layout using the Section and Column macros.

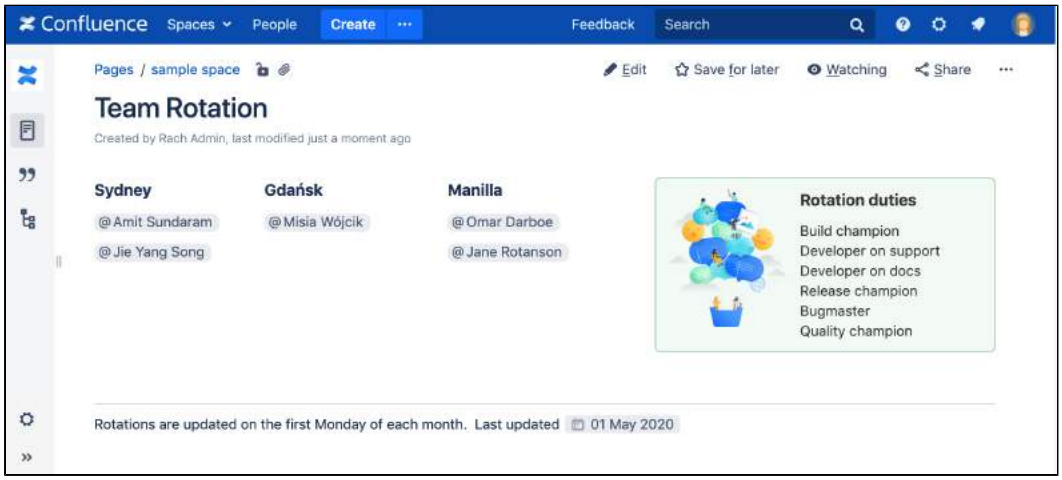

Want a simpler way to lay out your page? Try a page layout instead.

Add this macro to your page

To add the Section macro to a page:

1. From the editor toolbar, choose **Insert** > **Other Macros**.
2. Choose**Section**from the**Formatting**category.
3. Choose**Insert**.

You can then start typing into the macro body, add some Column macros, then publish your page to see the macro in action.

Screenshot: section and column macros in the editor

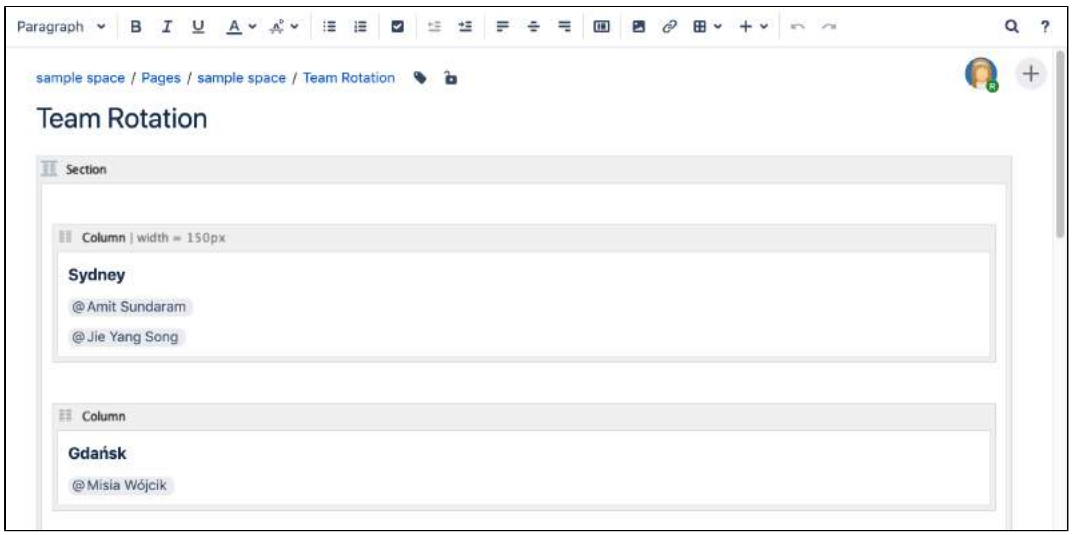

Change the macro parameters

Macro parameters are used to change the behaviour of a macro.

To change the macro parameters:

1. In the editor, click the macro placeholder and choose**Edit**.

2. Update the parameters as required then choose**Insert**.

Here's a list of the parameters available in this macro.

Parameter	Default	Description
Show Border (`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.

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`).

Other ways to add this macro

Add this macro as you type

Type**{**followed by the start ofthe macro name, to see a list of macros.

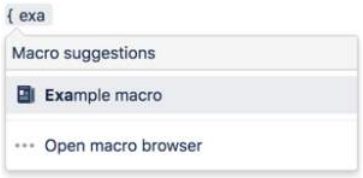

Add this macro using wiki markup

This is useful when you want to add a macro outside the editor, for example as custom content in the sidebar, header or footer of a space.

Macro name:`section`

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Macro body: Rich text, consisting of one or more Column macros.

```
{section:border=true}
{column:width=100px}
This is the content of *column 1*.
{column}
{column}
This is the content of *column 2*.
{column}
{section}
```


Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Space Attachments Macro

Add the Space Attachments macro to a page to display a list of all files attached to pages in the current space, or another space.

This is great for:

- design assets
- frequently used files, like forms
- image libraries.

It shows details of the file and the includes a link to the page the file is attached to. Filter the list by label or file extension to find the file you're looking for more easily.

On this page:

- Add this macro to your page
- To add the Space Attachments macro to a page:
- Change the macro parameters
- Other ways to add this macro

Screenshot: A page containing the Space Attachments macro, to provide quick access to files used in a project.

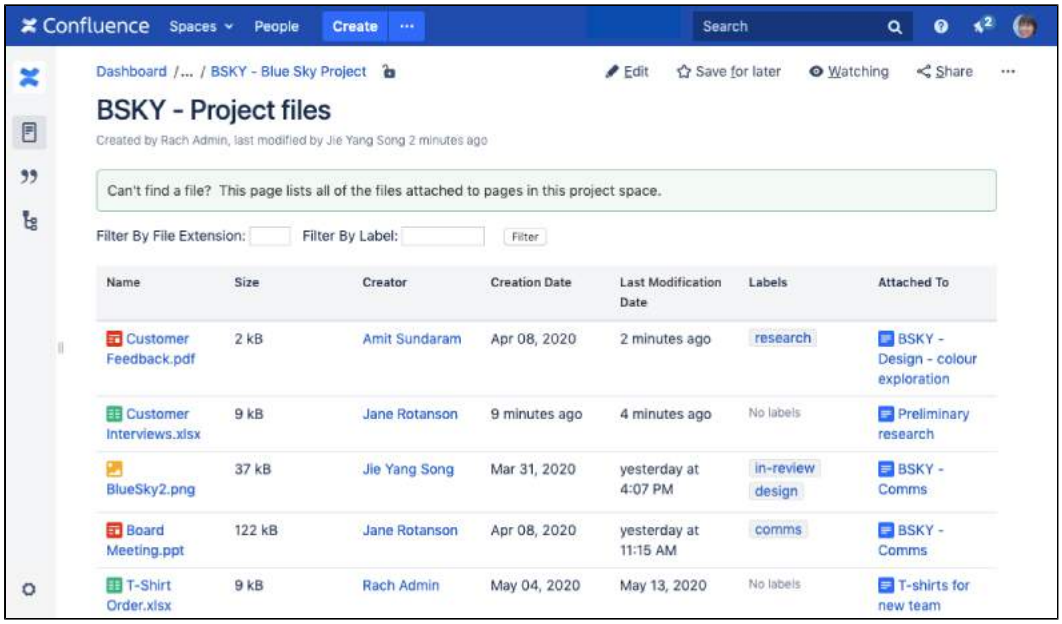

Add this macro to your page

To add the Space Attachments macro to a page:

1. From the editor toolbar, choose **Insert** > **Other Macros**.
2. Choose**Space Attachments**from the**Confluence content**category.
3. Enter any parameters.
4. Choose**Insert**.

You can then publish your page to see the macro in action.

Screenshot: Configuring the Space Attachments macro to only show files attached to pages in a specific space.

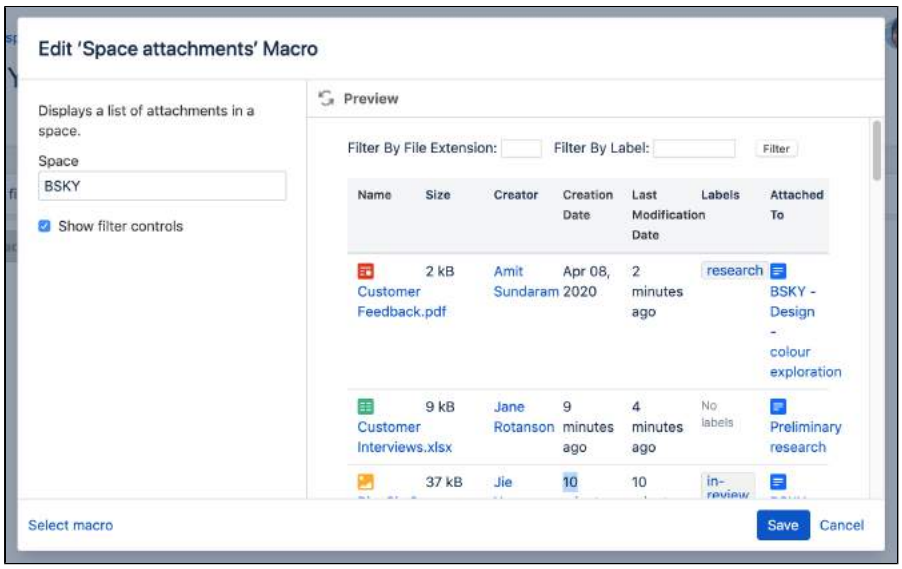

Change the macro parameters

Macro parameters are used to change the behaviour of a macro.

To change the macro parameters:

1. In the editor, click the macro placeholder and choose**Edit**.

2. Update the parameters as required then choose**Insert**.

Here's a list of the parameters available in this macro.

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.

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`).

Other ways to add this macro

Add this macro as you type

Type**{**followed by the start ofthe macro name, to see a list of macros.

Add this macro using wiki markup

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

This is useful when you want to add a macro outside the editor, for example as custom content in the sidebar, header or footer of a space.

Macro name: `space-attachments`

Macro body: None.

```
{space-attachments:showFilter=false|space=ds}
```


Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Space Details Macro

The Space Details macro displays the details of a Confluence space, including the space name, description, and more.

Add this macro to your page

We ended support for this macro in Confluence 7.0
The macro no longer appears in the macro browser and can't be added to a page.
Any macro already on a page will still work.

Change the macro parameters

Macro parameters are used to change the behaviour of a macro.

To change the macro parameters:

1. In the editor, click the macro placeholder and choose**Edit**.

2. Update the parameters as required then choose**Insert**.

Here's a list of the parameters available in this macro.

Parameter	Default	Description
Width of Table (`width`)	100%	The width of the space details table, specified as a percentage (%) of the page width.

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`).

Other ways to add this macro

Add this macro using wiki markup

This is useful when you want to add a macro outside the editor, for example as custom content in the sidebar, header or footer of a space.

Macro name:`space-details`

Macro body: None.

```
{space-details:width=50%}
```

Spaces List Macro

Add the Spaces List macro to a page to display a list of spaces on a page.

This is **legacy** macro. It was previously used in the dashboard, and has been known to cause issues large sites. Although limited, this macro is great if you want to:

- highlight spaces that have been created in the last 7 days
- provide a list of spaces, that can be filtered by space category.

On this page:

- Add this macro to your page
- To add the Spaces List macro to a page:
- Change the macro parameters
- Troubleshooting
- Other ways to add this macro

Screenshot: Page with a Spaces List macro showing cross-team spaces.

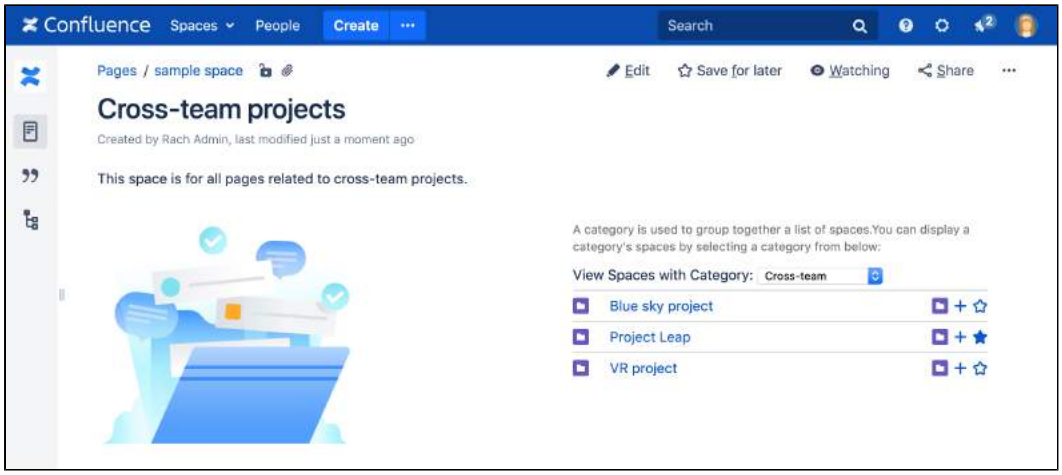

This macro works best with space categories. See Use Labels to Categorize Spaces to find out how to categorize spaces in your site.

Add this macro to your page

To add the Spaces List macro to a page:

1. From the editor toolbar, choose **Insert** > **Other Macros**.
2. Choose**Spaces List**from the**Confluence content**category.
3. Enter any parameters.
4. Choose**Insert**.

You can then publish your page to see the macro in action.

Screenshot: Configuring the spaces list macro to show spaces by category.

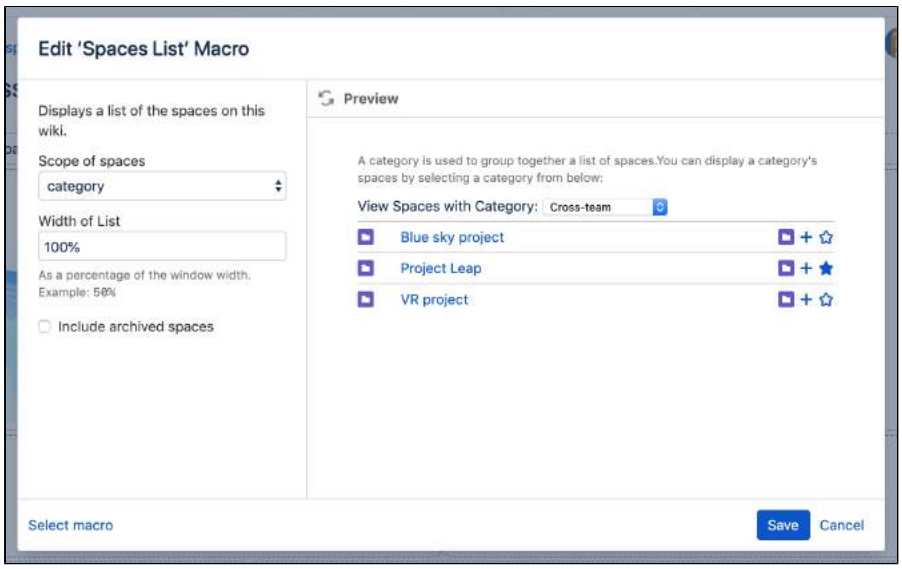

Change the macro parameters

Macro parameters are used to change the behaviour of a macro.

To change the macro parameters:

1. In the editor, click the macro placeholder and choose**Edit**.

2. Update the parameters as required then choose**Insert**.

Here's a list of the parameters available in this macro.

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. • **favorite** 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.
Include archived spaces (`includeArchivedSpaces`)	no	False	Include spaces that have been archived in the list. These are excluded by default.

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`).

Troubleshooting

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

This macro is known to cause issues in sites with large number of spaces. See

CONFSERVER-59804 - Using the Spaces List macro in a Confluence Page may overload the database
GATHERING IMPACT

.

This macro is provided by the **Dashboard Macros** system app.Disable the **Spaces** module if you want to prevent people using this macro.

Other ways to add this macro

Add this macro as you type

Type**{**followed by the start ofthe macro name, to see a list of macros.

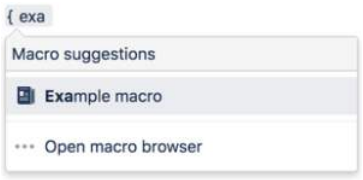

Add this macro using wiki markup

This is useful when you want to add a macro outside the editor, for example as custom content in the sidebar, header or footer of a space.

Macro name:`spaces`

Macro body: None.

```
{spaces:favourite|width=80%}
```


Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Status Macro

Add the Status macro to a page to highlight the status of a project, task, or item with a colored lozenge (rounded box).

This macro is great for indicating:

- items that are problematic or removed
- a successful project
- tasks that are in progress.

You can choose a solid or light background color for the lozenge and the text that appears inside the lozenge. The macro displays its current statusdirectly in the editor.

On this page:

- Using the Status macro
- To add the Status macro to a page:
- Change the macro parameters
- Other ways to add this macro

Screenshot: Status macros in a table highlighting important tasks, actions, and items.

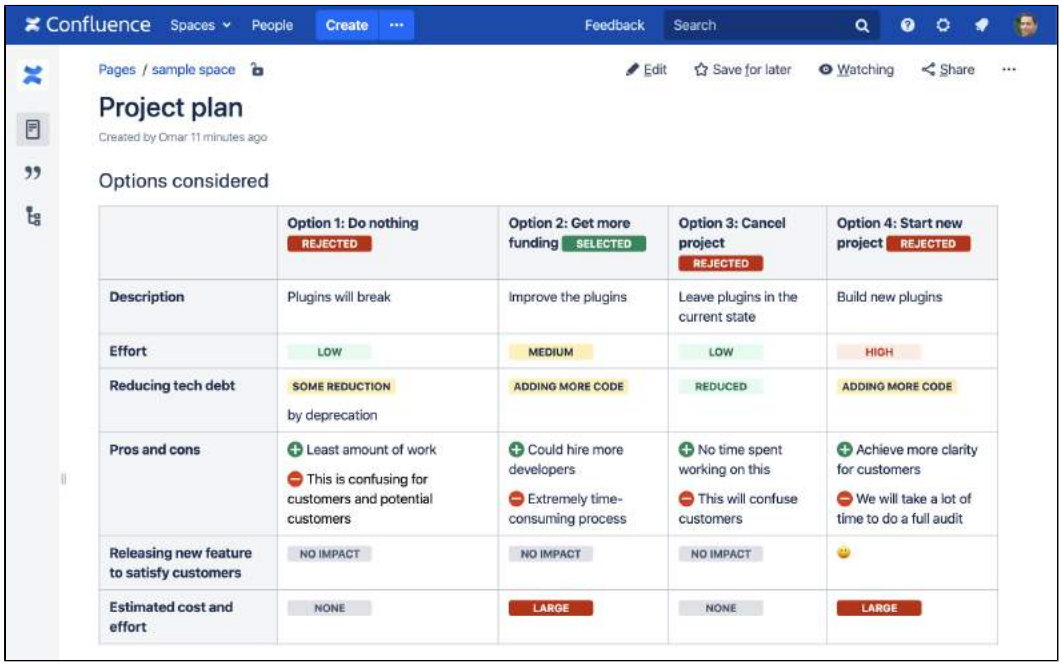

See the Status macro put to excellent use in How to build a release planning page in Confluence.

Using the Status macro

To add the Status macro to a page:

1. From the editor toolbar, choose **Insert** > **Other Macros**.
2. Choose**Status**from the**Confluence content**category.
3. Enter a status and choose a color.
4. Choose**Insert**.

You can then publish your page to see the macro in action.

Screenshot: Selecting lozenge colour in the Status macro.

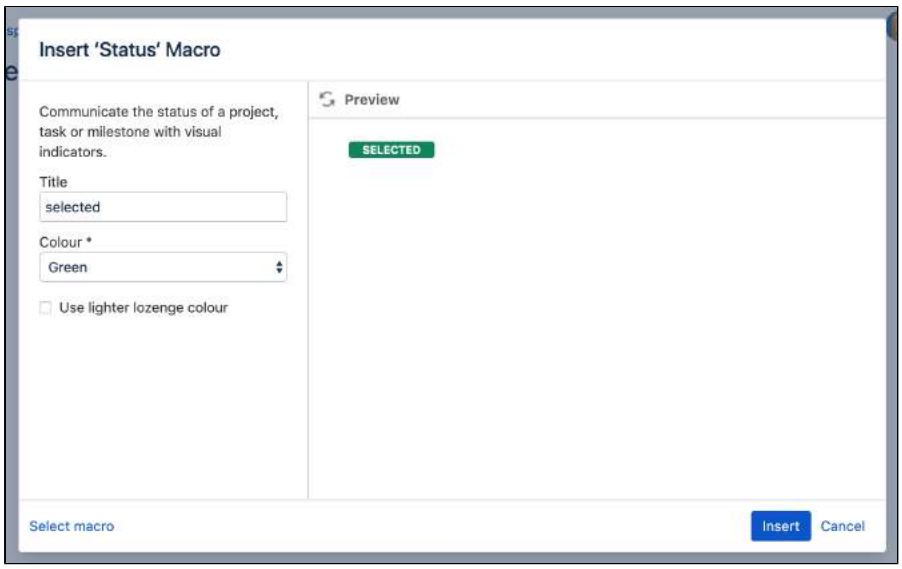

Change the macro parameters

Macro parameters are used to change the behaviour of a macro.

To change the macro parameters:

1. In the editor, click the macro placeholder and choose**Edit**.
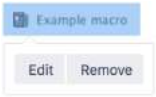

2. Update the parameters as required then choose**Insert**.

Here's a list of the parameters available in this macro.

Parameter	Default	Description
Color (`colour`)	Grey	The color of the lozenge. The following colors are available: **Grey** , **Red**, **Yellow**, **Green** and **Blue**.
Title (`title`)	The color that you select.	The text that will appear inside the lozenge. If you do not specify any text, the title will be the color of the lozenge, that is 'Grey', 'Red', 'Yellow', 'Green' or 'Blue'.
Use lighter lozenge color (`subtle`)	False	The lozenge background color and text color. The default lozenge style is a solid background color with white text. Select this parameter to use a lighter lozenge background color withcolored text.

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`).

Other ways to add this macro

Add this macro as you type

Type**{**followed by the start ofthe macro name, to see a list of macros.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

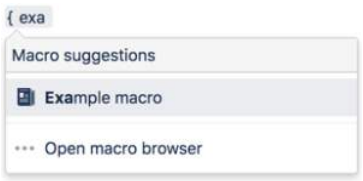

Add this macro using wiki markup

This is useful when you want to add a macro outside the editor, for example as custom content in the sidebar, header or footer of a space.

Macro name:`status`

Macro body: None.

```
{status:colour=Green|title=On track|subtle=true}
```

Do more with Confluence

For more customizable status macros check out these apps on the Atlassian Marketplace:

- Build Status Tracker for Confluence: Provide visibility on build status from Bamboo or Jenkins on your Confluence pages
- Spectrum Formatting Macros:Show a page status as a draft, as outdated, as action required or as official

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Table of Contents Macro

Add the Table of Contents macro to a page to help your readers skip directly to the information theyre looking for.

This macro is great for situations where:

- you have a large page with lots of information
- you want to build your headings into a neat table of contents.

This macro is popular because it helps you navigate lengthy pages. The macro organises your table of contents bynesting Heading 2 under Heading 1,and indenting progressively, like in the image below.

On this page:

- Using the Table of Contents macro
- Change the macro parameters
- Examples
- Other ways to add this macro
- Notes and known issues

Screenshot: Table of Contents macro neatly organises lengthy blog content.

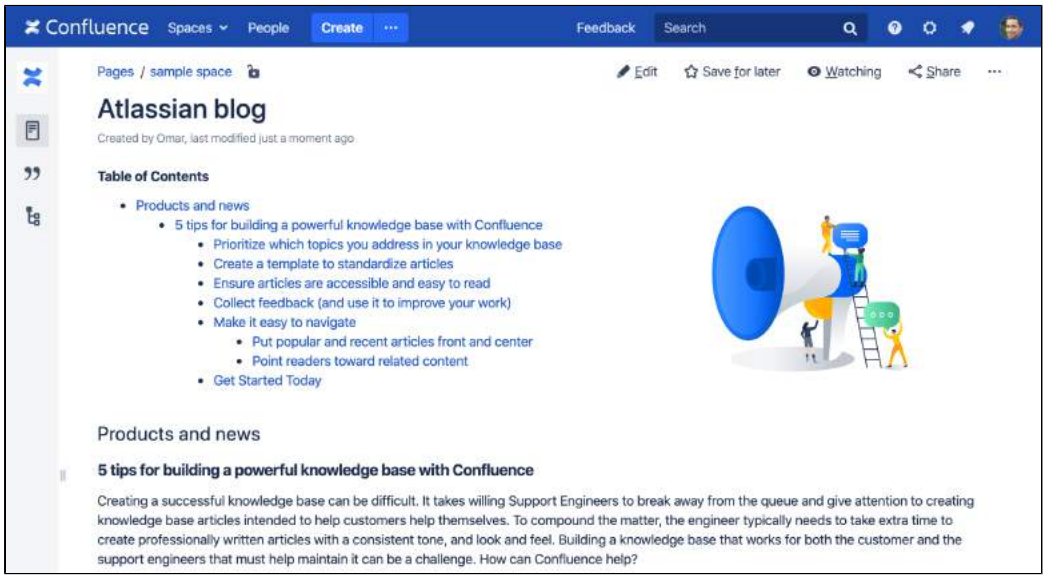

Using the Table of Contents macro

To add the Table of Contents macro to a page:

1. From the editor toolbar, choose **Insert** + > **Other Macros**.
2. Choose**Table of Contents**from the**Confluence content**category.
3. Enter any parameters.
4. Choose**Insert**.

You can then publish your page to see the macro in action.

Screenshot: Entering parameters for the Table of Contents macro.

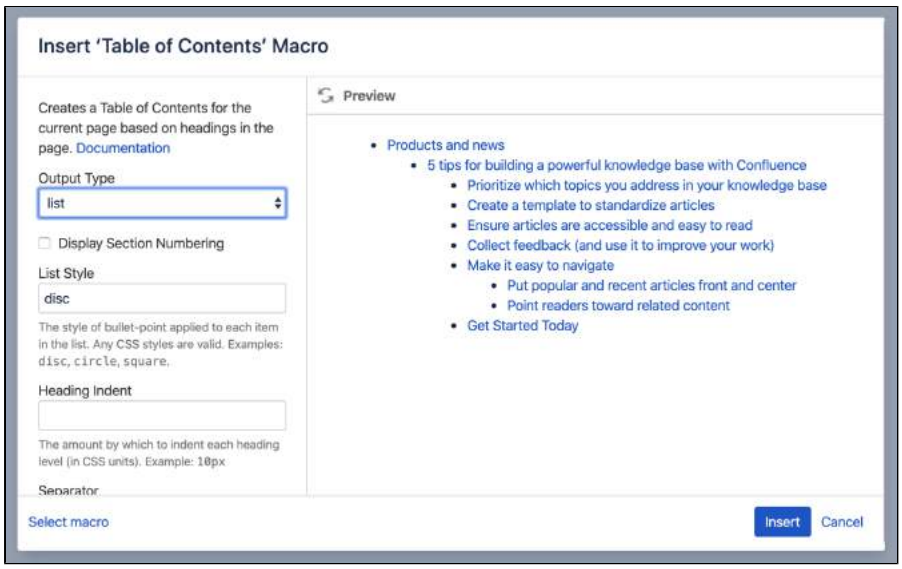

Change the macro parameters

Macro parameters are used to change the behaviour of a macro.

To change the macro parameters:

1. In the editor, click the macro placeholder and choose**Edit**.

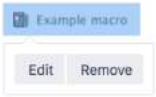

2. Update the parameters as required then choose**Insert**.

Here's a list of the parameters available in this macro.

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`)	default	Select the style of bullet point for each list item. Enter any valid CSS style, such as: • `default` -matches Confluence's default bullet style which uses a different style for each level • `none` no bullet • `disc` a filled circle • `circle` an open circle • `square` a square • `decimal` a numbered list (1, 2, 3, 4, 5) • `lower-alpha` a lower-case, alphabetical list (a, b, c, d, e) • `lower-roman` a lower roman numeral list (i, ii, iii, iv, v, vi) • `upper-roman` an upper roman numeral list (I, II, III, IV, V, VI)

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

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 headingsto 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.

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`).

Examples

The examples below are based on this table of contents:

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

- 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 'Favorite', 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]

Other ways to add this macro

Add this macro as you type

Type**{**followed by the start ofthe macro name, to see a list of macros.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

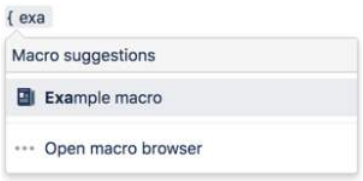

Add this macro using wiki markup

This is useful when you want to add a macro outside the editor, for example as custom content in the sidebar, header or footer of a space.

Macro name: `toc`

Macro body: None.

This example shows a list-type table of contents.

```
{toc:printable=true|style=square|maxLevel=2|indent=5px|minLevel=2|class=bigpink|exclude=[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=.*}
```

Notes and known issues

- 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.)
- The Table of Contents macro only displays page or blog post content. You can't use it to add a table of contents of headings in a comment for example.
- 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.
  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.
- There is a known issue where if you click a heading in the Table of Contents macro, then click the back button in your browser, you won't be returned to the table of contents (or to your previous page). As a workaround, use theTable of Content Zone Macro. See
  CONFSERVER-40462 GATHERING IMPACT and CONFSERVER-52497 GATHERING IMPACT
  for more information.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Do more with Confluence

Extend Confluence with one of the hundreds of other macros in the Atlassian Marketplace.Here's some specific to documentation:

- Scroll Office for Confluence- turn your Confluence pages pages into professionally styled documents
- Advanced Children Display for Confluence: combine Confluence's built-in children display and table of contents macros

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Table of Content Zone Macro

Add the Table of Content Zone macro to a page to display a table of contents froma defined section of the page.

This macro is great for:

- creating a table of contents from sections of a page
- having multiple table of contents throughout a very long page.

You'll have to create a table of contentsfor headings within the body of the macro.

On this page:

- Add this macro to your page
- Change the macro parameters
- Examples
- Other ways to add this macro
- Troubleshooting

Screenshot: The Table of Content Zone macro configured with flat Output Type.

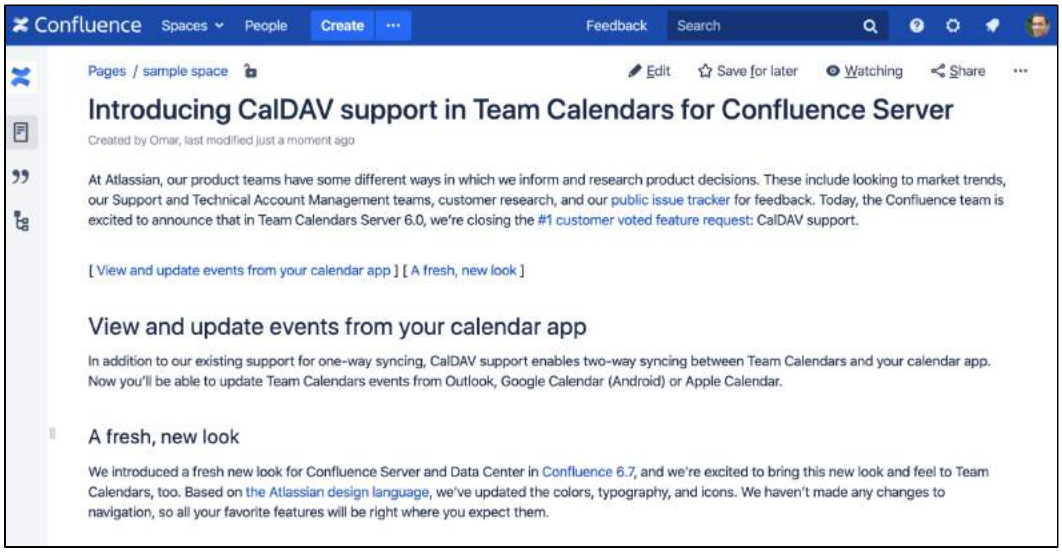

Add this macro to your page

To add the Table of Content Zone macro to a page:

1. From the editor toolbar, choose **Insert** + > **Other Macros**.
2. Choose**Table of Content Zone**from the**Confluence content**category.
3. Enter any parameters.
4. Choose**Insert**.
5. Add or paste your content into the macro body. The headings within the macro will be included in the table of contents.

You can then publish your page to see the macro in action.

Screenshot:Entering headings within the body of the Table of Content Zone macro.

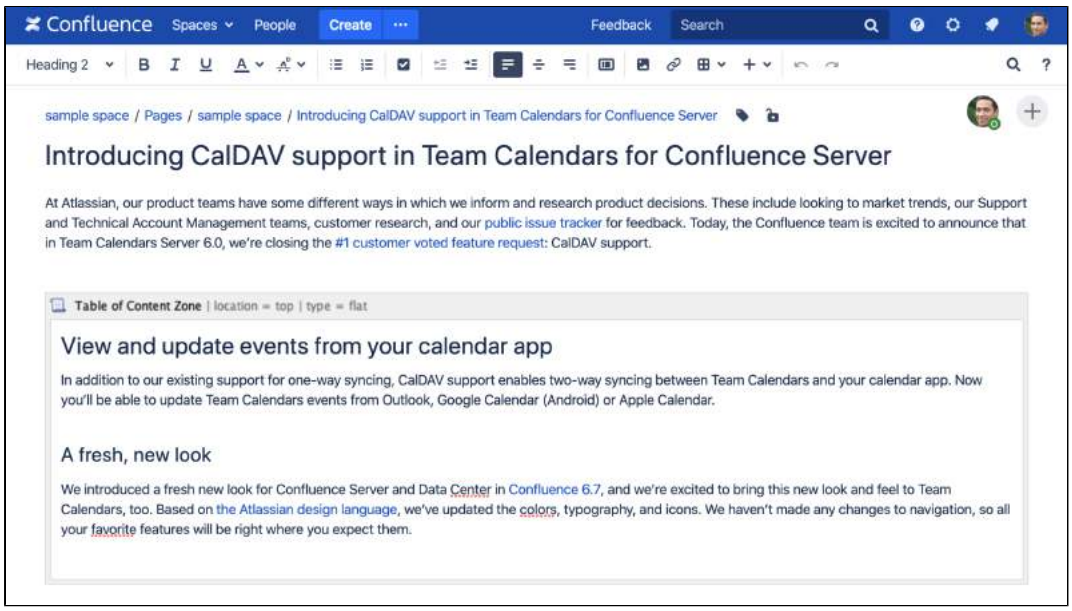

Change the macro parameters

Macro parameters are used to change the behaviour of a macro.

To change the macro parameters:

1. In the editor, click the macro placeholder and choose**Edit**.

2. Update the parameters as required then choose**Insert**.

Here's a list of the parameters available in this macro.

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. • **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.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

List Style (style)	default	Select the style of bullet point for each list item. Enter any valid CSS style, such as: • `default` -matches Confluence's default bullet style which uses a different style for each level • `none` no bullet • `disc` a filled circle • `circle` an open circle • `square` a square • `decimal` a numbered list (1, 2, 3, 4, 5) • `lower-alpha` a lower-case, alphabetical list (a, b, c, d, e) • `lower-roman` a lower roman numeral list (i, ii, iii, iv, v, vi) • `upper-roman` an upper roman numeral list (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 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 table of contents. 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 table of contents is set to print. If you clear this parameter, the table of contents will not be visible when you print the page.
CSS Class Name (class)		If you have a custom table of contents in your CSS style sheet, you can use this parameter to output the table of conents with the specified "class" attribute.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

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`).

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 "Favorite", 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

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

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

[Unknown Places] [Unknown Foods] [Unknown Foods]

Other ways to add this macro

Add this macro as you type

Type**{**followed by the start ofthe macro name, to see a list of macros.

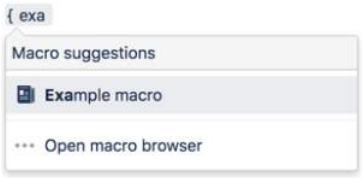

Add this macro using wiki markup

This is useful when you want to add a macro outside the editor, for example as custom content in the sidebar, header or footer of a space.

Macro name:`toc-zone`

Macro body: Accepts rich text.

```
{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}
```

Troubleshooting

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.

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.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Task Report Macro

Add the Task Report macro to a page to display a list of tasks. Filter the tasks by space, page, user, label, created date and more.

This macro is great for:

- meeting notes
- status reports
- project planning pages.

See Add, Assign, and View Tasksfor 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:

- Add this macro to your page
- Change the macro parameters
- Other ways to add this macro

Related pages:

- Add, Assign, and View Tasks

Screenshot: The Task Report macro showing incomplete tasks on a meeting notes page.

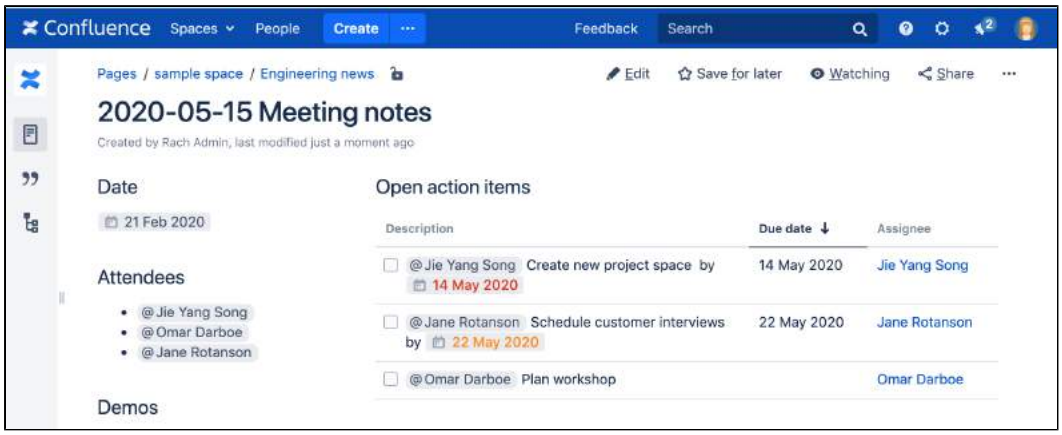

Add this macro to your page

To add the Task Report macro to a page:

1. From the editor toolbar, choose **Insert** > **Other Macros**.
2. Choose**Task Report**from the**Confluence content**category.
3. Enter any parameters.
4. Choose**Insert**.

You can then publish your page to see the macro in action.

Screenshot: Configuring the Task Report macro to show tasks from a specific space.

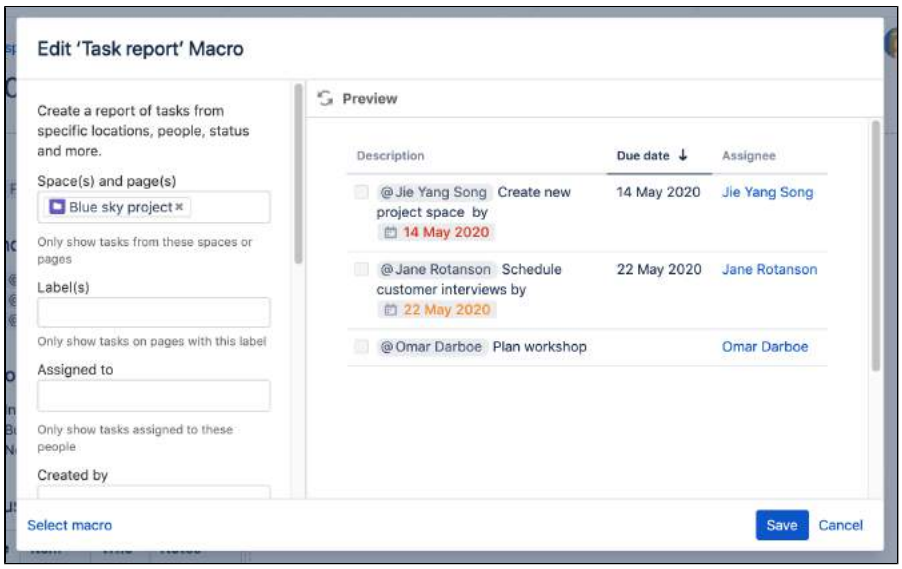

Change the macro parameters

Macro parameters are used to change the behaviour of a macro.

To change the macro parameters:

1. In the editor, click the macro placeholder and choose**Edit**.

2. Update the parameters as required then choose**Insert**.

Here's a list of the parameters available in this macro.

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 spaces or parent pages (and their child pages) 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 at least one of the specified labels (for example, 'label-a' OR 'label-b'). 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.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

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.

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`).

Other ways to add this macro

Add this macro as you type

Type **{** followed by the start ofthe macro name, to see a list of macros.

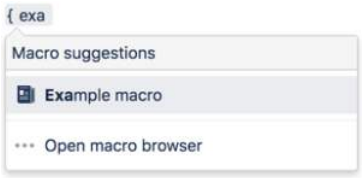

Add this macro using wiki markup

You can't use wiki markup to add this macro.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Team Calendar Macro

Add the Team Calendar macro to a page to display a calendar on a Confluence page, making it easy to track and manage events.

This macro is great for:

- displaying team events, project timelines, and milestones in project and team spaces
- sharing Jira issue due dates and sprint dates on sprint planning or retrospective pages
- publishing schedules and on-call rotations where everyone can see them.

On this page:

- Add the Team Calendar macro to your page
- Change the macro parameters
- Other ways to add this macro

Screenshot: The Team Calendar macro showing important project dates.

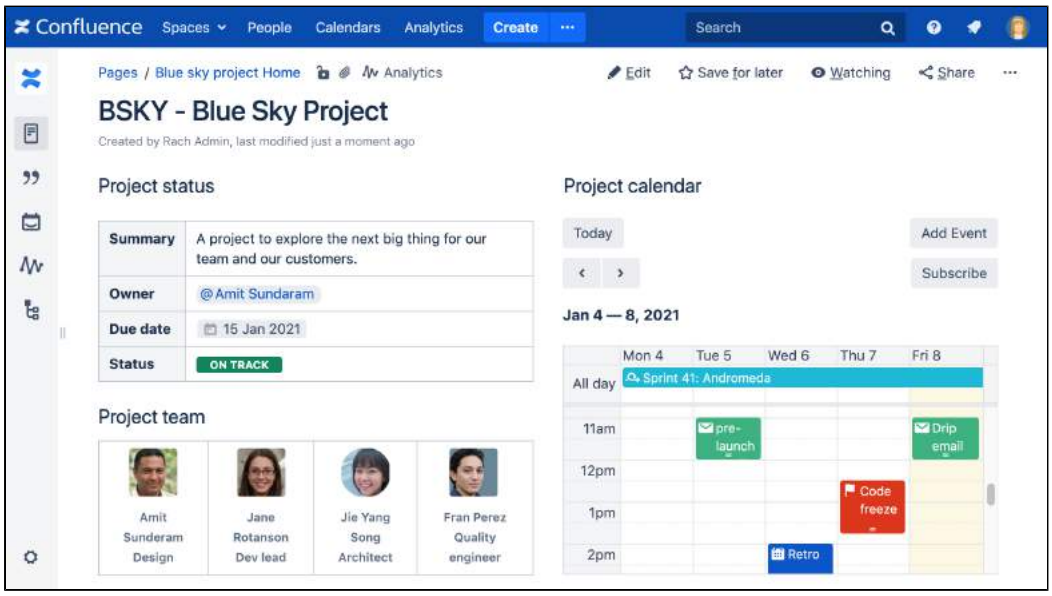

For general information about creating calendars, or subscribing to existing calendars, see Team Calendars.

Add the Team Calendar macro to your page

To add the Team Calendar macro to a page:

1. From the editor toolbar, choose **Insert** + > **Other Macros**.
2. Select **Team Calendar**from the**Reporting**category.
3. Select **Add Existing Calendar**.
4. Search for the calendar name
5. Select **Add**.

You can then publish your page to see the macro in action.

Screenshot: configuring the Team Calendar macro in the macro browser.

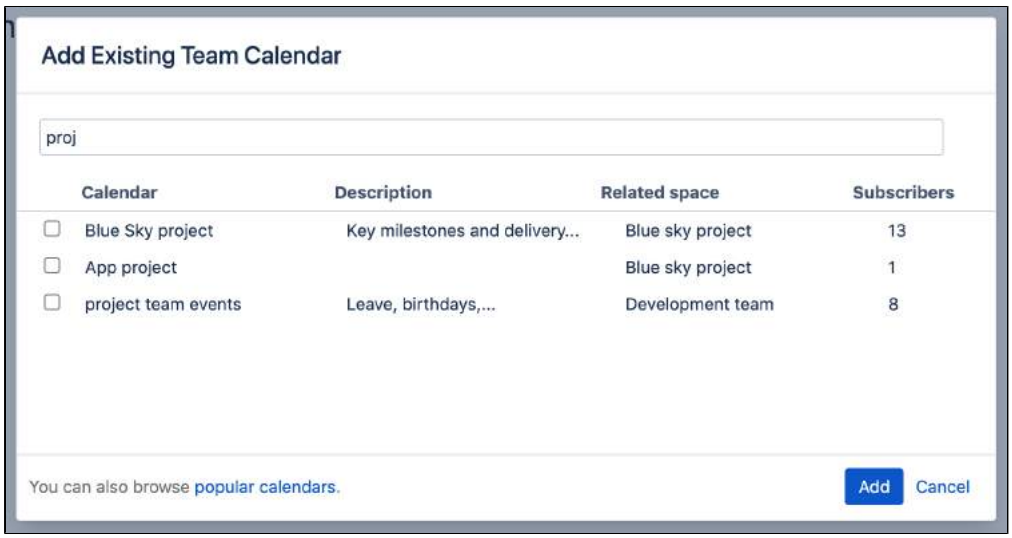

Change the macro parameters

Macro parameters are used to change the behaviour of a macro.

To change the macro parameters:

1. In the editor, click the macro placeholder and choose**Edit**.
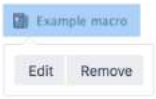

2. Update the parameters as required then choose**Insert**.

Here's a list of the parameters available in this macro.

Parameter	Required	Default	Description
Calendar	Yes	blank	The name of the calendar to display.
View (`defaultView`)	No	month	Format to display the calendar. Available values: • month • week • list • timeline
Width (`width`)	No	none	Width of the calendar in pixels. Leave blank, and the calendar will fill the available space.
Height	No	none	Maximum height of the calendar in pixels. Only applies to the timeline view.
Calendar legend	No	none	Location of the calendar legend in the month and week views. The legend shows the event types present in the selected calendar. Available values: • Right • Bottom • None

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Calendar ID (`id`)	Yes (Wiki markup only)	blank	One or more calendar IDs, separated by commas. This parameter is only required when adding the macro using wiki markup. To find out the ID for a calendar: 1. Choose**Calendars** from the Confluence header or space sidebar 2. Click ••• next to a calendar and choose **Embed** A dialog appears with the link to the calendar. The last part of the link contains the ID of the calendar. For example: `http://<your_site_url>/calendar/previewcalendar.action?subCalendarId=890143d3-5c7d-4f17-ad4b-70d8a4a4d345`

Other ways to add this macro

Add this macro as you type

Type**{**followed by the start ofthe macro name, to see a list of macros.

Add this macro using wiki markup

This is useful when you want to add a macro outside the editor, for example as custom content in the sidebar, header or footer of a space.

Macro name: `calendar`

Macro body: None.

```
{calendar:id=4f5f9524-f588-468e-a48c-668ea480a77c,49221951-cca9-424f-b1ee-42ec44452bcc|defaultView=month|width=400}
```


Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

User List Macro

Add the User List macro to a page to display a listof Confluence users, in a particular group.The macro can also indicate when users are online or offline.

This is a **legacy** macro, and can cause performance issues in very large sites. We recommend you don't use this macro to attempt to show all users in a site.

On this page:

- Limitations
- Add this macro to your page
- Change the macro parameters
- Other ways to add this macro

Screenshot: A page containing two User List macros.

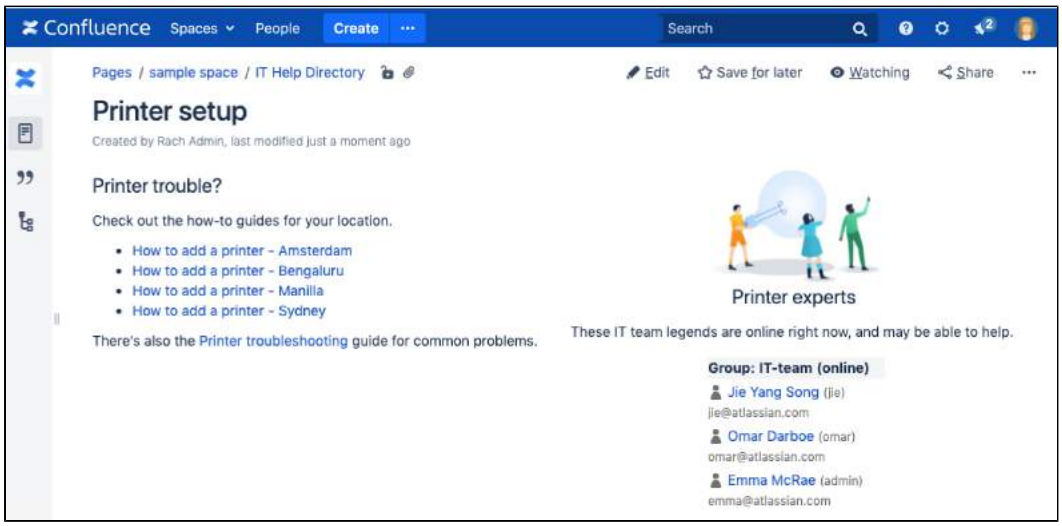

Limitations

The User List macro can be quite memory hungry in sites with lots of users. To prevent it causing out of memory errors in your site, we don't show this macro if there are more than 10,000 people in the groups specified. Your administrator can change this limit using the `confluence.extra.userlister.limit` system property.

Add this macro to your page

To add the User List macro to a page:

1. From the editor toolbar, choose **Insert** + > **Other Macros**.
2. Choose**User List**from the**Confluence content**category.
3. Enter the group you want to display.
4. Choose**Insert**.

You can then publish your page to see the macro in action.

Screenshot: Configuring the User List macro to show the members of a group.

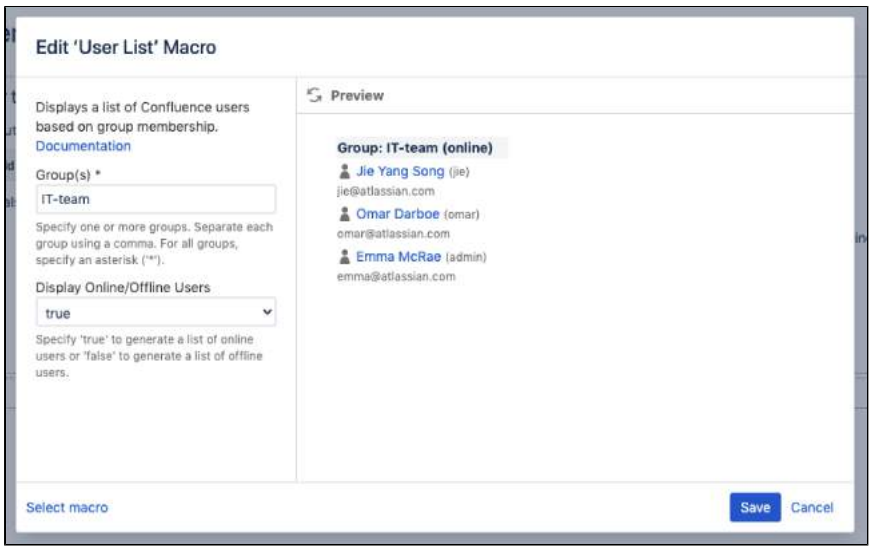

Change the macro parameters

Macro parameters are used to change the behaviour of a macro.

To change the macro parameters:

1. In the editor, click the macro placeholder and choose**Edit**.

2. Update the parameters as required then choose**Insert**.

Here's a list of the parameters available in this macro.

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. Seethis 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.

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`).

Other ways to add this macro

Add this macro as you type

Type**{**followed by the start ofthe macro name, to see a list of macros.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

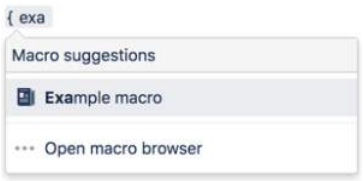

Add this macro using wiki markup

This is useful when you want to add a macro outside the editor, for example as custom content in the sidebar, header or footer of a space.

Macro name:`userlister`

Macro body: None.

```
{userlister:groups=confluence-users|online=false}
```


Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

User Profile Macro

Add the User Profile macro to a page to show profile information about a user.

This can be useful for:

- team space homepages
- project pages.

The macro will display any details the user has added to their profile, such as their phone number, role, department, and location.

On this page:

- Add this macro to your page
- Change the macro parameters
- Other ways to add this macro

Screenshot: Four User Profile macros on a project page.

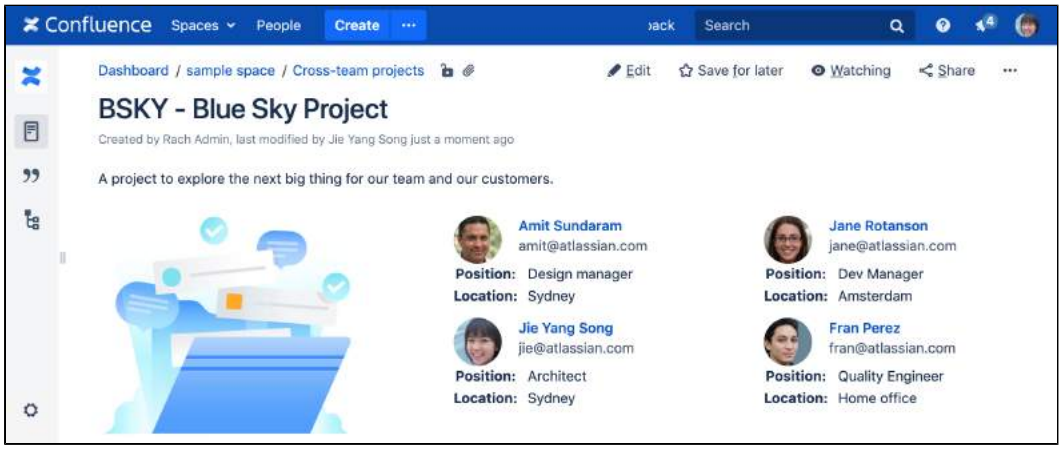

This macro is more useful if your profile is complete. Head to Your User Profile to find out how you can edit your profile details.

Add this macro to your page

To add the User Profile macro to a page:

1. From the editor toolbar, choose **Insert** + > **Other Macros**.
2. Choose**User Profile**from the**Confluence content**category.
3. Enter the username of the person you want to display.
4. Choose**Insert**.

You can then publish your page to see the macro in action.

Screenshot: Specifying a user in the User Profile macro.

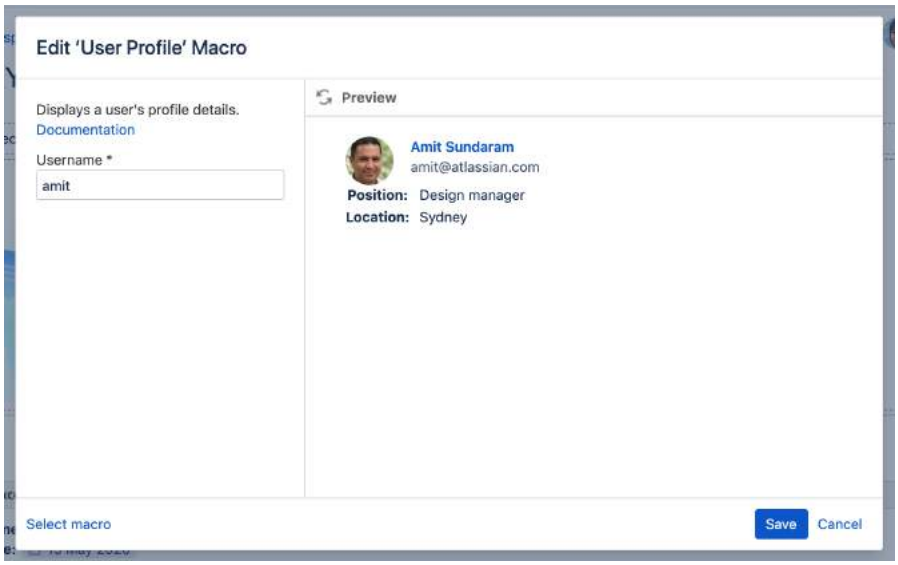

Change the macro parameters

Macro parameters are used to change the behaviour of a macro.

To change the macro parameters:

1. In the editor, click the macro placeholder and choose**Edit**.
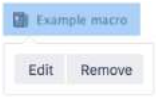

2. Update the parameters as required then choose**Insert**.

Here's a list of the parameters available in this macro.

Parameter	Default	Description
Username (user)	none	The username of the Confluence user whose profile summary you wish to show.

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`).

Other ways to add this macro

Add this macro as you type

Type**{**followed by the start ofthe macro name, to see a list of macros.

Add this macro using wiki markup

This is useful when you want to add a macro outside the editor, for example as custom content in the sidebar, header or footer of a space.

Macro name:`profile`

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Macro body: None.

```
{profile:user=admin}
```


Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

View File Macro

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 specific information about each macro, see:

- Office Excel Macro
- Office PowerPoint Macro
- Office Word Macro
- PDF Macro

On this page:

- Add this macro to your page
- Supported file types
- Editing files
- Importing content from Word documents

Add this macro to your page

To add a View File macro to a page:

1. From the editor toolbar, choose **Insert** > **Other Macros**.
2. Choose**Office Excel**, **Office PowerPoint**, **Office Word**, **or PDF**from the**Confluence content**category.
3. Select the file you want to display and enter any additional parameters.
4. Choose**Insert**.

You can then publish your page to see the macro in action.

Supported file types

To display an Office or PDF document in a page, use one of the followingView Filemacros in the macro browser:

- Office Excel(`.xls`and `.xlsx`)
- Office PowerPoint(`.ppt`and `.pptx`)
- Office Word(`.doc`and `.docx`)
- PDF (`.pdf`)

Editing files

You can edit files embedded with the View File macro using your preferred desktop application, then save the file back to Confluence automatically. Head toEdit files for instructions.

- **Word and Excel** choose the **Edit** button above the content.
- **PowerPoint and PDF** choose the edit icon on the viewer.

Importing content from Word documents

If you want to use the contents of your Word document to create a Confluence page, seeImport a Word Document into Confluence.

Widget Connector Macro

Add the Widget connector to a page to embed online videos, slideshows, photostreams and more.

This is great for bridging the gap between Confluence and other sites and services your team uses to get work done.

The macro currently supports content from these sites:

- YouTube
- Vimeo
- Twitter
- Google Docs, Sheets, and Slides
- Google Calendar
- Google Maps
- Wufoo
- Facebook
- LinkedIn
- Microsoft Stream
- Figma
- Spotify
- Prezi

It can also display content from these sites, once they have been added to theallowlist:

- Scribd
- Flickr (requires Flash)
- Slideshare (requires Flash)
- Viddler (requires Flash)

On this page:

- Add this macro to your page
- Change the macro parameters
- Examples
  - YouTube
  - Vimeo
  - Flickr
  - Twitter
  - Google Docs, Slides, and Sheets
  - Google Calendar
  - Google Maps
  - Facebook
  - LinkedIn
  - Microsoft Stream
  - Figma
  - Spotify
  - Prezi
- Other ways to add this macro

Add this macro to your page

To add the Widget Connector macro to a page:

1. From the editor toolbar, choose **Insert** > **Other Macros**.
2. Choose**Widget Connector**from the**Media**category.
3. Enter the URL you want to display.
4. Choose**Insert**.

You can then publish your page to see the macro in action.

Change the macro parameters

Macro parameters are used to change the behaviour of a macro.

To change the macro parameters:

1. In the editor, click the macro placeholder and choose**Edit**.
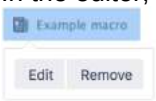

2. Update the parameters as required then choose**Insert**.

Here's a list of the parameters available in this macro.

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.

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`).

Examples

Every site is a little different, so we've put together some info on what you'll need to do to embed content from each site 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.youtube.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.youtube.com/watch?t=15&v=LhHKkodOPFo`. Paste in the short sharing URL to be sure it works.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

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.

Flickr

You can embed albums (formerly known as sets) and tags. You can't embed individual photos or user photostreams.

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.

The Widget Connector uses Flash to display this content. For security reasons, Flash is disabled in most modern browsers.

Twitter

You can embed single tweets, profiles, lists, and moments.

Single Tweets

To embed a single tweet, click**Copy link to tweet**.Add the Widget Connector macro to the page and paste the link into the URL field.

When theembedded Tweet is a reply, the parent Tweet will be displayed above the reply.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

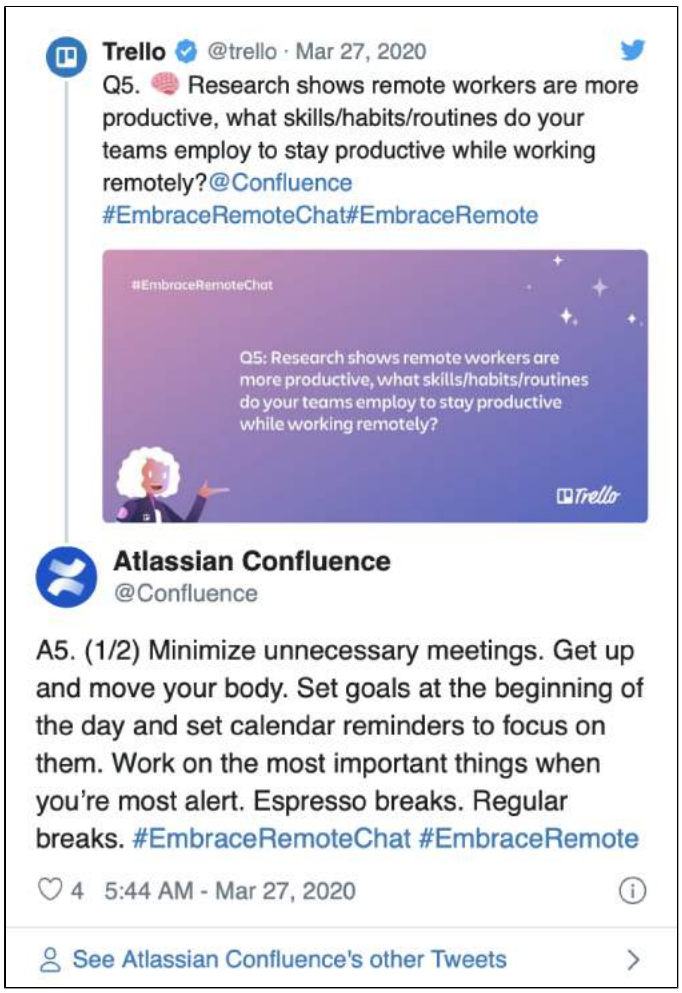

Profiles and Lists

Click on the Profile or List and copy the URL from the address bar of your browser. Paste the URL into the editor and Confluence will autoconvert the link and insert the macro for you.

Alternatively, you can**Copy link to profile**or**Copy link to List**. These links wont autoconvert, youll need to add the Widget Connector macro to the page and paste the link into the URL field.

Moments

Click on the Moment, copy the URL from the address bar of your browser, orselect **Copy link to this**.Add the Widget Connector macro to the page and paste the link into the URL field.If you see an **error rendering tweet**message, replace the word 'events' with 'moments' in the URL field.

Google Docs, Slides, and Sheets

In Google Docs, Sheets, or Slides,click **File** and select **Publish to the web**. Click **Publish** then copy the link. Paste it into the editor and Confluence will autoconvert the link and insert the macro for you.

Google Calendar

Open Google Calendar andselect**Settings**. Click the name of the calendar you want to embed.

You can only embed public calendars. To allow all visitors to see your calendar, open**Access permissions**a nd check the box next to**Make available topublic**.

In the**Integratecalendar**section, copy the**Public URL to this calendar**.Paste the URL into the editor and Confluence will autoconvert the link and insert the macro for you.

Google Maps

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

You can embed interactive mapsfor a business, address, place, Street View 360-degree panorama, and some (not all) search results (such as Bars, Banks, and Restaurants).

Business,address, and place

Copy the URL of a business, address, orplacefrom the address bar of your browser and paste it into the editor.Confluence will autoconvert the link and insert the macro for you.

Street View andsearch results

In Street View mode, or when viewing search results, click the menu and go to**Share or embed map**>**Embed a map**>**COPY HTML**. Paste the link into the editor and Confluence will autoconvert the link and insert the macro for you.You'll need to delete the additional iframe tags in theeditoror they will display alongside the map.

Facebook

On a public post, go to **Embed** > **Advanced Settings** and copy the **URL of this post**. Paste the URL into the editor and Confluence will autoconvert the link and insert the macro for you.If you can't see the post headline, description, or like count, change the Pixel Height and Pixel Width Values in the macro parameters.

LinkedIn

Only LinkedIn posts that an author has shared as Public can be embedded. Posts can include articles, images, and videos.Select**Embed this post**.

Copy the code without including the HTML iframe tags, like in the image below.Paste the code into the editor and Confluence will autoconvert the link and insert the macro for you.

```
<iframe src="https://www.linkedin.com/embed/feed/update/urn:li:share:6662864786231042048" height="633" width="504" frameborder="0" allowfullscreen="" title="Embedded post"></iframe>
```

Alternatively, click **Copy code**.Paste the code into the editor and Confluence will autoconvert the link and insert the macro.You'll need to delete the additional iframe tags in theeditoror they will display alongside the LinkedIn post.

Microsoft Stream

To share a video, select the**Share**button and copy the**Direct link to video**.Paste the link into the editor and Confluence will autoconvert the link and insert the macro for you. You can also paste the URL from the address bar on the video page.

Only people authorized to see a video will be able to view it.If you receive a playback error message,select the**Open in new window**buttonin the message to play the video in a new window.

Figma

Open a Figma file and click**Share**.In the Link Sharing settings, select**Anyone with the link can view**. Click**Copy link**.Paste the link into the editor and Confluence will autoconvert the link and insert the macro for you.

If youre using Figma in thebrowser, you can copy the URL from the address barand paste it into the editor. Confluence will autoconvert the link.

Hover your cursorover the embedded file to see options for full-screen mode and zooming in and out. You can click, hold, and drag the embedded file to see more of it.Clicking the file name link in the bottom-left of the embedded file will open the file directly in Figma.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

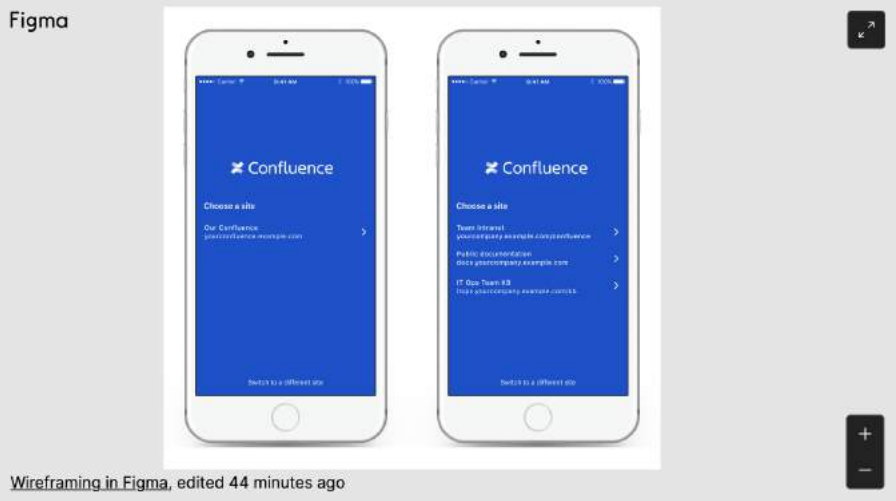

Spotify

You can embed a song, album, or artist.

Click**Copy Link**when using the web player. If youre using the desktop app, select**Share**>**CopyLink**. Pastethe link into the editor and Confluence will autoconvert the link and insert the macro for you.

Alternatively, in the web player, copy the URL from the address barand paste it into the editor.Confluence will autoconvert the link.

Prezi

Open a Presentation, Design, or Video, and**Copy**the link. Paste the link into the editor and Confluence will autoconvert the link and insert the macro for you.

Alternatively, from your Prezi dashboard, select**Share view link**on aPresentation or**Share**on a Design.**Copy**t he link, paste it into the editor and Confluence will autoconvert the link.

Troubleshooting

If the Widget Connector can't display content from the external site, the macro will look like this:

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 canraise an issueabout it to let us know.

Some sites must be added to the allowlist

The following sites need to be added to Confluence's allowlist before the macro can display any content. This is due to the way we need to connect to that site.

- Scribd
- Flickr
- Slideshare
- Viddler

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Some content requires Flash

The Widget Connector requires Flash for Flickr, Slideshare, and Viddler. This is blocked by most modern browsers due to security concerns. We don't recommend you enable the Flash plugin in your browser.

Other ways to add this macro

Add this macro as you type

Type**{**followed by the start ofthe macro name, to see a list of macros.

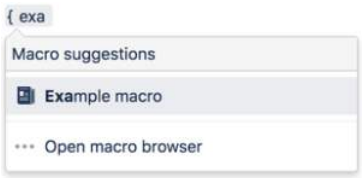

Add this macro using wiki markup

This is useful when you want to add a macro outside the editor, for example as custom content in the sidebar, header or footer of a space.

Macro name:`widget`

Macro body: None.

```
{widget:height=400|width=400|url=http://youtube.com/watch?v=23pLByj_q5U}
```


Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

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, favorite 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
- Save for later
- View and Revoke OAuth Access Tokens

Related pages:

- Watch Pages, Spaces and Blogs
- Create a Space
- Save for later

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 thePeople Directory, if you haven't set up yourpersonal space.

In your own profile, you can access account management features and update information about yourself, like your name, email address, and pass word. You can also view other users' profiles.

On this page:

- Find your user profile
- Edit your user profile
- Notes

Related pages:

- Set Your Profile Picture
- Create a Personal Space

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

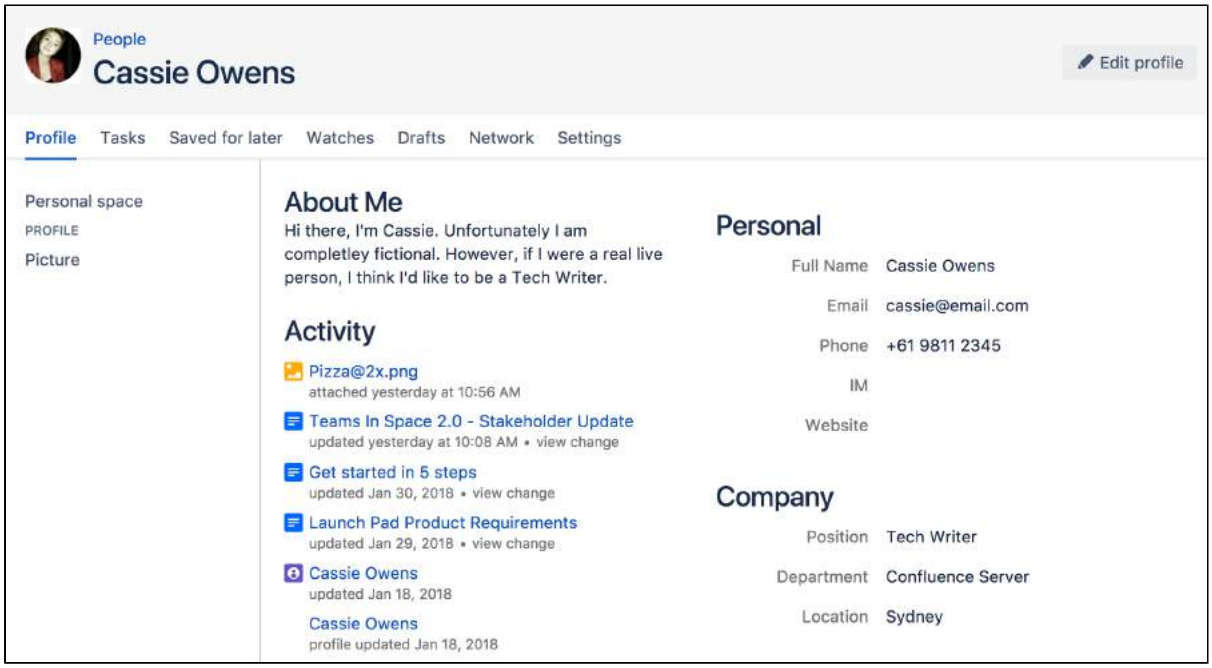

From your user profile, you can access the following:

Profile	• View andedit 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 aprofile picture(optional). • Change yourpassword.

Network	• View the recent activity of users that you are following via theNetwork view. • Follow other users from this view.
Saved for later	• View a list of pages yousaved for later.
Watches	• View a list of the pages and spaces you are currentlywatching.
Drafts	• Retrieve any pages you were in the process of editing. SeeDrafts.
Settings	• Edit yourGeneral Settings(homepage, language and timezone). • Subscribe toemail notifications. • View and revoke yourOAuth 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 organization.
Department	The name of your department or team.
Location	Your location. This can be your town, city, region or country.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

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.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

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

On this page:

- From within Confluence
- From the login page
- Don't see the Password option?

Related Pages:

- Your User Profile
- Set Your Profile Picture
- Create a Personal Space

From within Confluence

Change your password when you're logged in:

1. Choose your profile picture at top right of the screen, then choose**Profile**
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**

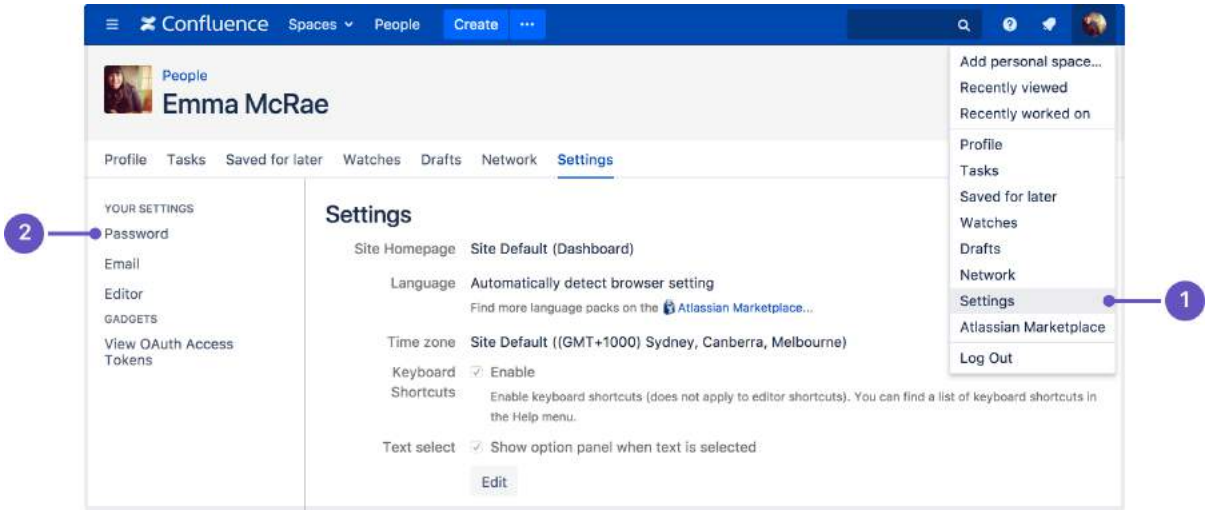

1. Location of **Settings**in the profile menu
2. Location of **Password** in the side bar

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.

Don't see the Password option?

You may not be able to change your password directly in Confluence if your login credentials are coming from another user directory, for example if Confluence is integrated with an LDAP directory or Jira for user management.

Talk to your administrator about where you should change 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.

On this page:

- General User Preferences
- More about Language
- Editor Preferences

Related pages:

- Your User Profile
- Set Your Profile Picture
- Create a Personal Space
- Autocomplete for links, files, macros, mentions and emojis

General User Preferences

To edit your general user settings:

1. Choose **your profile picture** at top right of the screen, then choose**Settings**
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

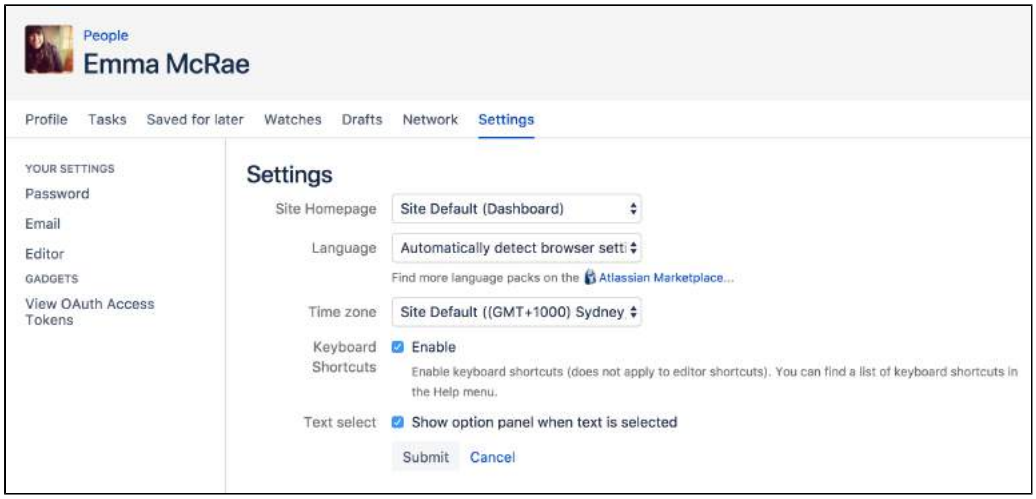

More about Language

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 prioritized list of languages. Confluence will use the first supported language in that list.Confluence administrators can disable this option by setting the `confluence.browser.language.enabled` system property to `false`.
- The default language for your site, as defined by your Confluence site administrator.

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

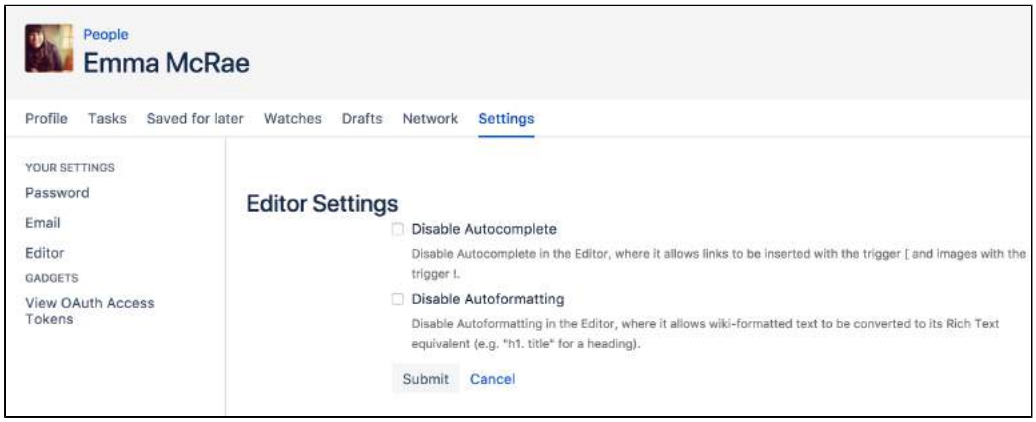

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Set Your Profile Picture

Your profile picture is used as the icon for yourpersonal space, to represent you in thePeople Directory, and to illustrate yourcomments. 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.

Related pages:

- Your User Profile
- Create a Personal Space
- Your profile and settings

This page is about Confluence Server and Data Center. If you use Confluence Cloud head over here to see how to update your personal profile.

Upload and adjust your profile picture:

1. Choose your profile picture at top right of the screen, then choose**Profile**
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**

Screenshot: Choosing a profile picture

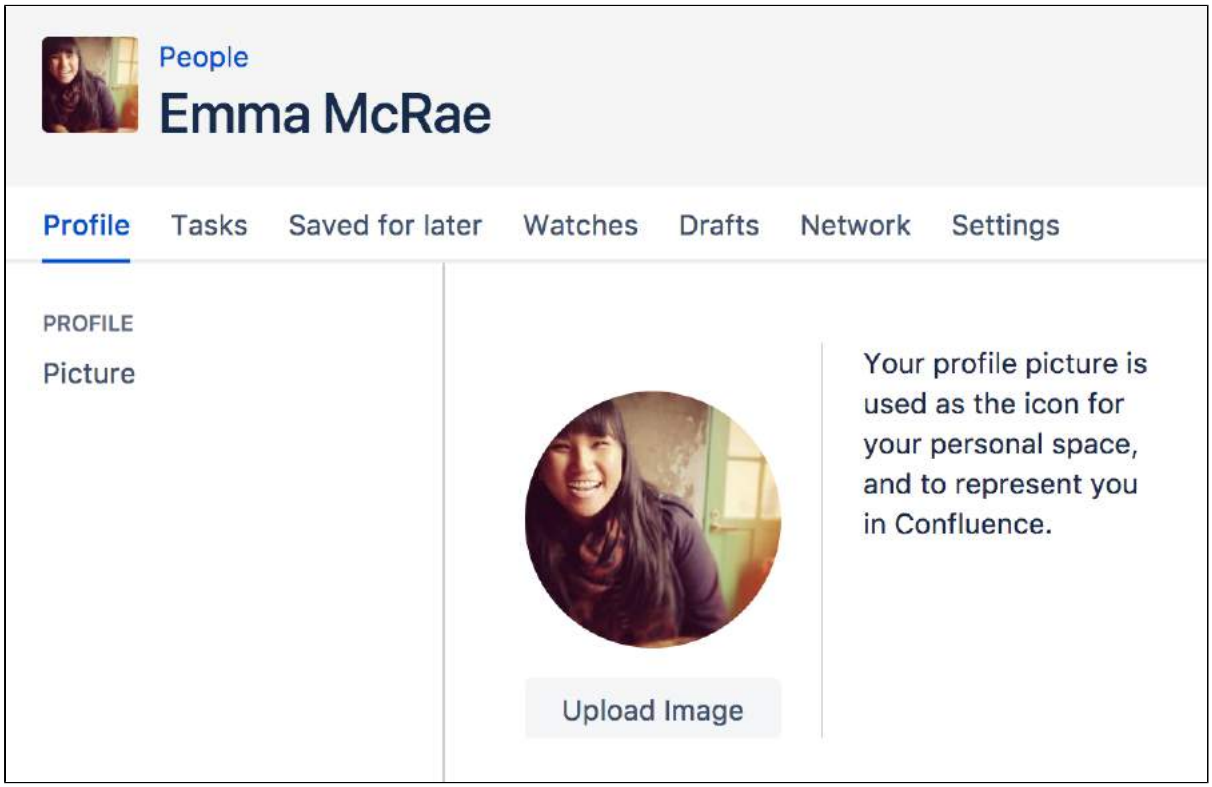

Screenshot: Resize and position your profile picture

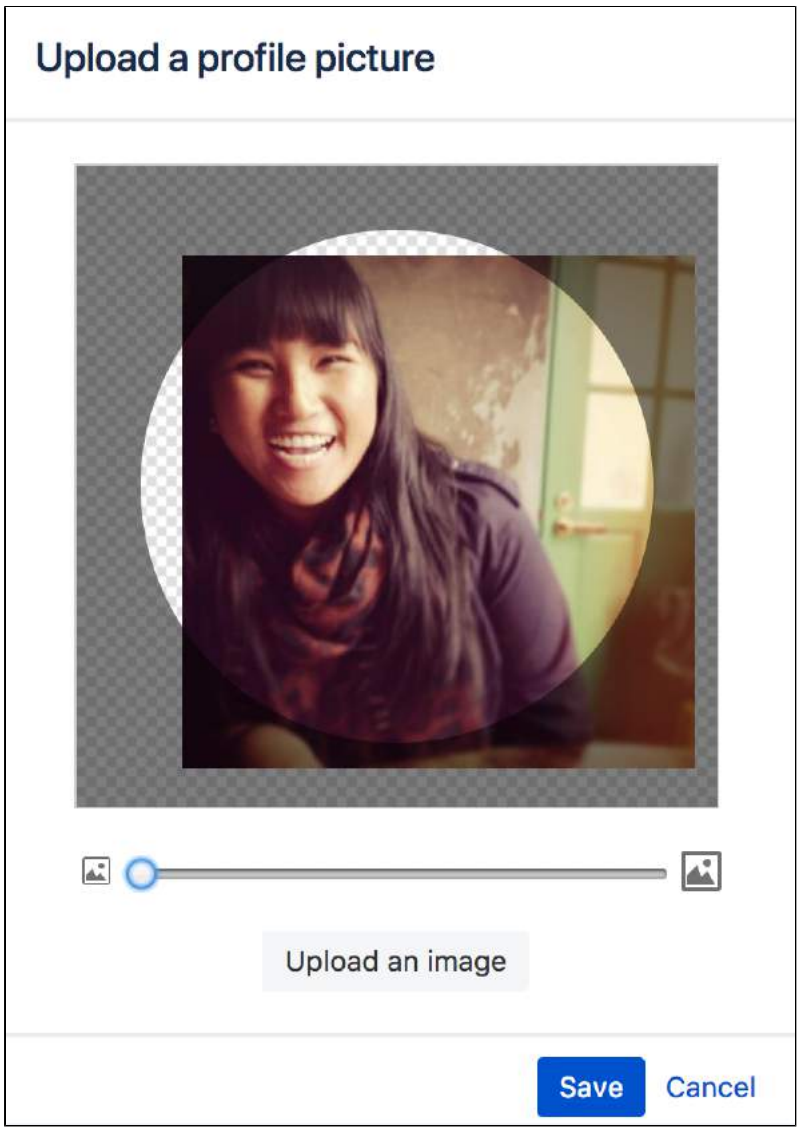

You can't remove your own profile picture, but you can upload a new one any time. Alternatively you can ask your admin to remove your profile picture for you.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

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 personalize your experience. and use an existing space home page as your landing page.

Related pages:

- Your profile and settings
- Your User Profile

To set your home page:

1. Choose **your profile picture** at top right of the screen, then choose**S ettings**
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

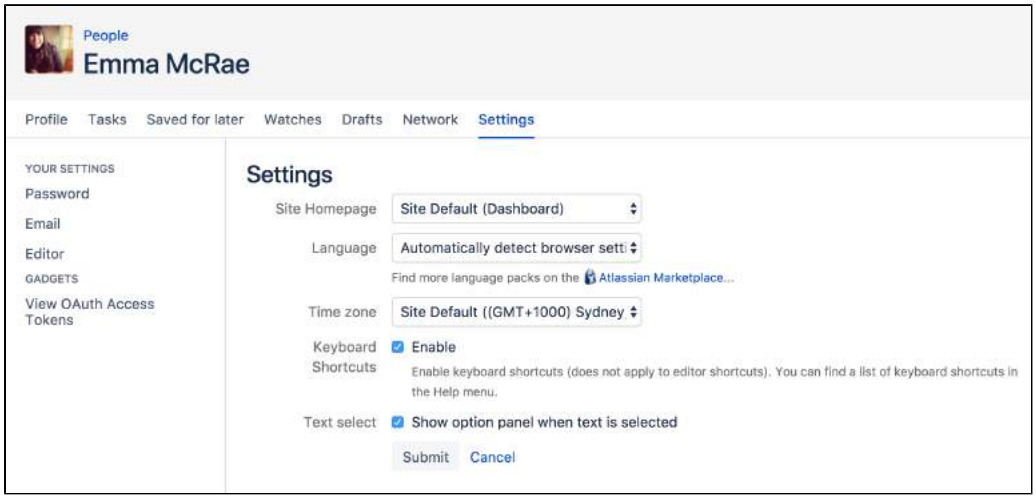

You can access the dashboard at any time using the dashboard URL. It'll look something like this:`https://yoursite.com/wiki/dashboard.action`.

Save for later

Saving pages for later helps you access them quickly from the dashboard or from your profile.

No time to read that page now? No problem, hit **Save for later** and it'll be waiting for you on the dashboard when you have more time. It's also a great place to store those pages that you use on a day to day basis.

Save for later was previously called Favorites.

On this page:

- Save a page for later
- Get back to your saved pages

Save a page for later

To save a page for later, hit the **Save for later** button at the top of the page.

The star icon will change to dark grey to indicate the page is saved. Hit the button again if you want to remove the page from the list.

Get back to your saved pages

To view your saved pages:

- Choose **Saved for later**on the dashboard sidebar.
- Choose**yourprofile picture**at top-right of the screen, then choose**Saved for later** there's a list of your saved pages, and the spaces that you've added to**My spaces**.

You can also use theFavorite Pages Macroto include a list of your saved pages on any page.

Screenshot: Viewing and removing saved pages from the dashboard

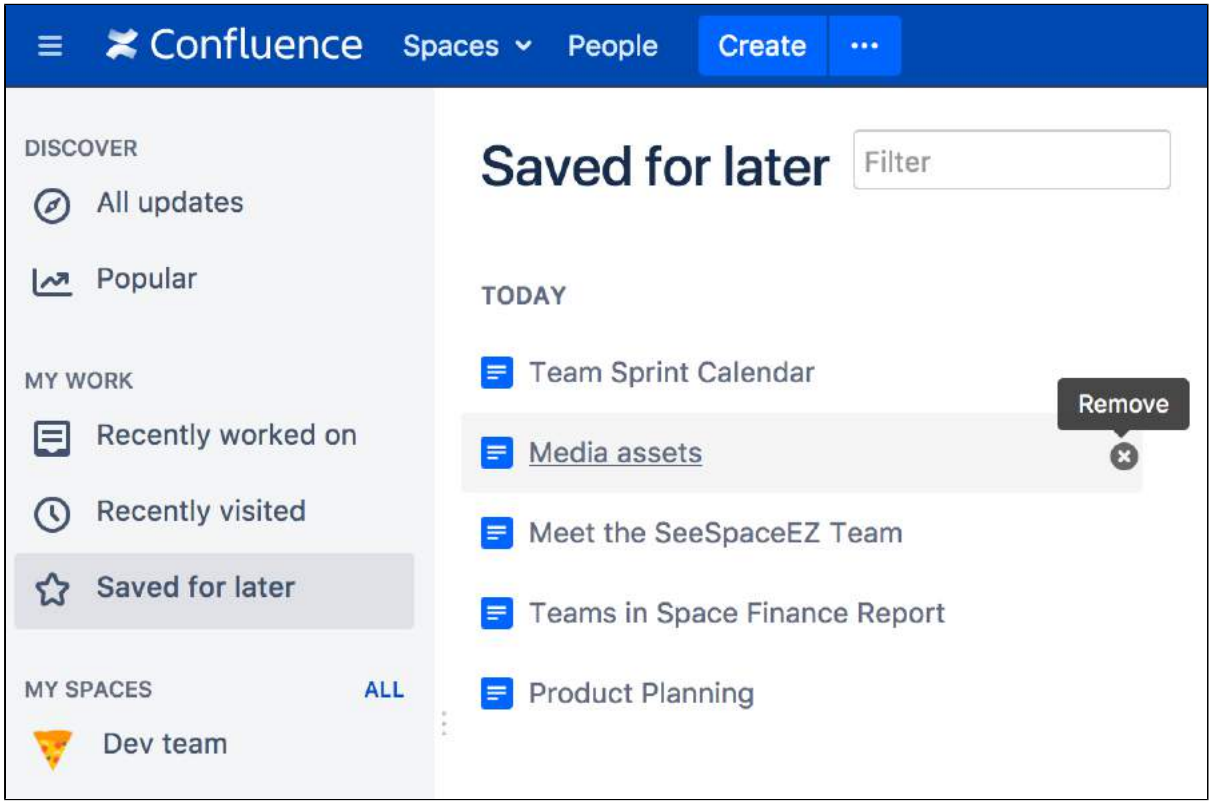

View and Revoke OAuth Access Tokens

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 toConfiguring OAuthfor 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.
    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.

On this page:

- View your OAuth Access Tokens
- Revoke your OAuth Access Tokens

Related pages:

- Configuring OAuth for Confluence Admins

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.
    If no access tokens have been set, then 'None specified' is shown.

Screenshot: Viewing your OAuth Access Tokens

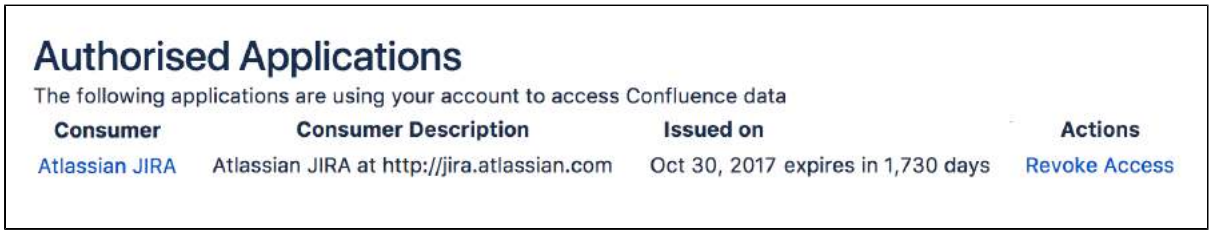

OAuth Access Token Details

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 customize 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.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

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:

- Work together with your teamon a page or blog and see their changes in real time
- Share a linkto a page or blog post via email
- Mentiona user when you write a page, blog post, comment, or add a task
- Likea 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*like* the content.

Other users can also find out about changes to content in Confluence bywat chingpages and spaces.

Another way to share Confluence content is byexporting it to other formatss uch as XML, HTML, Microsoft Word and PDF.

Related pages:

- Pages and blogs
- Watch Pages, Spaces and Blogs
- Export Content to Word, PDF, HTML and XML

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 editpages or blog posts
- Commenton a page or blog post or edit existing comments
- Update theiruser profile

On this page:

- Follow another user
- Access your network view
- Notes

Related pages:

- Network Macro
- Subscribe to a Network RSS Feed
- Email Notifications

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**yourprofile 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

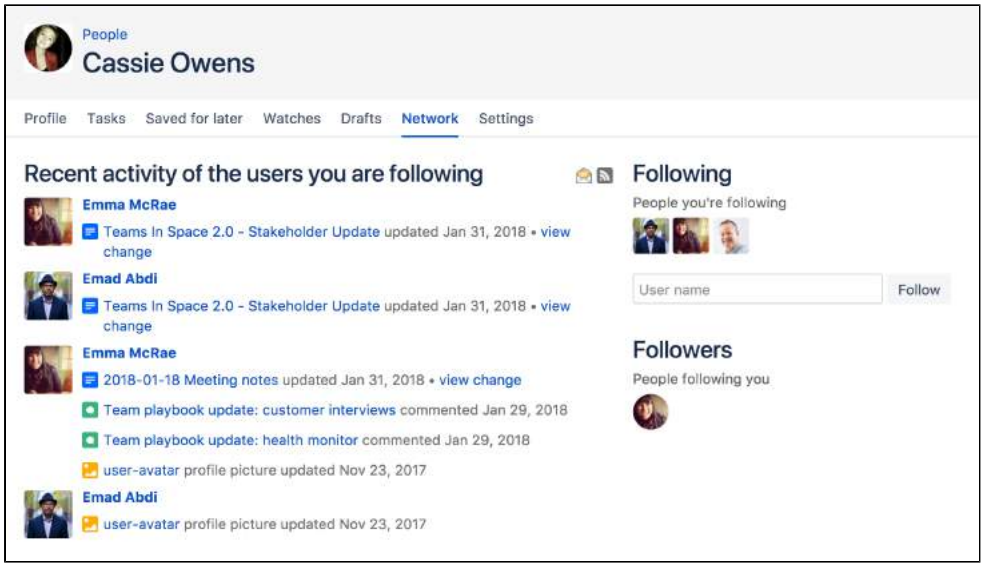

Notes

- **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.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

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.

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 thedashboard.

Disabling the 'like' feature

The*like*functionality is provided by a system app called the 'Confluence Like Plugin'. To remove the*like*functionality from your site,Disabling and enabling apps.

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`

On this page:

- Disabling the 'like' feature
- Disabling notifications when your content is 'liked'

Related pages:

- The dashboard
- Email Notifications
- Network Overview

Mentions

Mentions (often known as @mentions)are a useful 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.

Related pages:

- The Editor
- Autocomplete for links, files, macros, mentions and emojis
- Keyboard shortcuts
- Workbox Notifications

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.

Confluence will suggest people you've mentioned previously (after yourself, of course).

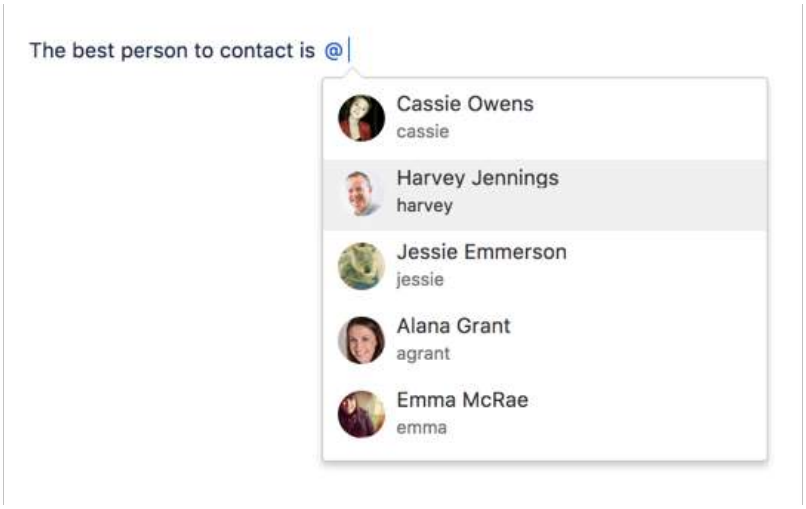

It then continues to suggest matches as you type. If you've not mentioned the person recently, we'll also include information about whether they've commented or contributed to the current page, to help you find the right person, fast.

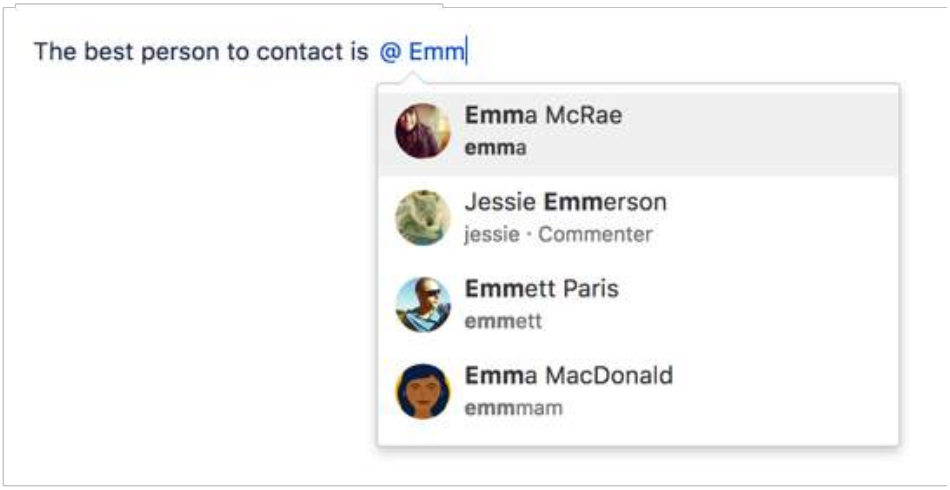

Use the Insert menu

If you'd rather use the **Insert**menu, choose **Insert** > **User Mention** then search for and select the user you wantto mention.

Notes

- **Disable mentions**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. SeeDisabling or Enabling a Plugin.
- **Mentioning groups** You can only mention individual users who have the 'Can Use' Confluence global permission. There's a feature request to allow mentions for groups:
  CONFSERVER-23015 - Extend 'Mentions' to work with groups as well FUTURE CONSIDERATION
- **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).
- **Mention notifications**- A notification is sent to a person the first time you mention them in the content of a page, but not for subsequent mentions. If you need to catch someone's attention, and you've already mentioned them on the page, try mentioning them in a comment. A notification is sent every time you mention someone in a page comment or inline comment, not just the first time.
- **Frequently mentioned people** - Confluence relies on your browser's local storage to remember the people you mention regularly. You may see different results if you switch devices, or don't allow local storage. Confluence doesn't indicate whether someone is a creator, contributor or commenter when they are also a recent mention (because we're grabbing them straight from your local storage, not the page itself).
- **Changing the mention name** - if you change the mention name in the editor (for example you backspace to remove their surname, or edit the mention link to change their full name to their preferred name) this will be treated as free text and won't be updated if the person changes their name, or is deleted from Confluence.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Share a Page or Blog Post

Use the **Share**button when viewing a page or blog post to email anyonea link to that page or blog post. You can either grab the short URL from the share dialog, or enter a Confluence user, group or email address.

To share a page or blog post by email:

1. Go to the page or blog post you wish to share.
2. Choose **Share**.
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 (or use the trash icons to remove people from the list).
4. Enter an optional message.
5. Choose **Share** to send the link via email.

Related pages:

- Workbox Notifications
- Create and Edit Pages
- Blog Posts
- Configuring a Server for Outgoing Mail
- Space Permissions Overview

In addition to an email, Confluence users will also receive a notification in their Confluence workbox. SeeWorkbox Notifications.

You can also share pages from inside the editor. Hit the + button in the editor to invite people to edit the page with you.

Notes:

- The option to add people is only available if your site has a mail server configured.
- Sharing or inviting someone to edit a page or blog post does not automatically grant any permissions - they will still need the appropriate Confluence permissions to access Confluence and view or edit the page.

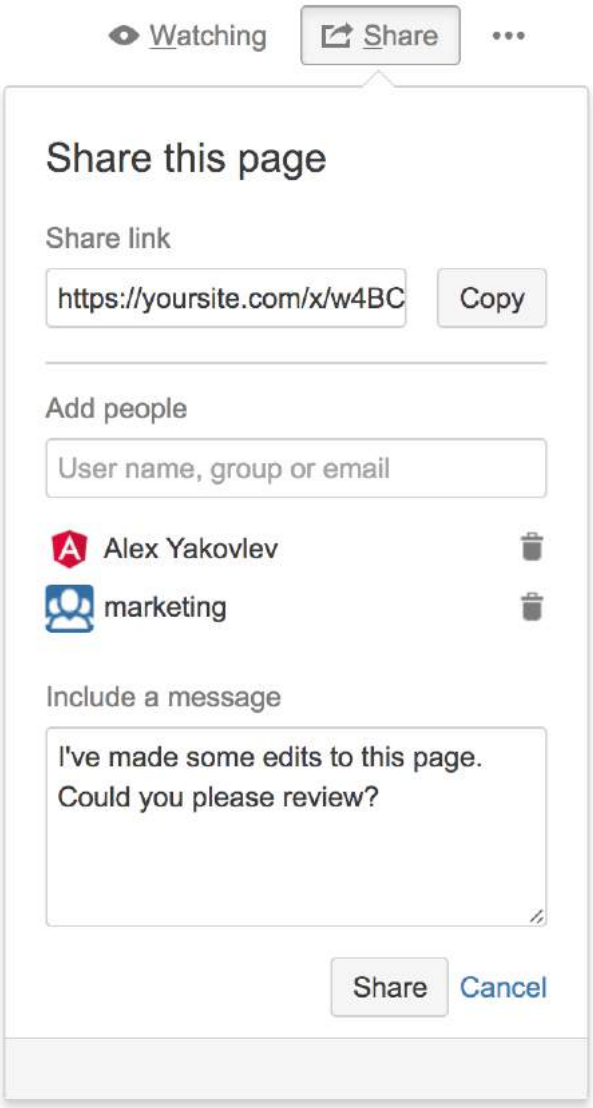

Comment on pages and blog posts

Comments are a great way to bring others into the conversation about a page or blog post.They allow you toremark on content, add important information, ask questions, and generally drive collaboration and teamwork.

You can add a commentat the bottom of any page or blog post, or add an inline commentby highlighting specific text on the page.

On this page:

- Add a page or blog post comment
- Add an inline comment
- View inline comments
- Resolve inline comments
- Rich comments
- Link to a comment
- Inline comments in the editor
- Comment permissions
- Notes

Related pages:

- Create and Edit Pages
- Blog Posts
- Share and Comment on Files
- Share a Page or Blog Post

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 notificationsabout 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)

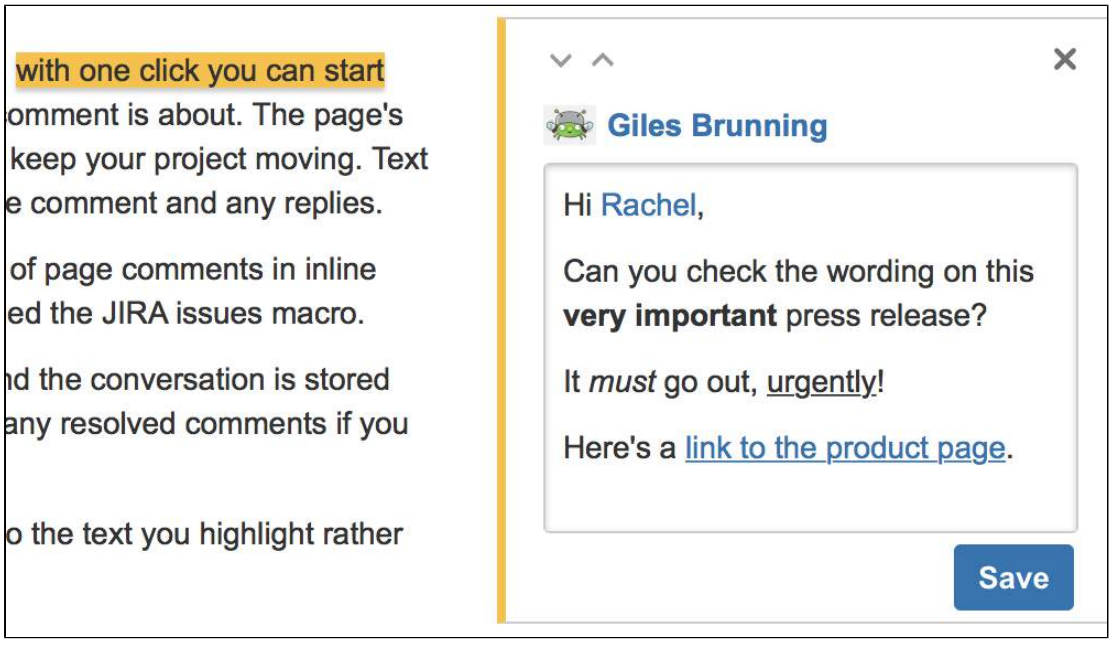

Screenshot: Create an inline comment by highlighting text.

View inline comments

Text will appear with a yellow highlightto indicate there is an inline comment; select 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.

An easy way to check whether there are any inline comments on a page is to select **View inline comments**fr om the page tools, or use the keyboard shortcut "v".

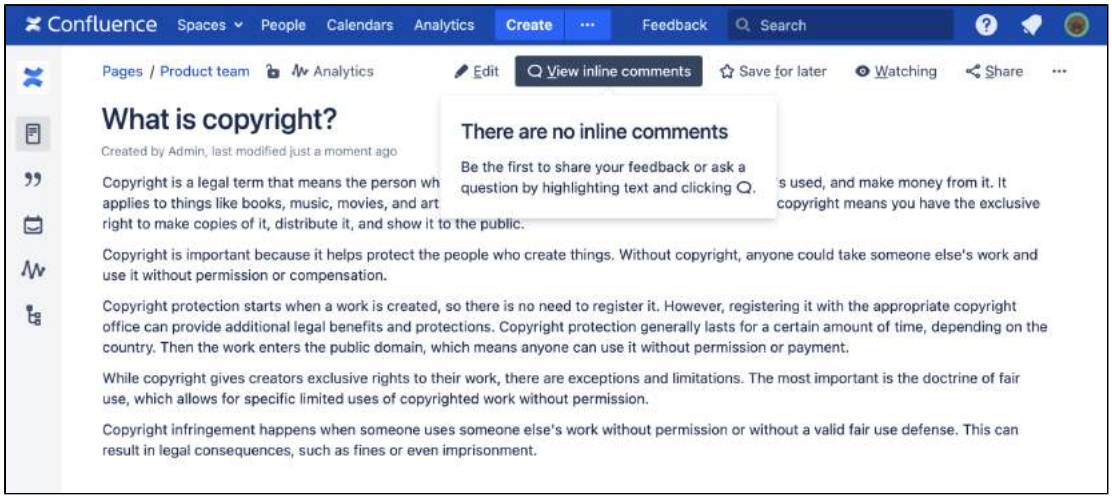

Screenshot: No inline comments on a page.

Resolve inline comments

Hit **Resolve** to hide the thread of inline comments once the conversation's finished. If you want to view resolved comments, select ••• >**Resolved comments.** To reopena resolved comment, choose**Reopen** at the bottom left.

Rich comments

Inline and page comments might look simple, but they support rich text (like**bold**,underline, and*italics),*bullete d 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. SeeLinksfor more information.

> If you don't see apopup when you highlight text, check thatText Selectis enabled in your profile settings.

Inline comments in the editor

When you are in the editor, there are limitations to how you can interact with inline comments. The below table summarises the differences in behavior when you are in viewing a page and editing it.

Feature	Viewing a page	Editing a page
View inline comments	✓	✓
Add and edit inline comments	✓	✗

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Resolve inline comments	✓	✓
Delete inline comments	✓	✓
Reply to inline comments	✓	✗
Like/unlike inline comments	✓	✗
Close/hide inline comments	✓	✓
Quick nav between inline comments	✓	✓
Link to inline comment	✓	✗

When an inline comment is added, resolved, or deleted on a page, this is reflected in the editor without you having to refresh the page. These changes may take up to a minute to reflect on the page.

Comment permissions

- **Add a comment** You need the 'Add Comments' permission in the space.
- **Edit a comment** You need the 'Add Comments' permission.Space administratorscan 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.

SeeSpace permissionsfor more information.There is no permission that controls comments across the entire site.

Members of the Confluence-administrators group can add, edit and delete comments, even if you remove their comment permissions in the Space permissions configuration.

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 changethe order of comments.
- Inline comments on text that is included on a page using the Include macro or Excerpt include macro won't be visible. They are only visible on the original page.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

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:

- Page / blog post 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 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 behavior is called 'autowatch'.

There's no daily digest for email notifications.You'll receive immediate emails for important notifications (like mentions and new pages), but when lots of changes are being made at the same time, you'll only receive a single email with all the changes (within a 10 minute window).

You **willnot**receive email notifications for content changes due to the output of a macro, because the page content itself hasn't been edited. We also don't send a notification when a comment is edited.

You need 'View' permission for thepage, blog post or space to receive notifications.

On this page:

- Watching a page or blog post
- Watching an entire space
- Watch for new blog posts in a space
- Watch all spaces on the site
- Watching for all new blog posts on the site
- Manage watches from your user profile
- Manage watches from the email message
- Autowatch and other notification options

Related pages:

- Manage Watchers
- Email Notifications
- Your User Profile

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 wayis 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.

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 when a blog post is updated, deleted or commented on.

To watch for new posts:

1. Go to a blog post in the space.
2. Choose**Watch**and select **Watch for new 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.

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**yourprofile 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

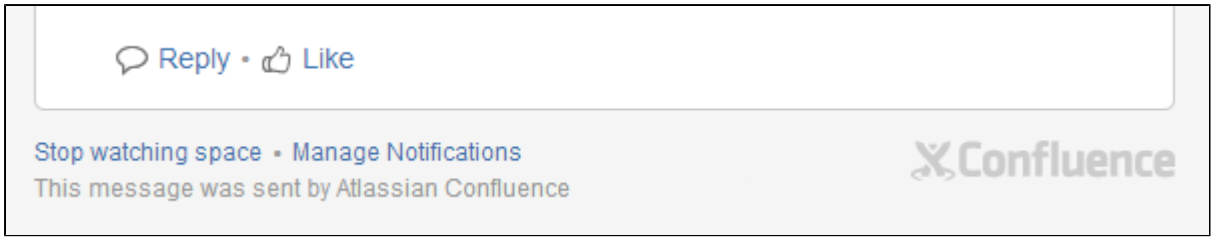

Autowatch and other notification options

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

By default, Confluence will assign you as a watcher of any page or blog post that you create or edit. This behavior is called 'autowatch'. You can turn autowatch on or off, and set other notification options, in the email settings section of your user profile. SeeEdit Your User Settings.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Manage Watchers

As a space admin, you may want to control who's notified about changes and updates to pages and posts within a space. You'll do this by managing the watchers of specific pages and posts, or of the entire space.

Take the example of a new member starting on your team, who should see when changes are made in the team space you can add them as a space watcher so they get updates when any page or post in the space is changed. If someone no longer needs to be notified, you can remove them as a watcher just as easily.

Related pages:

- Watch Pages, Spaces and Blogs
- Email Notifications
- 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' dialog 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:
   - Add someone as a watcher of the page, post, or space type their username in the relevant search box and hit**Add**
   - Remove an existing page, post, or space watcher choose the trash iconnext to their name

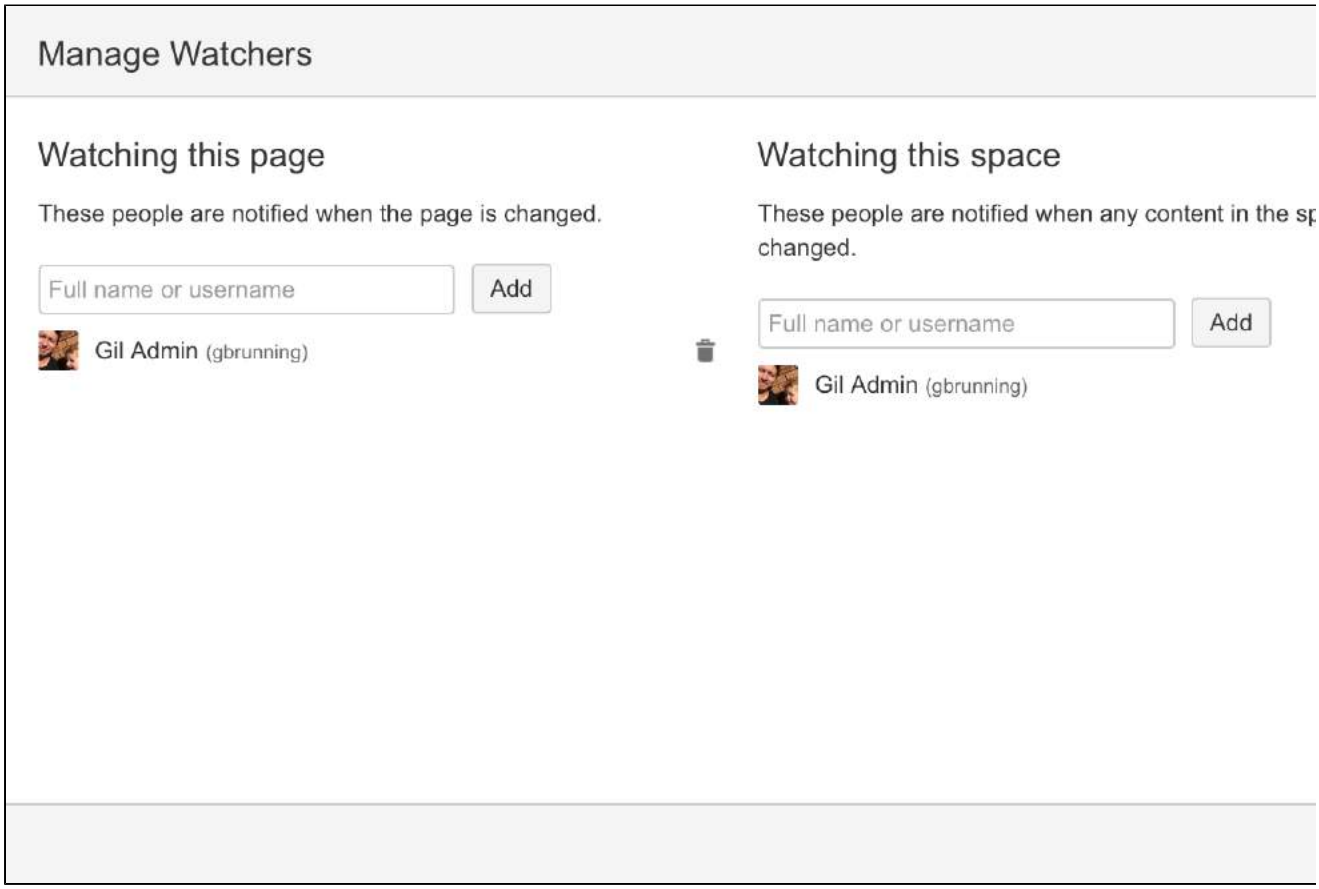

Email Notifications

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'll receive immediate emails for important notifications (like mentions and new pages), but when lots of changes are being made at the same time, you'll only receive a single email with all the changes within a short window (usually 10 minutes).

You can also subscribe to daily email reports and other notifications of various updates, as described below.

You'll only receive notifications for content that you have permission to view. Users that have been disabled by an administrator will not receive email notifications.

On this page:

- Subscribing to email notifications
- Notes

Related pages:

- Watch Pages, Spaces and Blogs
- Subscribe to RSS Feeds within Confluence
- Your profile and settings
- Edit Your User Settings

Subscribing to email notifications

You can subscribe to be notified when:

- A blog post is added or changed in a space that you have permission to view.
- Someone you're following makes an update in a space that you have permission to view.
- Someone follows you.

You can alsosubscribe to these summary reports:

- A daily report of the 30 most popular updates to all spaces that you have permission to view.
- 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**

Here's an explanation of all the email settings.

Setting	Description	Content
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.

Subscribe to daily updates	Receive daily 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. Up to 30 updates will be included, sorted by popularity.*	• 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.
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.
Subscribe to network	Receive email notifications for changes to content by all users that you are following, which you have permission to view.	• Space is created. • Page or blog post is created.
Subscribe to new follower notifications	Receive an email message when anyone chooses to follow you.	
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.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

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.
Subscribe to recommended updates	Receive a daily or weekly 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.

Notes

- **Mail server**: To enable Confluence to send email notifications, a System Administrator must configure an email server. SeeConfiguring a Server for Outgoing Mail.
- **Batching window:** System Administrators can change the batching window for changes and comments on the same page or blog post in the **Send batched notifications**scheduled job. Increase the time for fewer emails or reduce the time if more immediate notifications are essential in your site.
- **Recommended updates email**: Confluence Administrators can set the default options for therecommended updates notification.

  Choose the**cog icon** , then choose**General Configuration**

  .Click **Recommended Updates Email** in the left-hand panel. See Configuring the Recommended Updates Email Notification.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

- **Email threading**: Confluence will attempt to group all the email notifications about changes to a specific page together. Other notifications such as sharing a page, requesting access to a page, or recommended updates emails are intentionally not grouped. Not all mail clients support email threading, anddifferent email clients use different methods for threading emails. We've tested Confluence's email threading withApple Mail 10.3, Outlook 2011, Outlook 2016, GMail, Google Inbox and Outlook.com.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

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. Specialized 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, a browser extension, 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.

On this page:

- Confluence RSS feeds
- Remove an RSS feed

Related pages:

- Watch Pages, Spaces and Blogs
- RSS Feed Macro

Confluence RSS feeds

RSS feedsallow you to track updates to content within Confluence. You will need an RSS newsreader to read a feed.

You can create a customized 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 customized RSS feeds using the RSS Feed Builder Create a customized 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.

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.

Related pages:

- The RSS Feed Builder
- RSS Feed Macro
- Network Overview

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

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, yourNetworkview 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 customize your Confluence RSS feed (for example, use a label to filter your feed), use theRSS Feed builderinstead of the above instructions.

The RSS Feed Builder

Using the RSS feed builder, you can create customized RSS feeds to subscribe to changes within Confluence.

Wondering what an RSS feed is? See more information about RSS Feeds.

On this page:

- Build an RSS feed
- Notes

Related pages:

- Watch Pages, Spaces and Blogs
- Subscribe to RSS Feeds within Confluence

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 customized 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 customize your feeds.
Exclude these spaces	Exclude specific spaces from those already selected.
Sorted by	While you can choose to sort by creation or the date they were last updated, there is a known issue where the feed is always sorted by the last updated date. See ~~CONFSERVER-52542~~ CLOSED .
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 customize your feed.**You can use the RSS feed builder to track updates to labeled pages and comments on those pages. Here is an idea for customizing 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 labeled 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=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.

- **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. SeeCONF-21601for details and a workaround.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

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
- Customize 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

Customize your network RSS feed

Confluence does not provide a way of customizing 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.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

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 aJiraapplication such as Jira Software or Jira Service Management, 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

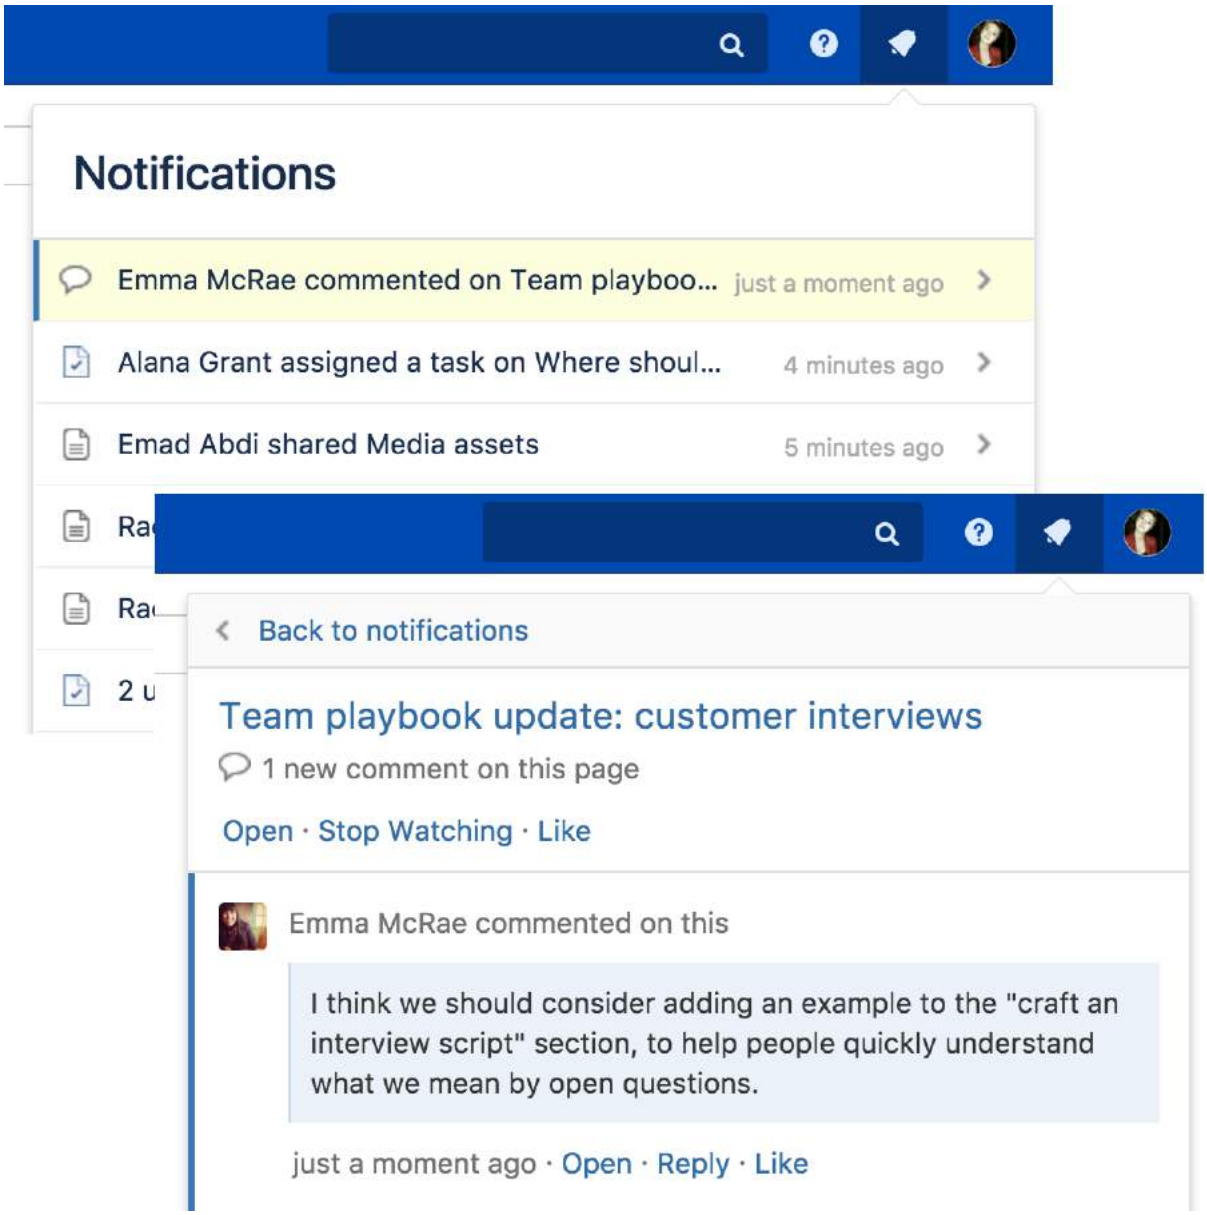

Which notifications are included?

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

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

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.
- TheConfluence 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 theworkbox for an individual user.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Analytics

This feature is available with a Confluence Data Center license.

Use analytics data to understand how yoursite,spaces, andpages are performing, and how active your users are. You can also export reports to share with your team.

On this page:

- How to access analytics
- Limit who can view analytics reports for specific spaces
- Analytics for Confluence marketplace app

Analytics data is useful for identifying:

- most viewed content
- usage and activity trends across your site
- frequent contributors and viewers
- spaces with few views that could potentially be archived.

If required, you can set who can view analyticson your site and spaces.

How to access analytics

Analytics is only available with a Confluence Data Center license.

Your administrator may have limited who can view analytics reports. SeeAdminister analytics for more information.

Access Analytics from the Confluence navigation

When you access analytics from the navigation, you getan overview of your site's usage and access to insights on spaces, users, and searches.

Go to Analytics in the Confluence header.

SeeView insights on your site for more information.

Access Analytics from a space

When you access analytics from a space, you get an overview of all the pages and attached files in that space.

Go to Analytics in the space sidebar.

SeeView insights on spaces for more information.

Access Analytics from a page

When you access analytics from a page, you can get a quick summary of insights on your page, such as total page views. From there, you can view more detailed page analytics, such as which users have viewed the page, which versions of the page they viewed, and more.

At the top of any page, clickAnalytics. ClickView analyticsto seethe detailed analytics about your page.

SeeView insights on pages for more information.

Limit who can view analytics reports for specific spaces

You can further limit who can view analytics reports for specific spaces. You need Space Administrator permission to do this.

To change who can see analytics reports in a space:

1. Go to the space and choose **Space tools** > **Permissions** from the bottom of the sidebar.
2. Select the **Analytics Permissions** tab.
3. Select **Viewing Analytics Restricted** from the drop down.
4. Enter the users and/or groups you want to allow to view analytics reports.
5. Save your changes.

Good to know:

- If a Confluence administrator has denied a group permission to view analytics, adding the group at the space level will not grant this permission.
- The space will still appear in the site analytics report (if the user has permission to see the space, and use analytics globally), but they will be prevented from viewing the space analytics report if they don't have space permissions.

Analytics for Confluence marketplace app

If you previously installed the Good Software Analytics for Confluence marketplace app, you can continue to use the app until it reaches end of life, but your version may not contain the same features.

Analytics for Confluence is now part of Confluence Data Center
To get access to the features described on this page upgrade to Confluence Data Center 7.11 or later. See our FAQ for all the details

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

View insights on your site

Get an overview of the usage of your site, and view insights on spaces, users, and searches.

To view the analytics report for your site:

1. Go to**Analytics**in the header.
2. Use the filters at the top to choose the date range for your data. You can also choose how you want to group the data, whether to display it by total number of users or unique users, and more.
3. Hover over the chart to view a summary of data for a specific date.
4. Below the chart, see the lists of your most popular spaces, most active readers, and most active contributors.

On this page:

- View insights on your site's users
- Export analytics reports

Related pages:

- Analytics

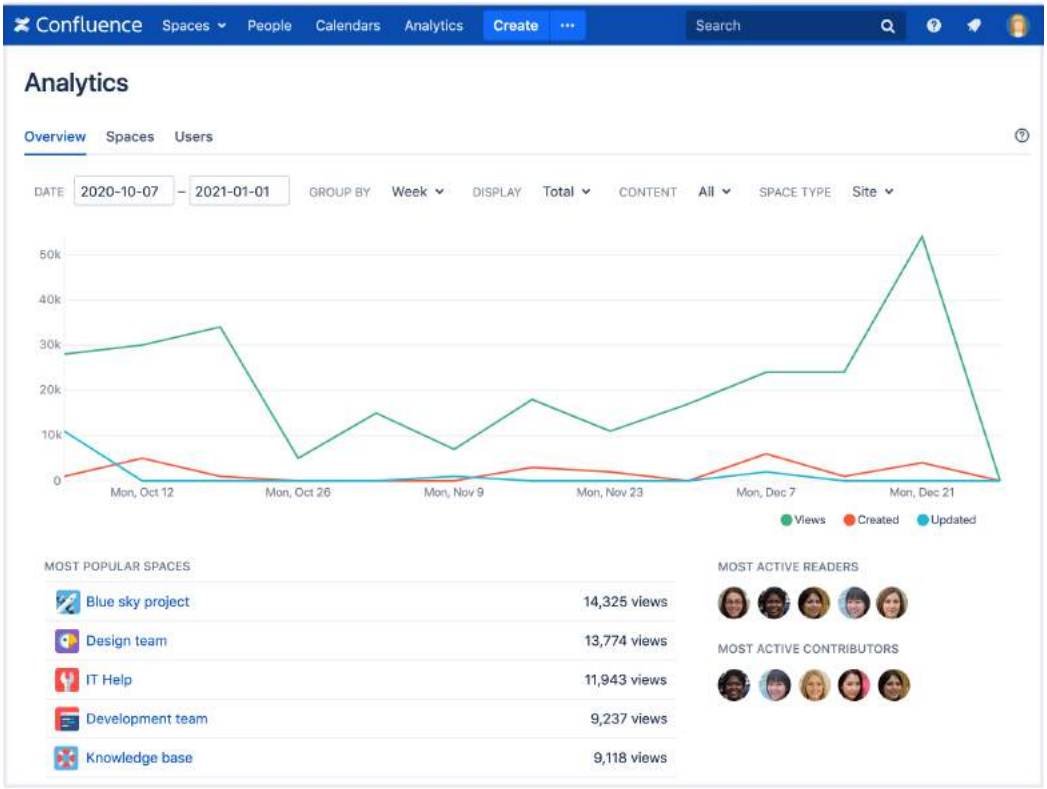

Screenshot showing the site analytics overview tab.

View insights on your site's users

To get an overview of and insights on how active users are on your site:

1. Go to **Analytics** in the header
2. Select the **Users**tab.
3. Use the filters at the top to choose the date range for your data. You can also choose the type of contentand space you want to get data on. For example, you can view users' activity on pages in site spaces.

Because sites generally contain many thousands of users, we display users in batches.Only the first 200 users (sorted by views) will appear in the list. Sorting the columns will query the whole list, not just the users displayed.

From the**Users**tab, you can view the following insights:

- **Name**- The name of the user (or an anonymized alias)
- **Created**-The number of pages has created for the date range
- **Updated**-The total number of times the user has edited pages for the date range

- **Comments**-The total number of times a user has commented on pages for the date range
- **Search**-The total number of times a user has performed a search for the date range
- **Views**-The total number of times a user has viewed pages for the date range

Your administrator may have enabled increased privacy mode, which replaces people's real names with an anonymized alias and avatar. Each alias represents an individual user, so you can still get an idea of engagement, but not attribute activity to particular people.

Export analytics reports

You can export and download some reports in Microsoft Excel format. Export is available on the Users and Search tabs.

1. Select the ellipses ••• on the top right of the page.
2. Select **Export Excel**.

Analytics for Confluence is now part of Confluence Data Center
To get access to the features described on this page upgrade to Confluence Data Center 7.11 or later. See our FAQ for all the details

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

View insights on spaces

Get an overview of how each space in your site is performing.

1. Navigate to a space.
2. Select**Analytics**in the space sidebar.
3. Use the filters at the top to choose the date range for your data. You can also choose the type of content and space you want to get data on.

Alternatively you can go to**Analytics**in the header, then choose a space.

On this page:

- View insights on the pages in a space
- View insights on the users of your space
- Export analytics reports

From the**Overview**tab, you can view the following insights:

- **Created**-The number of pagesthat have been created within the space for the date range
- **Updated**-The total number of times pageshave been edited within the space for the date range
- **Views**-The total number of views of pageswithin the space for the date range

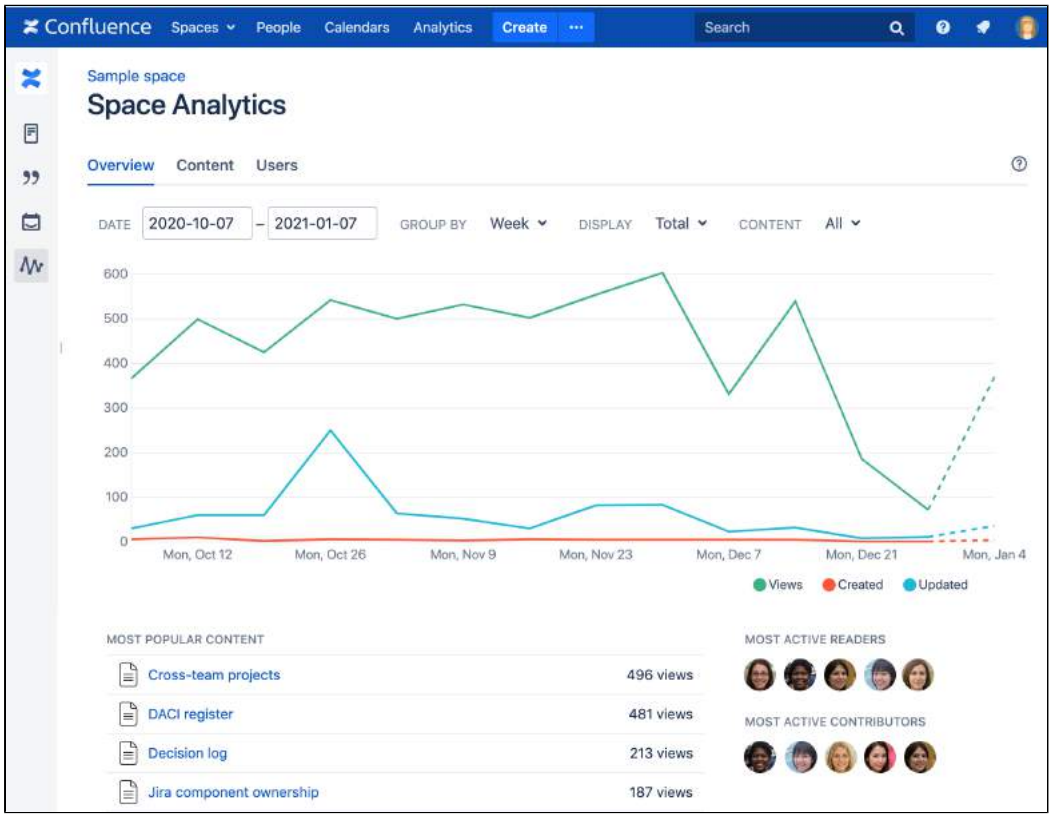

Screenshot showing the space analytics overview tab.

View insights on the pages in a space

Get an overview of how your space's pages and blogs are performing.

1. Select **Analytics** in the space sidebar.
2. Select the **Content**tab.
3. Use the filters at the top to choose the date range for your data and the type of content you want to get data on.

From the**Content** tab, you can view a table with:

- **Page name**-The name of thepage
- **Created**-The date the page was created
- **Last modified**-The last time the page was edited for the date range
- **Last viewed**-The last time a user has viewed the page for the date range
- **Comment activity**-The total number of comments on the page for the date range
- **Usersviewed**-The total number of unique users that have viewed the page for the date range
- **Views**-The total number of views for the page for the date range

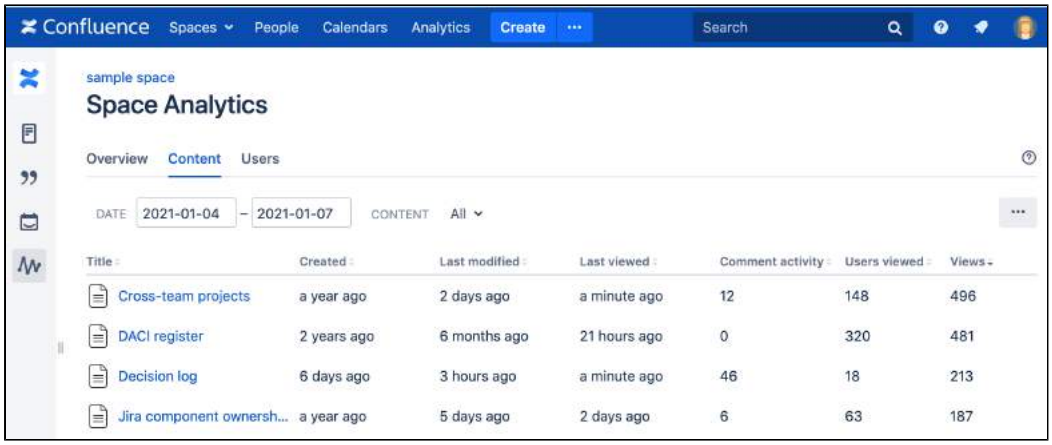

Screenshot showing the space analytics report content tab.

View insights on the users of your space

To get an overview of how active users are on your space:

1. Select **Analytics** in the space sidebar.
2. Select the **Users**tab.
3. Use the filters at the top to choose the date range for your data and the type of content you want to get data on.

From the**Users** tab, you can view a table with:

- **Name** - The name of the user (or anonymized alias)
- **Created**-The number of pages the user has createdfor the date range
- **Updated**-The number of times the user has edited pagesfor the date range
- **Comments**-The number of time the user has commented on pagesfor the date range
- **Views**-The total number of times the user has viewed pages for the date range

Your administrator may have enabled increased privacy mode, which replaces people's real names with an anonymized alias and avatar. Each alias represents an individual user, so you can still get an idea of engagement, but not attribute activity to particular people.

Export analytics reports

You can export and download reports in Microsoft Excel format. Export is available on the Overview, Content, and Users tabs.

1. Select the ellipses ••• on the top right of the page.
2. Select**Export Excel**.

Analytics for Confluence is now part of Confluence Data Center
To get access to the features described on this page upgrade to Confluence Data Center 7.11 or later. See our FAQ for all the details

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

View insights on pages

Get detailed analytics of your pages and blogs to determine how they're performing.

1. Go to the page whose analytics you'd like to view.
2. Select**Analytics**(above the page title)
3. Use the filters at the top to choose the date range for your data. You can also choose how you want to group the data and whether to display it by total number of users or unique users.
4. Hover over the chart to get the number of views for a specific date.

From the page's analytics, you can view a table with:

- **Name**- The name of the user (or anonymised alias)
- **Last version viewed**-The latest version of the pagethat the user has viewed
- **Last viewed**-The last time the user viewed the page
- **Views**-The total number of times the user viewed the page

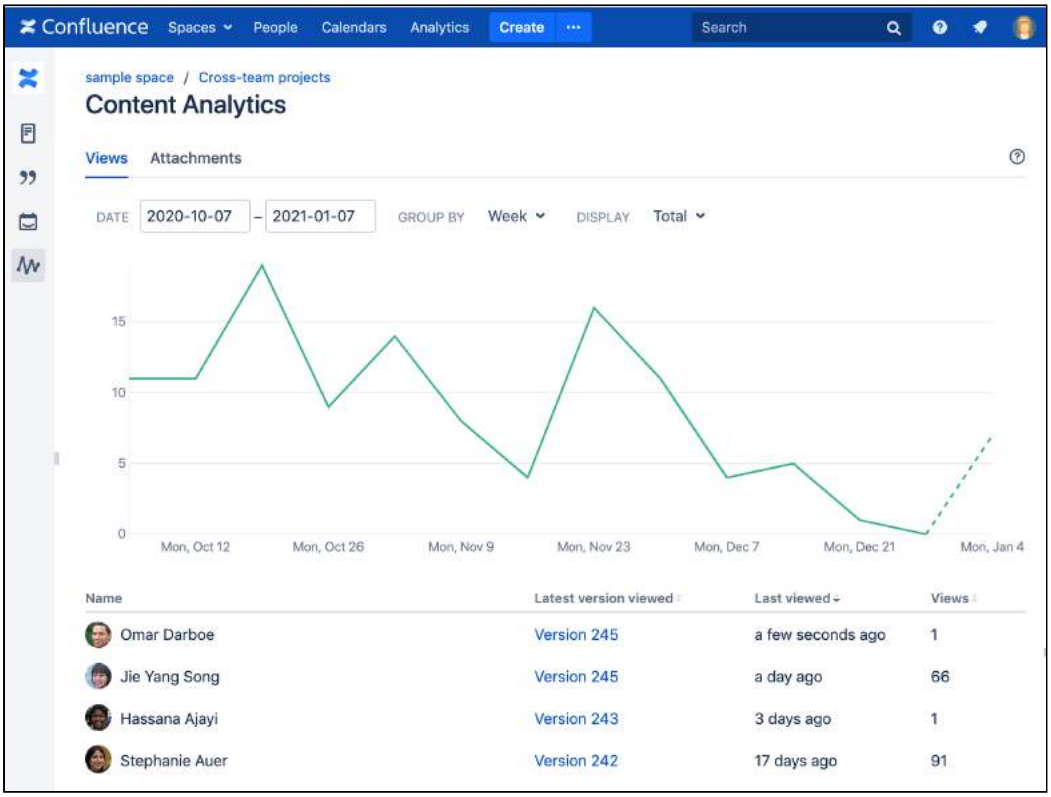

Screenshot showing the page analytics view tab, sorted by last view date.

Your administrator may have enabled increased privacy mode, which replaces people's real names with an anonymized alias and avatar. Each alias represents an individual user, so you can still get an idea of engagement, but not attribute activity to particular people.

View insights on attached files

Select the**Attachments**tab for insights on files attached to it.

The**Attachments**tab provides the following insights for each file on the page:

- **Name**-The filename name of the attached file
- **Last viewed**-The last time a user has viewed the attached file
- **Views**-The total number of views for the attached file

Analytics for Confluence is now part of Confluence Data Center
To get access to the features described on this page upgrade to Confluence Data Center 7.11 or later. See our FAQ for all the details

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

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.

On this page:

- How Confluence search works
- Start a search
- Filter your search results
- Search for admin options
- Advanced search
- Advanced search syntax

Related pages:

- Confluence Search Syntax
- Confluence Search Fields
- Ranking of Search Results
- Recently Viewed Pages and Blog Posts
- Search Results Macro
- Livesearch Macro
- Search the People Directory

How Confluence search works

When you enter a search term, Confluence looks for content in all spaces (including personal spaces), pages, mail, personal profiles, and space descriptions. It also looks at the content of some attached file types (Word, Text, PowerPoint, Excel, PDF, and HTML).

Search results are based on your Confluence permissions, so you'll only see content you're allowed to view.

Learn how Confluence ranks search results

Start a search

To search Confluence:

1. Click the search field in the top-right of Confluence to open the search panel.
2. Start typing your search term.

Results will appear as you type you don't need to hit enter.

We exclude comments from your search results unless you select the comment option from the type filter.

Screenshot: the search panel

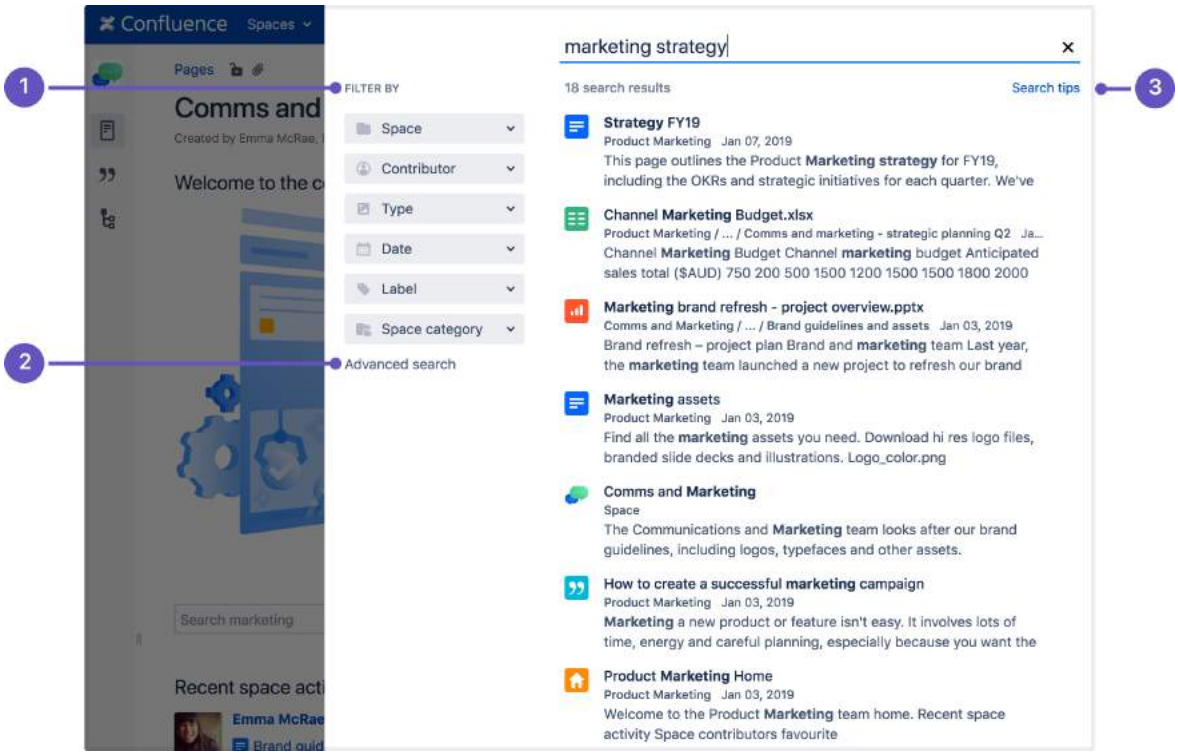

1. **Search filters** refine your results by space, contributor, type, date, label, or space category.
2. **Advanced search** go the the advanced search page.
3. **Search tips** get search help, and tips for refining your search.

Tip: Type **/** on your keyboard to quickly open the search panel.

Filter your search results

You can refine your search using interactive filters on the left-hand side of the search panel.

Search within a space

Use the **space** filter to find content within a particular space or list of spaces. The space you're currently in will appear at the top of the list by default. Start typing the space name and choose from the list of suggested spaces.

Click the toggle to search within archived spaces.

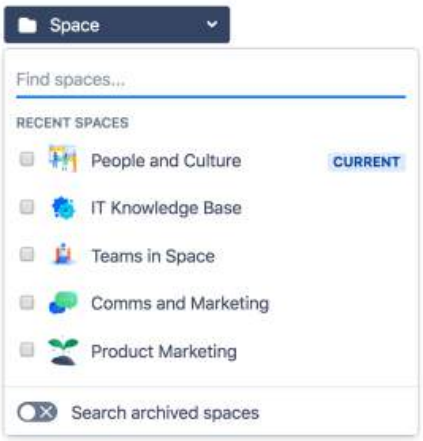

Search for content by a specific person

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Use the **contributor** filter to restrict your search to content modified (created, updated, or commented on) by particular people. Start typing the person's username, or part of their name, and we'll show you a list of possible matches. You can add as many people as you like.

> **Tip:** To search for your own work, click the **Contributor** filter, then select your profile from the drop-down menu. Your name appears here by default, so it's easy to find.

Filter by content type

Use the **type** filter to only show content of a certain type, such as pages, blog posts, comments or user profiles.

Search within a specific time frame

Use the **date** filter to search for content last modified (created, updated, or commented on) within a particular period of time.

To search within a specific date range:

1. Click **Advanced Search** on the left-side of the search panel.
2. In the **Last modified** section, choose**Custom**.
3. Select the date range from the drop-down date picker.

Search for content with a specific label

Use the **label** filter to search for content containing a specific label. Start typing the name of the label and choose from the list of possible matches.

Search within a space category

Use the space category filter to search within a group of related spaces.Start typing the category name and choose from the list of possible matches. You can browse existing categories from the Space directory.

> **Tip:** Space admins can organize spaces into categories. You can create space categories for departments, subject areas, office locations whatever works for your team. Learn how to create a space category

Search for admin options

As a Confluence admin, you can can quickly access admin options from the search panel.

Start typing what you want to do. We'll show the top three matching admin items at the top of your search results. You'll only see options you have permission to perform.

If you apply a search filter, admin items will no longer appear in your results.

Advanced search

The advanced search page allows you to add more search filters, such as creator, title, date range or ancestor page.

To use advanced search:

1. Click the search field at the top-right of Confluence.
2. Click **Advanced search**on the left-side of the search panel.
3. Type your keyword in the search field and hit enter.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Adding a filter from the advanced search page

To add a search filter from the advanced search page:

1. Click **Add a filter** on the bottom left of the advanced search page.
2. From the drop-down menu, select the relevant filter.

Screenshot: advanced search page

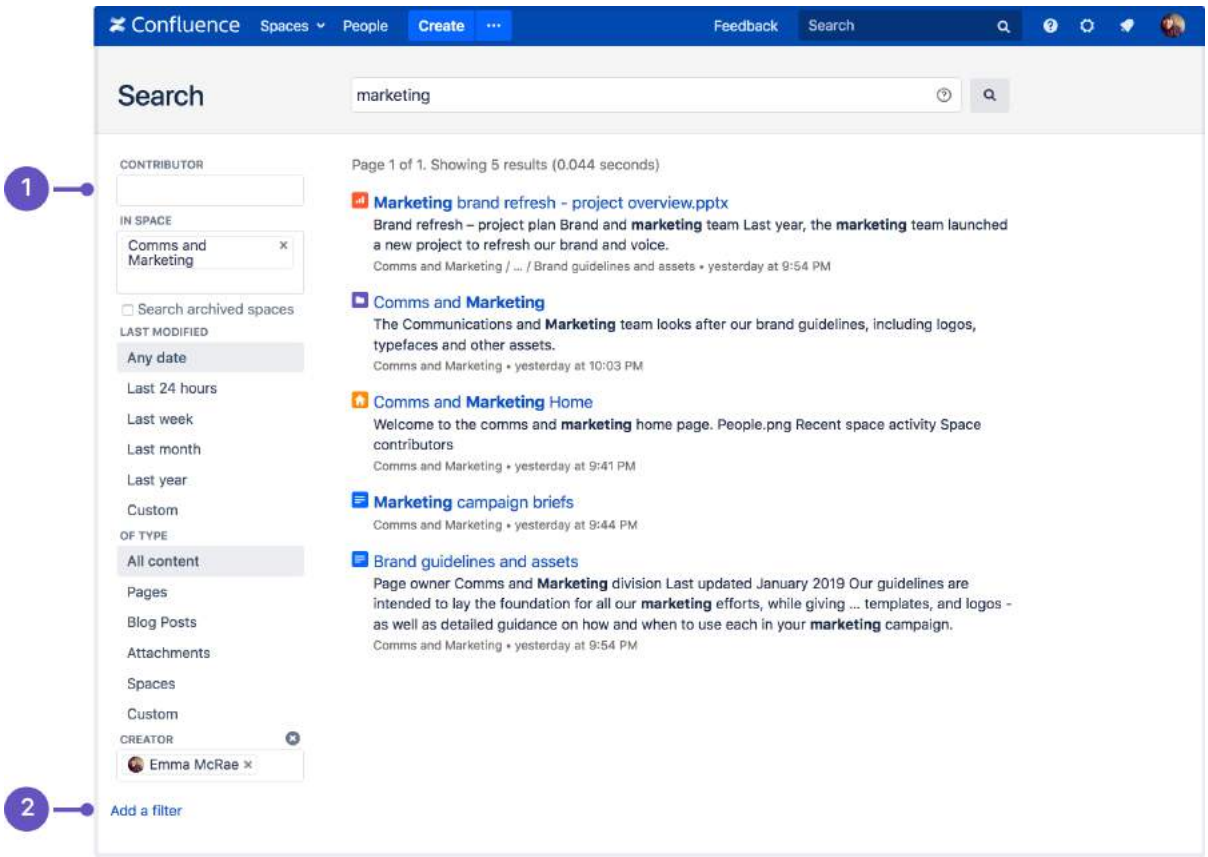

1. **Search filters** narrow your search by adding filters.
2. **Add a filter** add more filters for even more precise search.

You can choose from these additional filters:

- **Creator**Restrict your search to content created by a particular person. Start typing the person's username or part of their name andConfluence will offer you a list of possible matches.
- **Label** Only search for content containing specific labels.
- **With parent**Only search for direct children of a specific parent page.
- **With ancestor** Only search for pagesbelow a certain page in the hierarchy.
- **Created**Choose or enter a date to only show content created within a particular period of time.
- **Mentioning user**Only search for content that mentions a particular Confluence user.
- **With title**Only search within page or blog titles.

These filters are provided byConfluence Query Language (CQL).

Advanced search syntax

You can also refine your search using Confluence search syntax. These are words or symbols you enter into the search field to help narrow down your results. Learn more aboutConfluence Search Syntax.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Confluence Search Syntax

You can create an advanced search query using Confluence search syntax. These are special words and symbols you enter into the search field to narrow the focus of your search.

This page outlines the syntax supported by Confluence's search engine, Lucene.

On this page:

- How to use search syntax
- Search for an exact match
- Search using wildcards
- Exclude words from your search
- Combine search terms
- Search for nearby words (proximity search)
- Search within an alphabetical range
- Search for words spelled similarly (fuzzy search)
- Combining search operators
- Searching for macros
- Search specific fields in Confluence
- Confluence search fields

Related pages:

- Search
- Confluence Search Fields
- Search the People Directory
- Ranking of Search Results

How to use search syntax

To create a search query using Confluence syntax:

1. Click the search field at the top right of Confluence to open the expanded search panel.
2. Type your query using syntax supported by Confluence.

You can use multiple search words and operators in your query.

Screenshot: an example of a search query using Confluence search syntax

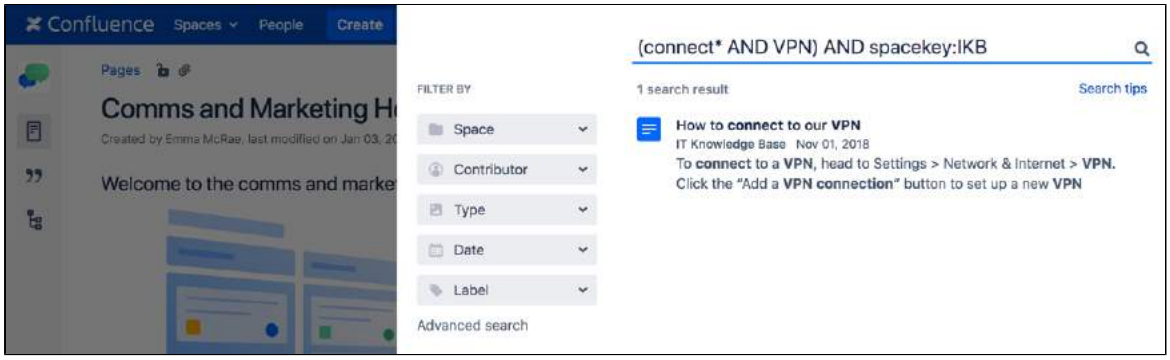

Search for an exact match

Use double quotes around your search term to find a specific word or phrase. For example "product roadmap" will search for content that contains the phrase 'product roadmap', or a phrase where 'product' and 'roadmap' are the major words.

```
"product roadmap"
```

Limitations with exact match search

Phrases with stop words

Confluence ignores common words (stop words) such as 'and', 'the', 'or', and 'it' even if they are included within double quotes.

For example, searching for "the IT budget" will only return pages containing 'budget', because 'the' and 'it' are stop words.

If you'd like to change this, vote on this improvement request:

CONFSERVER-14910 - Provide ability to override Lucene tokenisation and stemming and search for exact text (literal search) IN PROGRESS

Phrases with special characters

Confluence ignores all symbols, such as hyphens or underscores, even if they are included within double quotes.

For example, if you search for "DOC-8510", you get all pages containing 'doc' and '8510'.

Avoid using special characters, such as hyphens, in page or attachment names as they may not be found by Confluence search.

Search using wildcards

Wildcards replace one or more characters in your search. They can help expand your search. For example, the search below would find https://www.atlassian.com or http://www.atlassian.jp

```
http*.atlassian.*
```

Confluence doesn't support leading wildcards. This means searching for `*heesewill` not return cheese.

Wildcard	Description	Example

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Multiple characters	Use an asterisk (*) at the end of your word to replace multiple characters.	**print*** finds content containing 'printer', 'printing', 'prints' and so on.
Multiple wildcards	Use asterisks (*) to add more than one multiple-character wildcard	**r*c*** finds content containing 'react', 'recovery', 'refactor' and so on.
Single character	Use a question mark (?) to replace a single character in your search.	**b?tter** finds content containing 'butter', 'bitter', 'better', 'batter' and so on.

Leading wildcards

Lucene doesn't allow wildcards at the beginning of your search, but you can format your search as a regular expression as a workaround. For example, you can't search for `*hum*` or `?hum*`, as they begin with a wildcard, but you can search for /.*hum.*/ and find things like hum, human, and inhumane.

Exclude words from your search

Use NOT or minus (-) to exclude words from your search.

```
chalk NOT cheese
```

Operator	Description	Example
NOT	Use NOT (in capital letters) to exclude a word from your search.	**chalk NOT cheese** finds content containing 'chalk' but NOT 'cheese'
Minus (-)	Put a minus sign (-) in front of words you want to leave out.	**chalk butter -cheese** finds content containing 'chalk' and 'butter' but not 'cheese'

Combine search terms

Operator	Description	Example
OR	Use OR (in capital letters) to search for content that contains one of the terms.	**chalk OR cheese** finds content containing either 'chalk' or 'cheese'
AND	Use AND (in capital letters) to search for content that contains more than one search term.	**chalk AND cheese** finds content containing both 'chalk' and 'cheese'

You can also combine search terms and operators, for example:

```
(cheese OR butter) AND chalk
```

Search for nearby words (proximity search)

Use a tilde (~) 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 won't work, 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.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

```
"octagon post"~0
```

Search within an alphabetical range

Use 'TO' (in capital letters) to search for names that fall alphabetically within a specified range. For example:

```
[adam TO ben]
```

Note: You can't use the AND keyword inside this statement.

Search for words spelled similarly (fuzzy search)

Use a tilde (~) to find words spelled similarly, or to pick up misspellings.

For example, if you want to search for octagon, but you're not sure how it's been spelled, type the word followed by a tilde:

```
octogan~
```

Combining search operators

You can also combine various search terms together:

```
o?tag* AND past~ AND ("blog" AND "post")
```

Searching for macros

You can search Confluence pages to find where a macro is used. Start your search with `macroName:` and type the macro name after the colon. For example, to search for all excerpt-include macros:

```
macroName:excerpt-include*
```

Search specific fields in Confluence

Confluence data is stored in fields, for example title, label, type and so on. To search for content using a specific field, type the name of that field into the search box followed by a colon (:), and then the term you're looking for.

You can use multiple fields in the same query. For example, you could use the following query to find all blog posts containing the Excerpt Include macro.

```
type:blogpost AND macroName:excerpt-include*
```

Confluence will only look for the term directly after the colon. For example, the query below will search for 'some' in the title field and 'title' in the default fields:

```
title:some title
```

Use double quotes if you want to find multiple keywords:

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

```
title:"some title"
```

Confluence search fields

This table lists some common search fields, and shows examples of what to type into the search box. Fields are case sensitive, so make sure you type the field name exactly as it appears in the table below.

Field	Description	Examples

macroName	Searches for pages that contains a specific macro. Type the name of the macro in lowercase. You can use a wildcard to make sure Confluence finds the macro you're after.	`macroName:excerpt-include*`
		
		`macroName:jira*`

spacekey	Searches for content within a specific space, using the space key. Type the name of the space key in capital letters.
	
	You can add multiple spaces using brackets and commas.
		`spacekey:MARKETING`
		
		`spacekey:(IT, MARKETING)`

title	Searches for content with specific words in the title.	`title:"product roadmap"`

type	Searches for content of a particular type. You can use the following content types in your query:
	
	• page
	• blogpost
	• attachment
	• comment (only supported when using advanced search).
		`type:attachment`
		
		`type:blogpost`

labelText	Searches for content containing a specific label. If the label has a hyphen, include it within double quotes.	`labelText:roadmap`
		
		`labelText:"product-roadmap"`

For more information aboutsearch fields, seeConfluence Search Fields.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Confluence Search Fields

This page gives an overview of the Apache Lucene search fields used in Confluence.

On this page:

- Filter with CQL
- Searching for content in specific fields
- Confluence search fields
  - Personal Information
  - Pages
  - Blog
  - Attachments
  - Mail items
- Notes

Related pages:

- Search
- Confluence Search Syntax
- Search the People Directory
- Ranking of Search Results

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).

> 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.
>
> - 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:
>
> 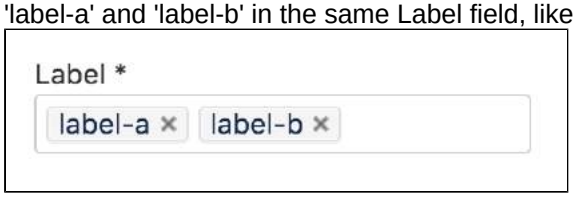
>
>
> - For an AND search, add more than one filter and specify a single value in each.
> To show only pages with label-a and label-b you'd put 'label-a' in one label field, then add a second Label field to the macro, and put 'label-b' in the second one, like this:
>
>
>
>
> Put simply, OR values are 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.

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. For example:

```
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.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

```
title:Some Heading
```

To learn more about using Confluence search fields in an advanced search query, head to Confluence Search Syntax.

Confluence search fields

Below are the fields which can be searched, listed by content type.

Personal Information

Name	Indexed	Stored	Tokenized	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	Tokenized	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	Tokenized	Notes
handle	true	true	false	
type	true	true	false	

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

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	Tokenized	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	Tokenized	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	

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

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.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Search the People Directory

The people directory displays a list of people who are authorized to log in to your Confluence site (they have the 'Can Use' global permission).

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.

On this page:

- View the people directory
- Search for people
- Follow people's activities
- Notes

Related pages:

- Create a Personal Space
- Editing your User Profile
- Set Your Profile Picture

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 apersonal 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 toNetwork 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 yournetwork view.

Screenshot: The people directory

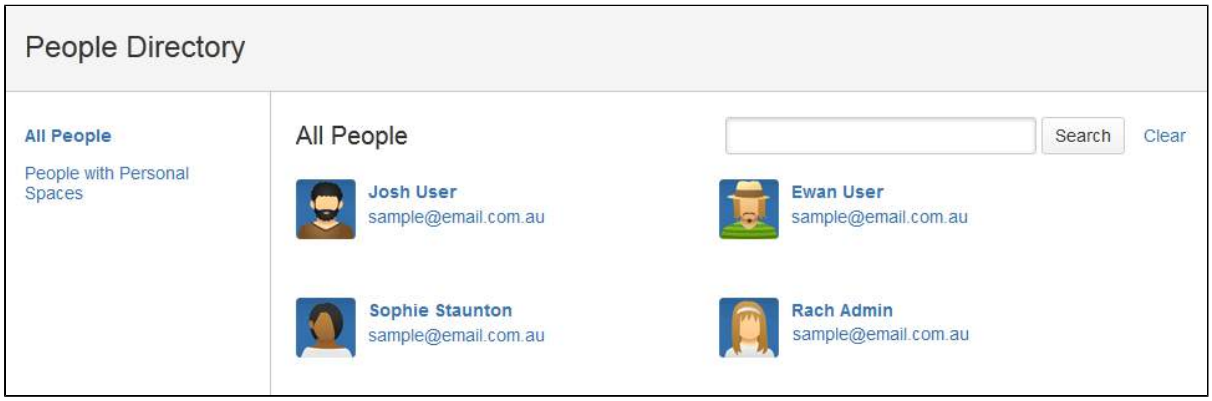

Notes

- Thepeople directoryuses the hCard microformat for simple integration with a variety of microformat-enabled tools. hCard is an open data format for representing people, companies, organizations, and places. Read more aboutmicroformatsandhCard.
- 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?showDeactivatedUsers=true
```

- Any user who does not have the 'Can Use' Confluence global permission won't appear in the People directory (for example, Jira Service Management customers who can view KB articles, but do not have a Confluence license).
- 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`showExternallyDeletedUsers`parameter to the URL. For example:

```
http://my.confluence.com/dopeopledirectorysearch.action?showExternallyDeletedUsers=true
```

- The Confluence administrator canhide the people directory. If it is hidden, you will not see the**People Directory**option.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

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**yourprofile 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.

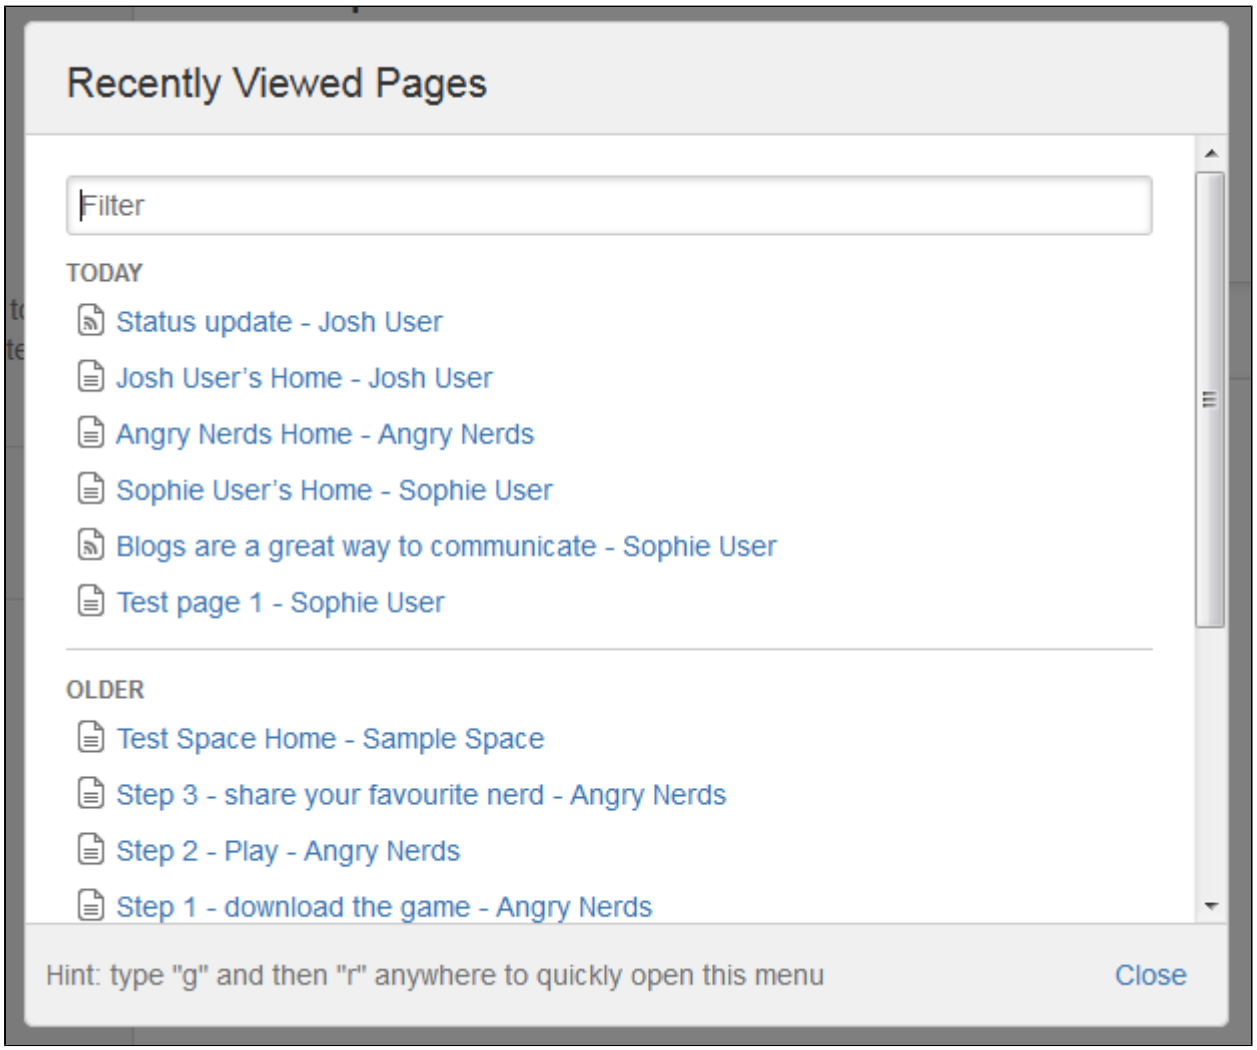

Ranking of Search Results

When you perform a search in Confluence, it is likely that there will be many pages or other content items that match your search terms.

Confluence will rank these matching items by evaluating their relevance. This should mean that the items most relevant to you will appear at the top of the search results list so that you can quickly select the item you need.

Below is an overview of the method Confluence uses to determine the relevance of the items returned by the search, that is, to rank the search results.

On this page:

- Summary of the ranking method
- Weighting by content type
- Weighting based on field type
- Weighting based on recency

Related pages:

- Searching Confluence
- Confluence Search Fields
- Confluence User's Guide

Summary of the ranking method

When displaying the search results, Confluence applies a weighting to each of the content items returned.

To come up with this single weighting, Confluence first combines three separate weightings based on the following factors:

1. The **content type** of the matching item this includes user profiles, pages, blog posts, images and other attachments, etc. More details below
2. The **field type** in which the matching term was found this includes title, name, body content, labels, etc. More details below
3. The **recency** of the matching item that is, when it was created or last modified; this could be today, yesterday, up to 1 week ago, up to 1 year ago, or over 1 year ago. More details below

The item with the heaviest final weighting will appear at the top of the list of search results. All of the other content items will appear below in descending order of weighting.

Weighting by content type

Content type	Weighting
User profile	1.5
Space	1.5
Page	1.5
Blog	1.3
Attachment (files, videos and images)	1
Comment	1
Mail	1

On Confluence, you're most likely to be searching for knowledge articles, work done by a particular team or person, or ways to collaborate. Our ranking logic prioritises content types such as user profiles, spaces, and pages that are most suitable for these tasks.

Simple example

If your search result returns 3 items a page, a blog, and a comment and they are the same in every other way, then they will be ranked in the order shown above.

Weighting based on field type

Field type	Weighting
Title	2
Content	1
Unstemmed title	1
Label	0

Search results that match the title field are twice as important and weighted twice as highly as matches in the body content.

Simple example

If you search for a user's name, the search results will rank the person's user profile above a page that only contains their name in the content. This is because the profile contains the name in the title field. This example assumes the results are the same in every other way.

Weighting based on recency

Last activity	Weighting
Today	2.01-2.05
Yesterday	1.92-2.01
Up to 1 week ago	1.52-1.92
Up to 1 month ago	1.46-1.52
Up to 3 months ago	1.36-1.46
Up to 6 months ago	1.25-1.36
Up to 1 year ago	1.11-1.25
Beyond a year	1-1.11

Recency is based on when an item was created or last modified, whichever happened more recently. Search gives a higher weighting to recently updated items because it assumes this content is likely to be more relevant than idle or older content.

When a content item has not been modified in over a year, we say that it is in a state of "decay". In this state, we assume it is less relevant and it is de-prioritised in the search results in favour of content modified within the last year, and even more so within the last week.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Simple example

- If two documents match in all other ways, then the newer one with be shown first.
- If the two documents being compared are both older than a year then their relative age does not matter.

Confluence uses the Apache Lucene search engine library. Lucene's score calculation has a number of additional terms, not mentioned in the above example. We have simplified the above explanation of search ranking for purposes of illustration. If you are interested, you can see more information in the Lucene documentation.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

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.

On this page:

- Levels of permission
  - Global permissions
  - Space permissions
  - Page restrictions
- How do permissions and restrictions interact?
- Check who can view a page

Related pages:

- Confluence Groups
- Global Permissions Overview
- Space Permissions Overview
- Page Restrictions
- Configuring Confluence Security

Levels of permission

There are three levels of permissions in Confluence:

- global permissions
- space permissions, and
- page restrictions.

Global permissions

Global permissionsare site-wide permissions, and are assigned by a Confluence administrator or system administrator.

Global permissions cover things like whether a user can log in or create a space. They don't really interact with space permissions or page restrictions.

For full details, check out theGlobal Permissions Overview.

Space permissions

Every space has itsown independent set of permissions, managed by the space administrators, which determine the access settings for different users and groups.

They can be used togrant or revoke permissionto view, add, edit, and delete content within that space, and can be applied to groups, users, and even to anonymous users (people who aren't logged in) if need be.

For full details, check out the Space Permissions Overview.

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 group 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 for viewing or editing by default**, but you canrestrict either viewing or editingto certain users or groups if you need to. Page restrictions can be applied to published or unpublished pages and blog posts (drafts).

Don't forget, every page in Confluence lives within a space, and space permissions allow thespace adminto 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.

For full details, check outPage Restrictions.

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 is 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.

Check who can view a page

To check who can view a page, go to••• >**People who can view**. This will show a list of people, including administrators, who have space permission to see the space, and are not prevented by page restrictions from seeing the page. SeeCheck who can view a page for more information.

The diagram below shows the points at which someone could be prevented from viewing a page.

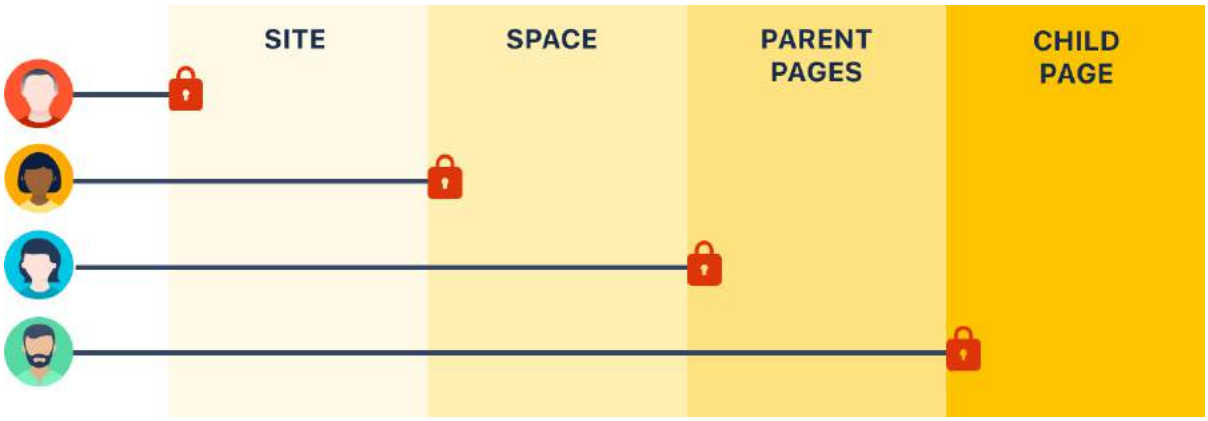

- **Site** - they don't have permission to log in to Confluence.
- **Space** - they don't have permission to view the space.
- **Parent page** -the current page has a view restriction preventing them from viewing.
- **Child page** -the current page is inheriting a view restriction from another page higher up in the page hierarchy, preventing them from viewing.

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 will be visible, but they'll get a "page not found" message. The space key is not revealed in the link URL.
- 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.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Confluence Groups

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

If you're a space admin, you can assign a set ofspace 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.

On this page:

- Default Confluence groups
- Anonymous users
- Unlicensed users from linked applications

Related Pages:

- Space Permissions Overview
- Assign Space Permissions
- Page Restrictions
- Confluence Groups for Administrators

Default Confluence groups

There are two default groups in every Confluence instance and, beyond that, Confluence administrators are free to set up and edit groups in any way they see fit.

The two default groups in Confluence are:

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

Overlapping permissions

Space permissions are additive. If a user is granted permissions as an individual or as a member of one or more groups, Confluence will combine these permissions together. This is sometimes known as their effective permissions.

Sasha is a member of the `confluence-users` group and the `developers` group. The `confluence-users` group has 'export' permission, but does not have 'restrict' permission. The `developers` group has 'restrict' permission but does not have 'export' permission.

By being a member of these two groups, Sasha can restrict *and* export content. The permissions do not conflict, they combine to determine what Sasha is allowed to do in this space.

Anonymous users

People who don't log in when they access Confluence are known as 'anonymous' users. By default, anonymous users don't have access to view or change any content in your Confluence site, but Confluence admins can assign permissions to anonymous users if it's required.

Unlicensed users from linked applications

If you're using Confluence as a knowledge base for Jira Service Management, your Jira Service Managementadministratorcan choose to allow all active users andcustomers (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 way 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.

Check who can view a page

Confluence is open by default,however because of the layers of space permissions and page restrictions that can be applied, it isn't always obvious who can see your page.

If you want to share a page with someone in a different team, for example, it's useful to know whether they have adequate permissions to see it *before* you share.

On this page:

- Check who can view a page
  - Why can these people view?
  - What should I do if someone can't see my page?
  - How do I make my page completley private?
- Disable the People who can view option

Related pages:

- Permissions and restrictions
- Global Permissions Overview
- Space Permissions Overview

Check who can view a page

To check who can view a page:

1. Go to ••• >**People who can view**.
2. A list of all the people who can view the page will appear.
3. Start typing a name to filter the list.

People who can view is available in both Confluence Server and Confluence Data Center.

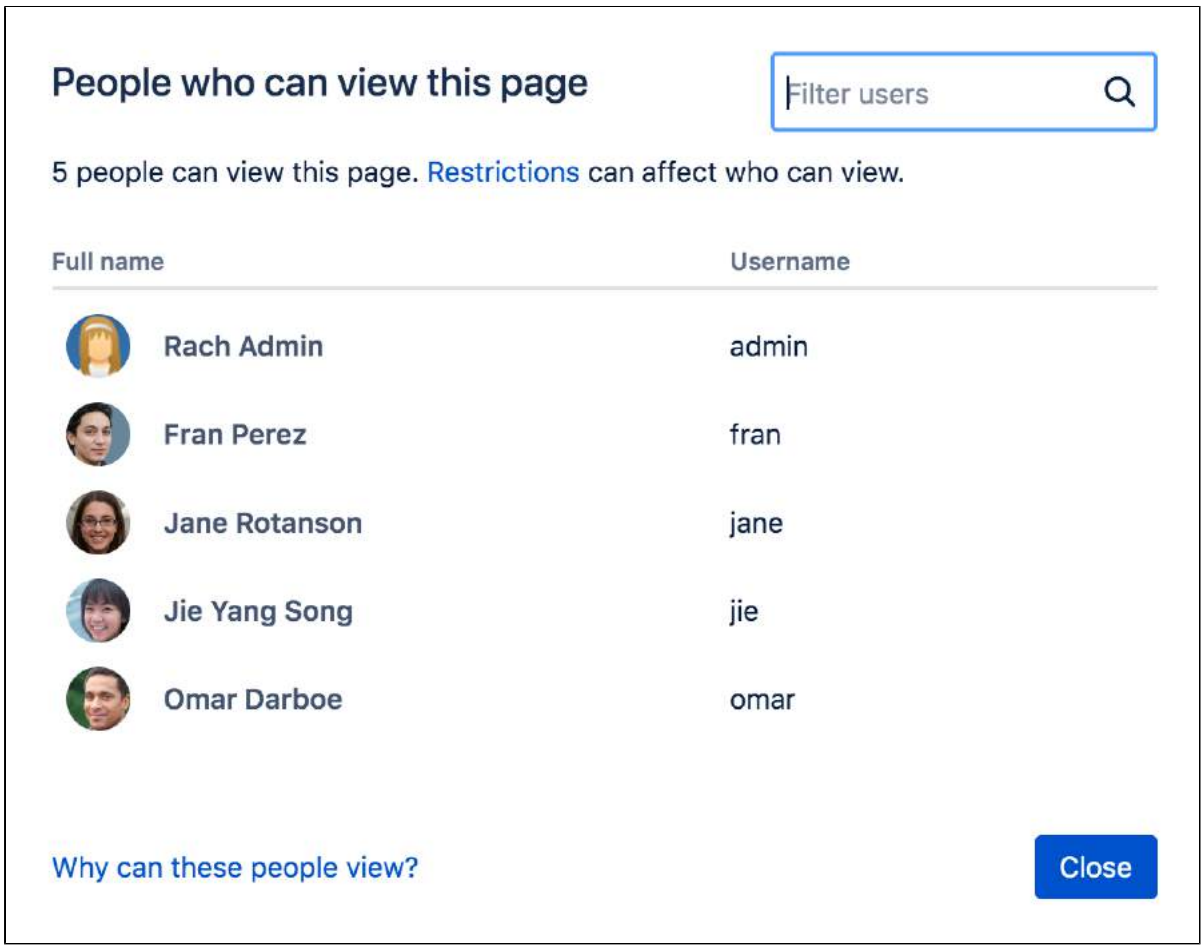

Screenshot: the People who can view dialog showing a list of 5 people.

Why can these people view?

The list includes everyone who:

- is allowed to view the space, via space permissions, and
- is not prevented from viewing the page by page restrictions.

The list also includes members of the confluence-administrators super group, who can always see all spaces and pages. This means that even if you restrict a page to yourself, it's possible that at least one other person can see it.

What should I do if someone can't see my page?

If you want to share your page with a person, and they're not listed, you'll need to work out what is preventing them from viewing the page.

If you restricted the page, add them to the page restrictions, and then check the People who can view list again.

If they still aren't listed, it's likely they don't have permission to see the space. You'll need to contact a space admin to help with this.Go to**Space Tools**>**Overview**for a list of space admins.

How do I make my page completley private?

The short answer is, you can't. Confluence is designed to be open. It's for sharing work with your team. While you can restrict a page to yourself,it's important to note that:

- Space administrators always have the ability to remove page restrictions, even from pages they can't see.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

- People with Confluence Administrator and System Administrator global permissions can't see your pages by default, but they can grant themselves space administrator permission to the space.
- Members of the confluence-administrators super group can see all spaces and pages. This group grants the highest possible permission in Confluence.Some organisations use this group heavily, while others don't use this group at all.

Disable the People who can view option

People with Confluence Administrator or System Administrator permissions can disable the **People who can view** option, if they don't want people to be able to see a list of user's names.

To disable the People who can view option:

1. Go to <base-url>/admin/plugins/gatekeeper-plugin/global/configuration.action
2. Deselect**Show 'People who can view' option**and save the change.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Inspect permissions

Inspect permissions is an advanced permissions feature, and is only available in Confluence Data Center.

Confluence is open by default, but because of the layers of group, individual and anonymous permissions that can be applied, it can be challenging to find out exactly who can do what.

Inspect permissions helps you:

- troubleshoot permissions problems
- audit who can do what in your site
- apply the permissions granted to a user or group in one space to multiple spaces.

It reveals a person's**effective permissions**, combining everything we know about their permissions in a way that can be easily interpreted.

On this page:

- Troubleshoot permissions problems
  - Inspect permissions via space administration
  - Inspect permissions via the Global administration
  - Inspect permissions for a group
  - Permissions explanations
- Audit permissions
- Bulk apply permissions
  - Bulk apply permissions for a group
  - Bulk apply permissions for a user
- Troubleshooting and known issues
  - General cache problems
  - Export options limited in Internet Explorer 11
  - Inspect permissions for a group only shows direct permissions
  - Excluding spaces with no permissions can take a long time in large sites

Related pages:

- Permissions and restrictions
- Global Permissions Overview
- Space Permissions Overview

Troubleshoot permissions problems

Often you will need to find out why someone can or can't do something in Confluence. By inspecting permissions you can drill right down to the root of the problem.

For example, someone reports to you that a teammate can see a space they shouldn't be able see. By inspecting permissions you can work out exactly what group memberships, for example, might be contributing.

Inspect permissions via space administration

You need space admin permissions for the space you want to troubleshoot.

To find out why someone can or can't view this space:

1. Go to the space and choose**Space tools**>**Permissions**from the bottom of the sidebar
2. Go to the **Inspect permissions** tab
3. Enter the person's name or username.
4. Leave the **Page** field blank (unless you need to investigate a specific page in this space).
5. Choose **Show**.

A table showing the person's effective permissions in this space will appear. Click one of the icons to go to the detail view, and find out exactly why they do or don't have that permission.

If you choose to specify a particular page, the permissions explanations will also include information about any page restrictions. The icons will represent just that page, not the user's permissions for the entire space.

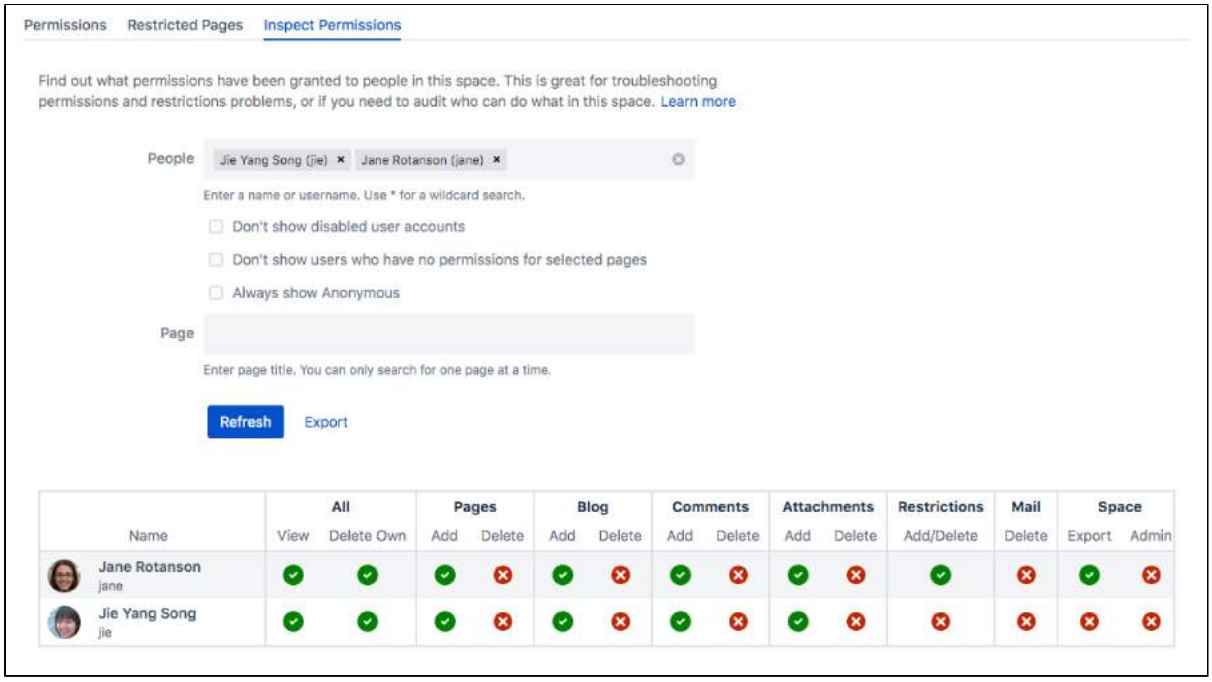

Screenshot: Inspect permissions tab in Space Tools showing permissions for two users.

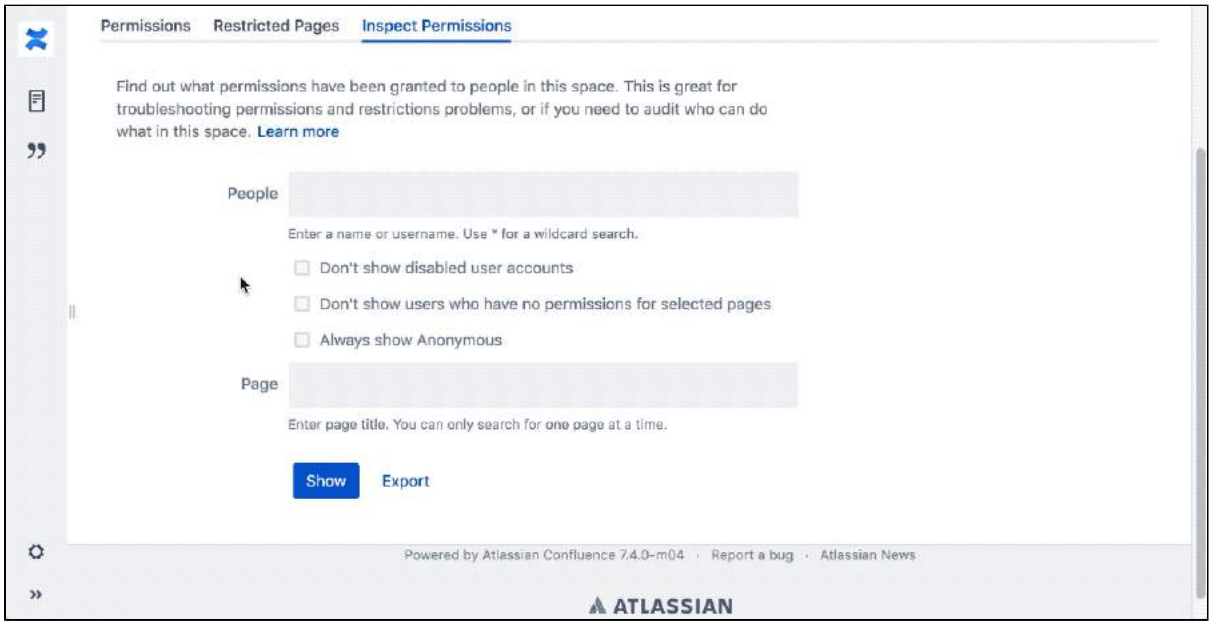

Animation: Inspect permissions tab in Space tools showing permissions explanations for a user

Inspect permissions via the Global administration

You need Confluence Administrator or System Administrator global permission to do this. You don't need to have permission to view the space itself.

To find out why someone can or can't view this space:

1. Go to > **General Configuration**>**Inspect permissions**.
2. Enter the person's name or username.
3. Enter the spaces you want to view.
4. Choose **Show**.

A table showing the users and spaces you searched for will appear. Click the link to see the detailed view, then click the icons to find out exactly why they do or don't have that permission.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

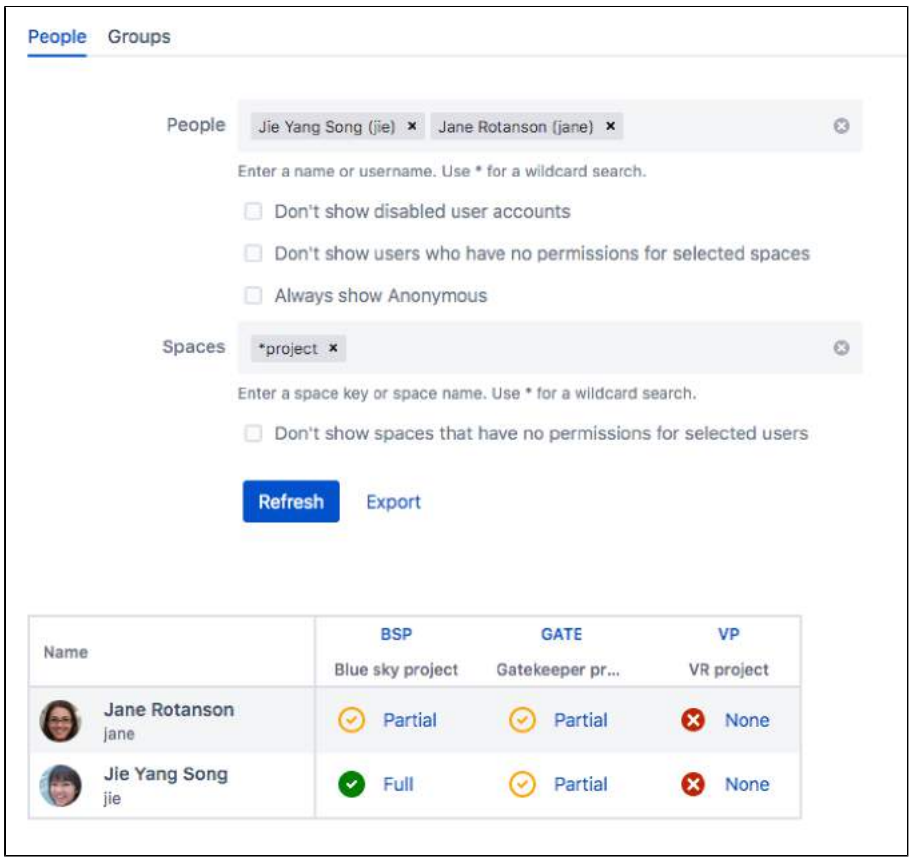

Screenshot: Inspect permissions in Global administration showing permissions for two users and three spaces.

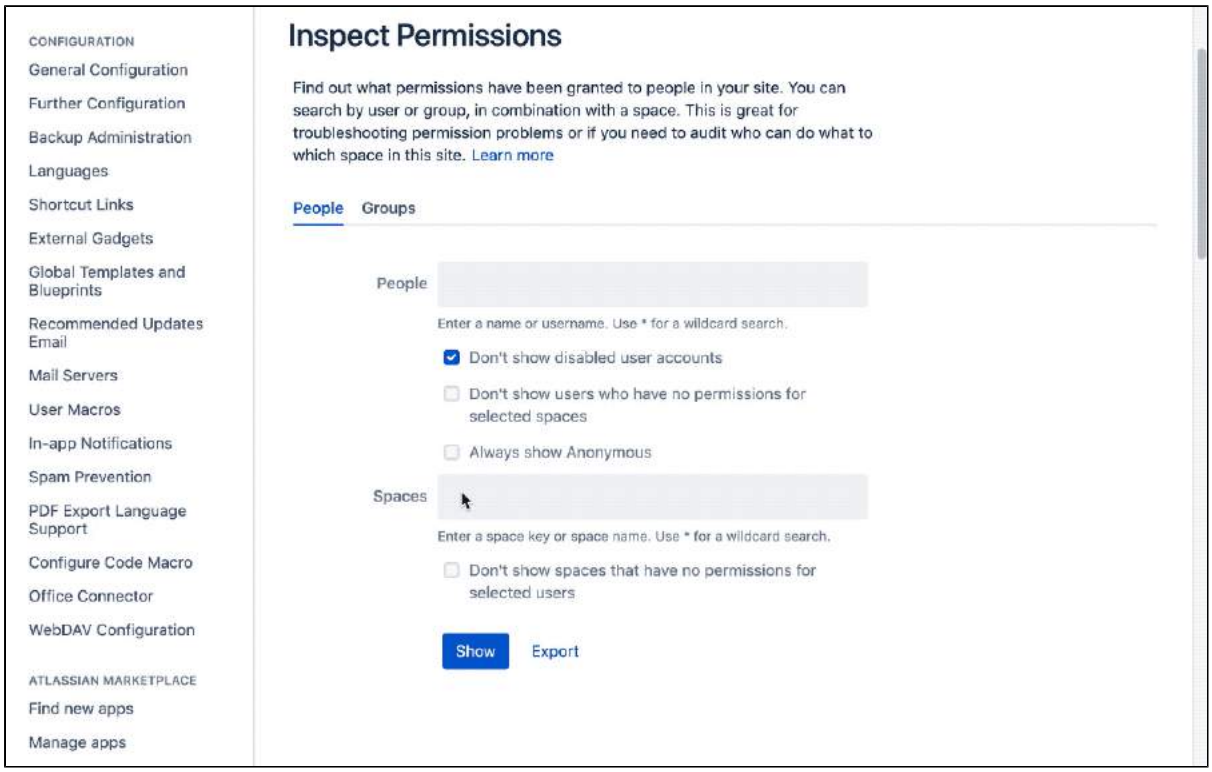

Animation:Inspect permissions in Global administration showing searching for all spaces containing the word "project" and then viewing permissions explanations for a user.

Inspect permissions for a group

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

You can also inspect permissions for a specific group. This can only be done via the global administration.

To inspect permissions for a group:

1. Go to > **General Configuration**>**Inspect permissions**.
2. Choose the **Groups** tab
3. Enter the group name.
4. Enter the spaces you want to view.
5. Choose**Show**.

A table showing the groups and spaces you searched for will appear. Click the link to see the detailed view.

When inspecting permissions for a group, be aware:

- We don't indicate when a group does not have the Can Use global permission, as we do for users.
- We don't show effective permissions for the group, as we do for users. We only show permissions directly granted to that group (not granted via membership of a parent group. This is only an issue if your external user directory has nested groups).

Using wildcards in your search

When searching for users, groups, or spaces, you can use * as a wildcard. For example if you wanted to search for all spaces that contain the word project in their title, type `*project`and select **Search for *project** from the suggestions dropdown.

Permissions explanations

The detail view shows the effective permissions for a single user or group in a space. Click each icon to see a detailed explanation, as shown here.

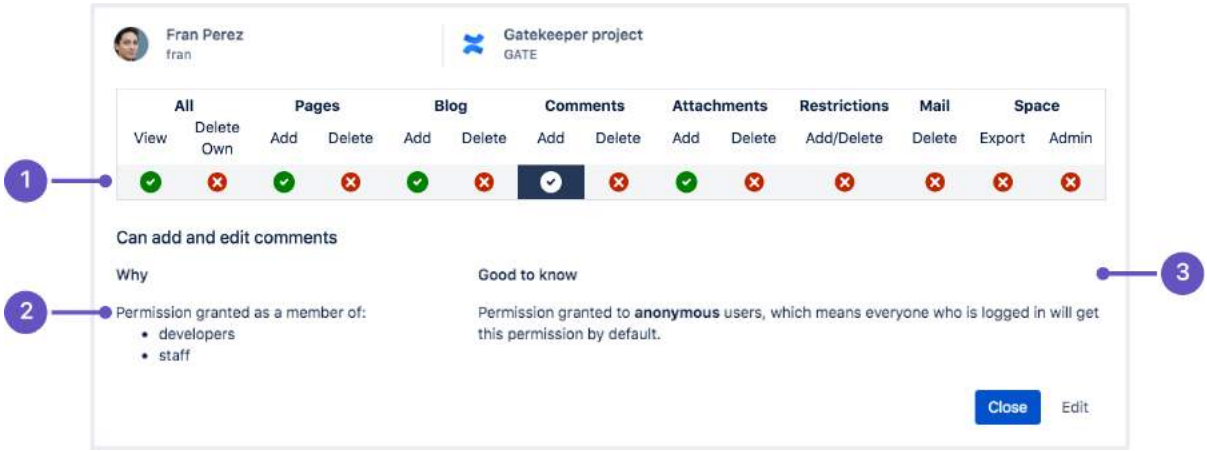

1. **Icons** - icons indicate whether the user or group can or can't do this action.
2. **Explanation** - explains why the user can or can't do this action.
3. **Good to know information** - provides additional information that may become relevant, for example that space administrators can grant themselves permissions they don't currently have.

The purpose of these explanations is toprovide a simple reason why someone can or can't do something in a space.The messages are designed to be short, and present the most relevant information first.

The table below contains a more detailed explanation of every message, including the conditions that trigger the message.

Explanations about users

Explanation	Message appears when...

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Permission granted as an individual	The user is listed in the space permissions for that space.
Permission granted as a member of:	The user is a member of a group, that is listed in the space permissions for that space. If the user is a member of multiple groups that are listed in the space permissions, we will list all of them.
Permission granted as an individual and as a member of:	The user is listed in the space permissions for that space, and is also a member of a group that is listed in the space permissions for that space.
Permission granted to**anonymous**users, which means everyone will get this permission by default, including people who are not logged in.	The permission is granted to anonymous users in this space, and your site is public (you have granted anonymous users the Can use global permission). Logged in users can't have fewer permissions than anonymous users.
Permission granted to**anonymous**users, which meanseveryone who is logged in will get this permission by default.	The permission is granted to anonymous users in this space, but your site is not public (anonymous users do not have the Can use global permission). Logged in users can't have fewer permissions than anonymous users. This is sometimes used as a shortcut way to provide 'everyone' with space permission, without making the site itself public.
No permission granted as an individual or as a member of a group.	The user isn't listed in the space permissions for that space, they are not a member of a group that is listed in the space permissions for that space, and anonymous has not been granted any permissions.
This person doesn't have the**Can use** global permission, so they can't log in to Confluence.	This user exists in the user directory, but doesn't have a Confluence license seat. They are not a member of confluence-users or another group that has the 'Can use' global permission.
This person is a**Confluence administrator**socould grant themselves this permission.	The user, or a group they're user is a member of, has Confluence Administrator global permission. This means they can recover permissions for a space they don't have permission to see, and then change the permissions for that space. Unlike members of the confluence-administrators super group, they can't see the space by default.
This person is a**space admin**, so could grant themselves this permission.	The user has space admin permissions in this space. This means they can modify permissions for this space, and could grant themselves any permissions they don't currently have.
This person is a**space admin**, so can edit restrictions. They can also remove all restrictions from pages they dont have permission to edit or view.	The user has space admin permissions in the space. This means they can always change page restrictions (even if they don't have the Restrict permission), and can access a list of all restricted pages in the space, and remove all restrictions from these pages.
This person has**Delete own**permission so can delete their own pages, blog posts, comments, and attachments.	The user can delete pages, blog posts, comments, and attached files that they have created. They can't delete pages, blog posts, comments, and attached files created by other users unless they also have Delete permission in the space. For example the user can delete a page they created, but they cannot delete a page their team mate created unless they also have the Delete Page space permission.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

To add and delete restrictions this person also needs the **Add page** permission.	The user has Restrict permission, but does not have the Add page permission. Applying a page restriction is considered editing the page, so both permissions are required.
This person is a member of the **confluence-administrators super group**. This means they can view all pages, including restricted pages. While they can't edit existing pages, they can add, delete, comment, restore page history, and administer the space.	The user is a member of the confluence-administrators group. This is a default group that has significant privileges in Confluence, beyond that provided by the Confluence Administrator or System Administrator global permission.
This person can't log in because their **account is disabled**.	This user exists in the user directory, but their account has been disabled. They don't have a Confluence license seat. This is usually because the person has left the organisation.
Restrictions on <page title> prevent this person from viewing the page.	A page restriction has been applied to the page. The user, or a group they're a member of, are not listed in the page restrictions dialog, so they have 'no access'. This message only appears when you inspect permissions for a specific page in a space.
Restrictions on <page title> allow this person to view, but prevent them from editing the page.	A page restriction has been applied to the page. In the page restrictions dialog, either everyone can view or the user or group they're a member of can view, but only specific users or groups can edit. This message only appears when you inspect permissions for a specific page in a space.
Restrictions on <page title> allow this person to view and edit the page.	A page restriction has been applied to the page. The user, or a group they're a member of, are listed in the page restrictions dialog, and they have 'View and edit' access.
Restrictions on <page title> allow this person to view the page.	A page restriction has been applied to the page. The user can view the page, either because the page restriction allows everyone to view, and only some people to edit, orthe user, or a group they're a member of, are listed in the page restrictions dialog, and they have "View only" or 'View and edit' access. This message only appears when you inspect permissions for a specific page in a space, and click on a permission that doesn't require permission to edit the page, for example View, or Add comment.

Explanations about groups

These messages appear when you select a group as the entity to inspect. You need to be a Confluence Administrator to inspect permissions for groups.

Permission explanation	Message appears when...
Permission granted to all members of this group.	The group is listed in the space permissions for that space.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Permission granted to**anonymous** users, which means everyone will get this permission by default, including people who are not logged in.	The permission is granted to anonymous users in this space, and your site is public (you have granted anonymous users the Can use global permission). Logged in users can't have fewer permissions than anonymous users.
Permission granted to**anonymous** users, which meanseveryone who is logged in will get this permission by default.	The permission is granted to anonymous users in this space, but your site is not public (anonymous users do not have the Can use global permission). Logged in users can't have fewer permissions than anonymous users. This is sometimes used as a shortcut way to provide 'everyone' with space permission, without making the site itself public.
No permission granted to this group.	The group isn't listed in the space permissions for that space, and anonymous has not been granted any permissions.
This group doesn't have the**Can use**global permission, so people in this group may not be able to log in to Confluence.	This group exists but does not have the 'Can use' global permission. This is very common. Often one group, such as confluence-users is used to grant a Confluence license seat, and additional groups used only to manage space permissions. This is only an issue if a user is not a member of another group that grant them Can use permission.
Members of this group are**space admins**, so could grant themselves this permission.	The group has space admin permissions in this space. This means members of this group can modify permissions for this space, and could grant themselves or this group any permissions they don't currently have.
Members of this group are**space admins**, socan edit restrictions. They can also remove all restrictions from pages they dont have permission to edit or view.	The group has space admin permissions in the space. This means members of this group can always change page restrictions (even if they don't have the Restrict permission), and can access a list of all restricted pages in the space, and remove all restrictions from these pages.
Members of this group have**Delete own**permission so candelete their own pages, blog posts, comments, and attachments.	Members of this group can delete pages, blog posts, comments, and attached files that they have created. They can't delete pages, blog posts, comments, and attached files created by other users unless they also have Delete permission in the space. For example the user can delete a page they created, but they cannot delete a page their team mate created unless they also have the Delete Page space permission.
confluence-administrators is a**super group**. Members of this groupcanview all pages, including restricted pages. While they can't edit existing pages, they can add, delete, comment, restore page history, and administer the space.	This is a default group that has significant privileges in Confluence, beyond that provided by the Confluence Administrator or System Administrator global permission.
To add and delete restrictions people in this group also need the **Add page**permission.	The group has Restrict permission, but does not have the Add page permission. Applying a page restriction is considered editing the page, so both permissions are required.
Delete ownis only available to members of this group who have logged in.	This message appears when Delete Own permission is granted to a group that doesn't have Can use global permission. It is just a reminder that people must be able to log in to delete their own content.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

This permission information may be incomplete because this group is a member of one or more parent groups. Permissions granted to parent groups are not shown here.	This message appears when the group is nested, that is, it's a member of another group. This hierarchy of groups comes from your external user directory. We don't show any permissions that a group is inheriting from a parent group. You'll need to inspect these parent groups seperately.

Explanations about anonymous

These messages appear when you select Anonymous as the entity to inspect. Anonymous in this instance means people who have not logged in to Confluence.

Permission explanation	Message appears when...
Permission granted to anonymous users.	Permission granted to the anonymous entity listed in the space permissions for that space.
No permission granted to anonymous users	No permission granted to the anonymous entity listed in the space permissions for that space.
Anonymous users don't have**Can use**global permission. People must log in to use Confluence. All logged in users will inherit this permission by default.	The permission is granted to anonymous users in this space, but your site is not public (anonymous users do not have the Can use global permission). Logged in users can't have fewer permissions than anonymous users. This is sometimes used as a shortcut way to provide 'everyone' with space permission, without making the site itself public.
Although**Delete own**permission can be granted to anonymous users, it has no effect.	The Delete Own permission is assigned to anonymous users. Because you need to be logged in for us to know who you are, and what you have created, Delete Own is never available to anonymous users, even when granted.
Restrictions on <page title> prevent anonymous users from viewing the page.	A page restriction has been applied to the page. Anonymous users have 'no access'. This message only appears when you inspect permissions for a specific page in a space.
Restrictions on<page title>allow anonymous users to view, but prevent them from editing the page.	A page restriction has been applied to the page. In the page restrictions dialog, everyone can view, but only specific users or groups can edit. This message only appears when you inspect permissions for a specific page in a space.
This permission can't be granted to**anonymous**users.	This is a reminder that some permissions, such as Space Admin, and Restrict are never available to anonymous users.
Anonymous access is enabled globally.	This is a reminder that your site is public. Because you have granted anonymous users the Can use global permission, people do not need to log in to access Confluence.

Audit permissions

If you need to regularly check who can do what in your site, for example for compliance or regulatory reasons, you can inspect permissions to conduct an audit.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

To export permissions information for all users and all spaces in your site:

1. Go to > **General Configuration**>**Inspect permissions**.
2. Choose your search criteria - you can search for particular users, groups, or spaces, include disabled accounts, or exclude users who have no permissions.
3. Choose **Show**. This will give you an idea of how large your query is.
4. Choose **Export**.
5. Keep the default separator. Comma separated is great for opening in most spreadsheet applications.
6. Choose how you want spaces to be listed. This depends on how much information you want to include. You can include just the space key, or the title and description too.
7. Choose **Export**.

The CSV file will be immediatley downloaded in your browser. This can take a few minutes, depending on the size of your query.

This file can be *extremely* large in sites with many users and spaces. You could use the wildcard search feature to limit the number of users to be included in each export.

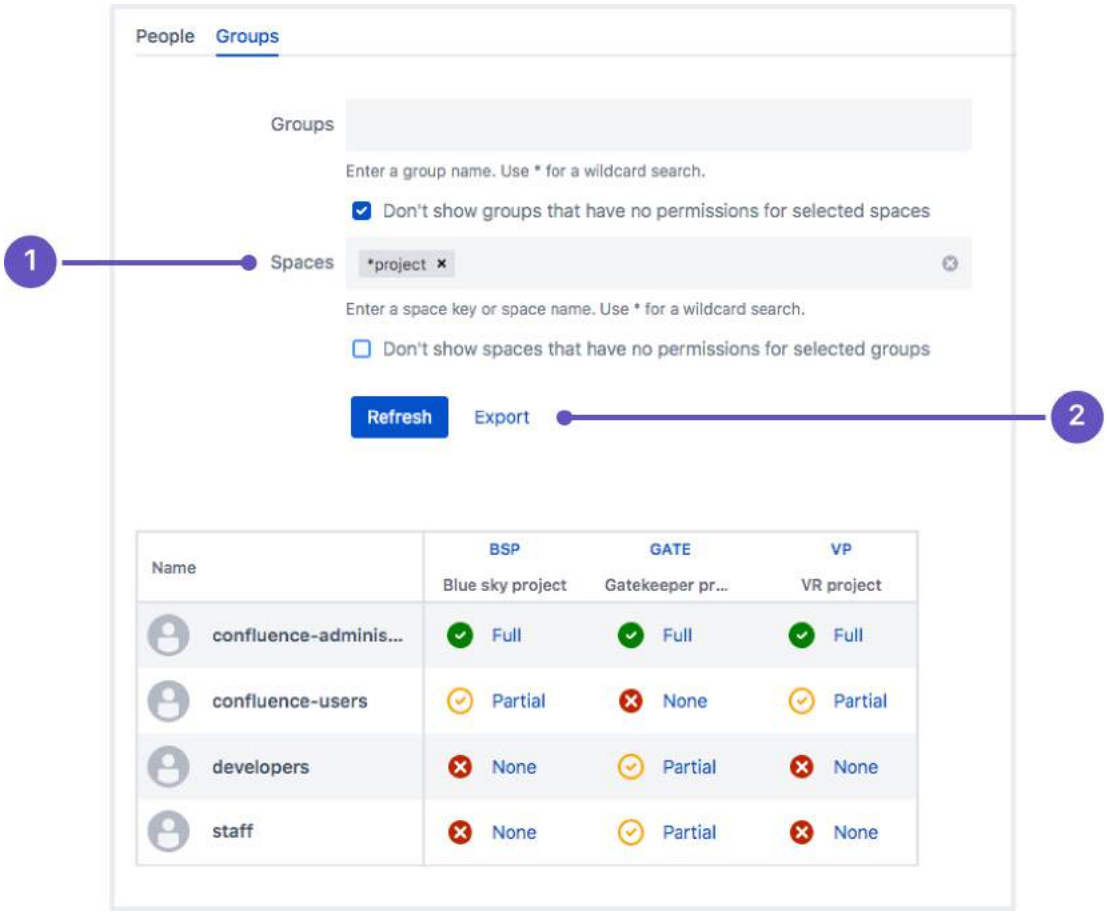

1. **Wildcard search** - use a wildcard to narrow your search.
2. **Export** - export the results to CSV for auditing or further analysis.

This example shows the output for one user, and three spaces.

- A row will be created for each of the 14 space permissions that can be granted.
- A column will be created for each space in your query. These are identified by space key, but you can choose to include the space name and description in the export if you require it.
- T and F indicates whether the user has this permission (true) or they do not have this permission (false)

Username	Display Name	Permission	SPACE1	SPACE2	SPACE3

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

user1	Sample User 1	view-space	T	T	F
user1	Sample User 1	remove-own	T	T	F
user1	Sample User 1	page-add	T	T	F
user1	Sample User 1	page-remove	T	F	F
user1	Sample User 1	blog-add	T	F	F
user1	Sample User 1	blog-remove	T	F	F
user1	Sample User 1	comment-add	T	T	F
user1	Sample User 1	comment-remove	T	F	F
user1	Sample User 1	attachment-add	T	T	F
user1	Sample User 1	attachment-remove	T	F	F
user1	Sample User 1	mail-remove	T	F	F
user1	Sample User 1	page-restrict	T	F	F
user1	Sample User 1	space-export	T	F	F
user1	Sample User 1	space-admin	T	F	F

In this example Sample User 1:

- Can do everything, including administer Space1
- Can view, add, and delete their own content in Space2
- Can't view Space 3 at all.

Bulk apply permissions

Bulk applying permissions is useful when:

- You create a new group, and want to give that group permissions to a number of existing spaces.
- You need to grant someone permissionsas an individual for a number of existing spaces.
- You have just created several new spaces, and want to use permissions from an existing space as a template.

We recommend using groups as an efficient way to manage permissions in your site. When someone new starts on your team, we would recommend making them a member of appropriate groups, over using the bulk add permission options to grant them permissions to all their spaces as an individual.

Bulk apply permissions for a group

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

To give a group the same permissions in multiple spaces:

1. Go to > **General Configuration**>**Inspect permissions**.
2. Go to the **Groups** tab.
3. Enter the name of the group.
4. Enter the name of the space that you want to use as a starting point.
5. Choose **Show**, then click to see the detailed view.
6. Choose **Edit**.
7. Make sure the correct permissions are selected. These are the permissions you'll be copying.
8. Enter the names of the spaces you want to copy these same permissions to. You can add as many spaces as you like.
9. Choose **Apply**.
10. Check the confirmation dialog to make sure all spaces were successfully updated.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

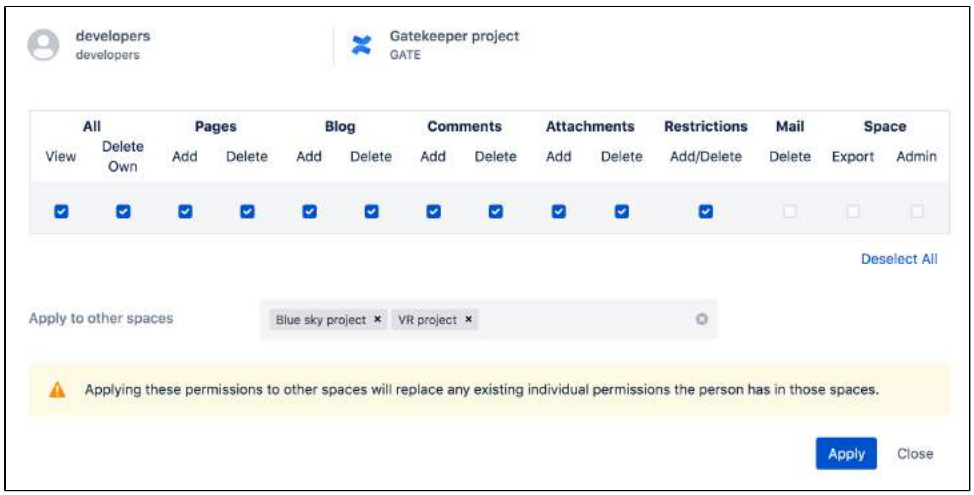

Screenshot: detail view of permissions for a group and space, with several spaces listed in the 'Apply to other spaces' field.

Bulk apply permissions for a user

Before you do this, we recommend you inspect permissions to find out what permissions the user already has for the spaces you plan to update, paying particular attention to how the permissions are granted (individually, or via a group). In many cases the best approach is to change the user's group membership, or the space permissions granted to a group, rather than bulk applying changes to the individual.

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

To give a user permissions in multiple spaces:

1. Go to > **General Configuration**>**Inspect permissions**.
2. Enter the name of the user.
3. Enter the name of the space that you want to use as a starting point.
4. Choose**Show**, then click to see the detailed view.
5. The user's effective permissions are shown. These permissions are a combination of any individual and group permissions. Choose**Edit**.
6. Make sure the correct permissions are selected. These are the permissions you'll be copying.
7. Enter the names of the spaces you want to copy these same permissions to. You can add as many spaces as you like.
8. Choose**Apply**.
9. Check the confirmation dialog to make sure all spaces were successfully updated.

As a general rule, we recommend managing permissions using groups. If you do decide to bulk apply permissions for a user, there are some things to be aware of:

- Permissions can only be granted to the user as an individual using this method.
- Any permissions the user has as a member of a group will be unchanged. For example if they are a member of a group that has Export permission, and you bulk apply permissions that do not have Export permission, they will still have Export permission. Changing their individual permissions can't override their group permissions.
- The checkboxes, when you click Edit, reflect the user's current effective permissions - that is the combination of all the permissions they already have as an individual or member of a group. When you click Apply, you'll apply these exact permissions as an individual. You may actually be doubling up on permissions the user already has, as a member of a group.

Troubleshooting and known issues

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

General cache problems

Confluence caches permissions information, which helps to make sure results are returned quickly when you check who can view a page, or inspect permissions. We continually update the cache as people and permissions change. However, we know of two scenarios where the cache is not correctly updated - when you import a site, and when you add, disable, or change the order of your user directories.

If this happens, the best way to force Confluence to rebuild the cache is to disable the **Inspect permissions - gatekeeper** plugin, then re-enable it.Alternatively, you can restart Confluence, as the cache is built on startup.

Export options limited in Internet Explorer 11

There's a known issue with the export dialog in Internet Explorer 11. The dialog is known to intermittently freeze if you select either of the dropdown menus. As a workaround, use the default values for the **Separator** and **List spaces by** fields, or use another browser to complete the export.

Inspect permissions for a group only shows direct permissions

If your external directory has nested groups (a group is a member of another group), and you inspect permissions for a group, you'll only see permissions granted directly to that group (not effectively granted by being a member of a parent group). If you search for a user, we'll always show the effective permissions, including those granted by parent groups.

Excluding spaces with no permissions can take a long time in large sites

In the global Inspect Permissions screen, if you select"Don't show spaces that have no permissions for selected users" and don't specify any other filters (such as specifying users, groups, or spaces), the query can take several minutes to return any results. This is particularly true in sites with a very large number of users and spaces.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Permissions best practices

There are several different strategies you can use for managing permissions in your site. The larger your site grows, the more important it is to make sure that your permissions strategy can scale with your organisation.

Granting permission to a space on an individual by individual basis may work well for small teams, but rapidly becomes unwieldy when your user base grows to thousands of people.

On this page, we provide our recommendations for the best ways to manage common permissions scenarios.Most of the advice boils down to:

- Keep Confluence as open as possible, it's designed to be open by default.
- Use groups over individual permissions wherever possible, to avoid headaches in the future.

On this page:

- Give people access
  - I want everyone in my organisation to be able to log into Confluence
  - I want everyone in my organisation to be able to view a space
  - I want to give people in my team access to our space
  - I want to give my team access to all our project spaces
  - I want all the spaces in my site to have the same permissions
  - I want to give external people access to my space
- Lock things down
  - I want to check what a person can access in Confluence
  - I need to prevent someone from accessing Confluence
  - I need to prevent specific people from viewing a space
  - I want to prevent people from seeing my work in progress
  - I want to prevent people seeing part of a space
  - I want share one page but keep the rest of the space private
- Delegate administration tasks
  - I want to delegate space administration to a specific group of people
  - I want to control who can create spaces
- The big questions
  - What permissions should I give people?
  - What should I do when someone leaves my team?
  - What should I do when someone leaves my organisation?

Related pages:

- Permissions and restrictions
- Global Permissions Overview
- Space Permissions Overview
- Inspect permissions

Give people access

I want everyone in my organisation to be able to log into Confluence

The best way to achieve this is to make everyone a member of a group that has permission to log in to Confluence, such as the default **confluence-users** group.

See Adding or Removing Users in Groups for information on how to add people to groups.

When new people join your organisation, add them to this group to grant permission to use Confluence.

If you don't want to use an existing group, you can create a new one. The process is much the same, but you will need to explicitly grant this group global permission to use Confluence.

To change the global permissions for a group:

1. Go to > **General Configuration** > **Global permissions**.
2. Choose **Edit Permissions**.
3. Enter your new group name in the **Grant browse permission to** field.
4. Make sure the **Can use** checkbox is selected.
5. Save your changes.

I want everyone in my organisation to be able to view a space

The best way to do this is to grant space permissions to a group that all users are a member of, such as the default **confluence-users** group.

If your site is **not** public (anonymous users do not have the 'Can Use' global permission, everyone must log in to use Confluence), you can also use the anonymous permission as an 'everyone' shortcut. This is useful if your groups setup is complex, and there isn't a single group that everyone is a member of. If you plan to make your site public in future however, it's best to avoid this workaround.

I want to give people in my team access to our space

Think about whether your space really needs to be private. If not, you can grant permission to a group that all users are a member of, such as confluence-users.

If it does need to be private, and your team is only going to be using this one space, it might be appropriate to grant permissions as individuals. That way you don't need to ask a Confluence Administrator to add people to groups. SeeAssign Space Permissions.

However, if your team needs access to multiple spaces, using a group is definitely the way to go, as it will save you a lot of time in future when people join or leave your team.SeeAdding or Removing Users in Groups .

I want to give my team access to all our project spaces

The best way to do this is to create a group, and grant that group permissions in each project space. When people join or leave your team, you only need to change the group membership, you don't need to edit the space permissions for multiple spaces.SeeAdding or Removing Users in Groupsfor more information.

It might be more work to set up now, but it will help you in the long term.

I want all the spaces in my site to have the same permissions

First, you should change the default space permissions, so that when a new space is created, it automatically gets your desired permissions.

To change the default space permissions:

1. Go to > **General Configuration** > **Space permissions**.
2. Choose **Edit permissions**.
3. Add groups, and grant permissions in the usual way then **save** your changes.

Any new spaces created will get these permissions by default.

For existing spaces, it is a little more laborious. You'll need to go to the space permissions screen in each space, and set your desired permissions manually.

If you have Confluence Data Center you can slightly speed up this process by applying the permissions from one space to multiple spaces. This is done on a group by group or user by user basis. There is no way to copy an entire set of permissions from one space to another. SeeInspect permissions.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

I want to give external people access to my space

If you don't want to make your site public, but you need to give people outside your company, such as a customer or contractor, access to your site, you will need to create user accounts for these people. We recommend creating a group specifically for these people, so that it is easy to remove their access later when it is no longer needed.

Your company is hosting a huge event, and you want to be able to collaborate with staff at Super Events, an external events company, in Confluence, rather than relying on email.

1. Create a group called `super-events-staff`
2. Grant this group the **Can use** global permission.
3. Create a user account for each Super Events person you'll be working with.
4. Add these users to the `super-events-staff` group only. Remove them from `confluence-users`, and any other default group to make sure they don't see any spaces they shouldn't. They should only be a member of `super-events-staff`.
5. Create a space for the event, and set the permissions for your internal staff.
6. Give the `super-events-staff` group permission to add pages, comments, and attachments only.
7. Send the space URL and usernames to your contacts at Super Events, and start collaborating.

By confining these users to a single group, they won't see any spaces, or other content that they don't have permission to see, such as Confluence Questions. However, they will be able to see things like the people directory.

Lock things down

I want to check what a person can access in Confluence

In Confluence Server there is no easy way to do this. You will need to find out which groups the user is a member of, and then manually check the permissions for each space.

In Confluence Data Center you can Inspect Permissions to find out what a user can view.

I need to prevent someone from accessing Confluence

The best way to do this is to disable the person's user account. They will not be able to log in.SeeDelete or Disable Usersto find out how to do this.

I need to prevent specific people from viewing a space

If you have Confluence Data Center, Inspect permissions for the person and the space, to find out exactly how they are being granted permission. If you have Confluence Server, youwill need to see what groups have permission, then manually check if the person is a member of that group.

If their permission was granted as an individual, simply go to the space permissions and change their permissions. If their permission was granted via a group, you'll need to decide whether to remove them from the group, or to change the whole group's permissions.

I want to prevent people from seeing my work in progress

First,check who can view your page. It may be that only you, or your team can see the page due to space permissions.

If you do need to lock it down further,the simplest way to do this is restrict the page, so that only you, or your team, can view it.SeePage Restrictionsto find out how to do this.

Once you're ready to share your work, remove the restrictions. A notification won't be sent when you remove the restrictions. Notifications are only sent at the point you publish the page (this means that if you restrict a page to yourself, and publish it, anyone who is watching the space for new pages won't ever get a notification).

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

I want to prevent people seeing part of a space

The simplest way to do this is to use Page Restrictions. This is particularly useful when the pages are a work in progress, and will eventually be opened up for more people to view at a later date.

In this example, a user wants to keep all the pages relating to a sensitive new project private, until the information can be shared with the whole organisation.

Here's what they would do:

1. Create a page called "Secret project" and restrict it to just the people working on the project.
2. Create or move any pages relating to the project to be a child of "Secret project". The view restriction will be inherited.

This approach is not foolproof. It requires people to remember to create future sensitive pages under the restricted parent page, and to avoid moving pages to a parent that is unrestricted. If the content is sensitive, and will always be restricted, consider moving it to a different space, and use space permissions to control who can see the pages.

I want share one page but keep the rest of the space private

This can be tricky, and introduces complexity that may be a problem later, because you are forcing Confluence to work in a way that is opposite to the way it is intended to be used.

Essentially you would need to organise your page hierarchy so that all pages are restricted, except the one you want to share. You would then change the space permissions to open up the space.You can then check who can view a page to make sure you've achieved the desired result.

In this example, a user wants to keep the work in their personal space private, but make their "What I'm working on" page available for their manager and team to view.

Here's what they would do:

1. Create a page called "Private work" and restrict this page to themselves. Only they can see this page.
2. Move all the pages in the space that should remain private to be a child of "Private work".
3. Create a page called "Open work". Move the "What I'm working on" page to be a child of this page.
4. Change the space permissions so that their manager and team can view the space.

Important things to note:

- This approach is not foolproof. It requires the user to remember to create future sensitive pages under the restricted parent page, and to avoid moving pages to a parent that is unrestricted.
- Any blog posts or other non-page content created in the space would be visible, because the page restrictions only apply to pages that are a child of "Private work".

Delegate administration tasks

I want to delegate space administration to a specific group of people

The best way to do this is to create a specific space administrators group. The benefit of using a group is that you can easily add and remove members, without needing to touch the space permissions for the spaces themselves.

If you need to create a sensitive space, that these people shouldn't be able to view or administer, simply edit the space permissions for that space, and remove the group's permissions.

- Create a group, and give it a recognisable name like space-administrators.
- Add the people you want to be space admins as members of this group.
- Grant this group space admin permissions in the default space permissions, so all new spaces will be created with this permission.
- Go to every existing space and manually grant thisgroup space admin permissions. If you have Confluence Data Center you can use Inspect permissions to speed up this step.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

I want to control who can create spaces

You can set which groups or individuals can create spaces in Global Permissions.

If you choose to limit who can create spaces, we recommend granting this permission to a group of champions, who can handle requests, create the spaces, and work with stakeholders to set up their space permissions in the most appropriate way for your organisation. These people don't need to be Confluence Administrators, they just need the Create Space global permission.

The big questions

What permissions should I give people?

This is going to depend on your organisation, and the type of work you are doing in Confluence. If collaboration is your goal, we recommend giving people full Add, Delete, and Restrict permissions, and granting Space Admin permissions to a handful of people, who can act as champions in the space, to perform tasks like creating templates, or customising the view.

In some industries you may need to prevent people from deleting or restricting content, for auditing or compliance reasons. If this is the case for your organisation, consider updating the default space permissions so that all new spaces are created with your ideal permissions.

The main use-case for your Confluence site also has an impact on how you will structure your permissions. Find out about using confluence forTechnical Documentation,Knowledge Base articles,your Intranet, or Softw are Teams.

What should I do when someone leaves my team?

If most spaces in your site are open, chances are you don't need to do anything. However it's good practice to change the person's group memberships to match their new role. This might happen automatically, via your external user directory, or you may need to search for the user, and change their group memberships manually.

Once you've changed their group memberships, if you're a Confluence administrator and you have Confluence Data Center you canInspect permissions to check what spaces the person still has access to, then edit their permissions for each space on the fly, to remove any individual permissions.

What should I do when someone leaves my organisation?

If someone leaves your organisation, usually you would disable their user account, either in Confluence, or in your external user directory.

You may want to tidy up any individual permissions they've been granted (just to reduce the number of people listed in your space permissions screens), but unfortunately there's no easy way to do this. If you're a Confluence administrator, and you have Confluence Data Center,you canInspect permissionsto check what spaces the person still has access to, then edit their permissions for each space on the fly, to remove any individual permissions.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Team Calendars

Team Calendars for Confluence is now part of Confluence Data Center
To get access to the features described on this page upgrade to Confluence Data Center 7.11 or later. Cant upgrade yet? Depending on your current Data Center version, you can access these features by installing the latest version of the app (at no cost). See our FAQ for all the details

Team Calendars provides one place for fast-moving teams working in Confluence toconnect their schedules of leave, Jira projects and events, or whatever use case your team has.

- Team Calendars Quick Tour
- Create, Add, and Edit Calendars
- Restrict a Calendar
- Watch a Calendar
- Share Calendars
- Export Team Calendars Content to Other Calendars
- Delete or Remove a Calendar

Create, organize, and find events easily. Team Calendars offers great tools for planning and managing your events, whether they're travel, conferences, birthdays, or JIRA issues and sprints.

- Add Events
- Add Jira Events
- Event Types
- Custom Event Types
- Reminders
- Embed Calendars on Confluence Pages

Calendar subscriptions provide a simple and effective way to stay up to date with holidays, conferences, and more.

You can subscribe to team calendars from many of your your favourite third-party calendar apps:

- Subscribe to Team Calendars from Microsoft Outlook
- Subscribe to Team Calendars from Outlook on the web
- Subscribe to Team Calendars from Apple Calendar
- Subscribe to Team Calendars from Apple iOS Calendar
- Subscribe to Team Calendars from Google Calendar
- Subscribe to Team Calendars from Android Calendar
- Subscribe to Team Calendars from Thunderbird

Alternatively, subscribe to a third-party calendar in Team Calendars:

- Subscribe to Outlook Calendar from Team Calendars
- Subscribe to Apple Calendar from Team Calendars
- Subscribe to Google Calendar from Team Calendars
- Subscribe to Opsgenie Calendars from Team Calendars
- Subscribe to PagerDuty Schedules from Team Calendars
- Subscribe to Teamup Calendars from Team Calendars

Team Calendars Quick Tour

Team Calendars allows you to create calendars for yourself and your team, and view other calendars from your organization, all in one place. This page provides you a few essential steps to get started.

On this page:

- Open Team Calendars
- Create and add Calendars
- Add events to a calendar
- Choose your view
- Embed calendars on Confluence pages

Related pages:

- Create, Add, and Edit Calendars
- Add Events
- Event Types

Open Team Calendars

There are two places you can view and work with Team Calendars in a Confluence space, or on your 'My Calendars' page.The My Calendars page presents a roll-up of all of the Team Calendars you're currently subscribed to.

- If you're looking for your team's calendar, or you'd like to create a calendar for your team, go to your team's Confluence space and choose**Calendars** in the sidebar.
- If you'd like to create a calendar for yourself, or subscribe to your Google Calendar, choose**Calendars**in the Confluence headerto open your 'My Calendars' page.

The My Calendars page

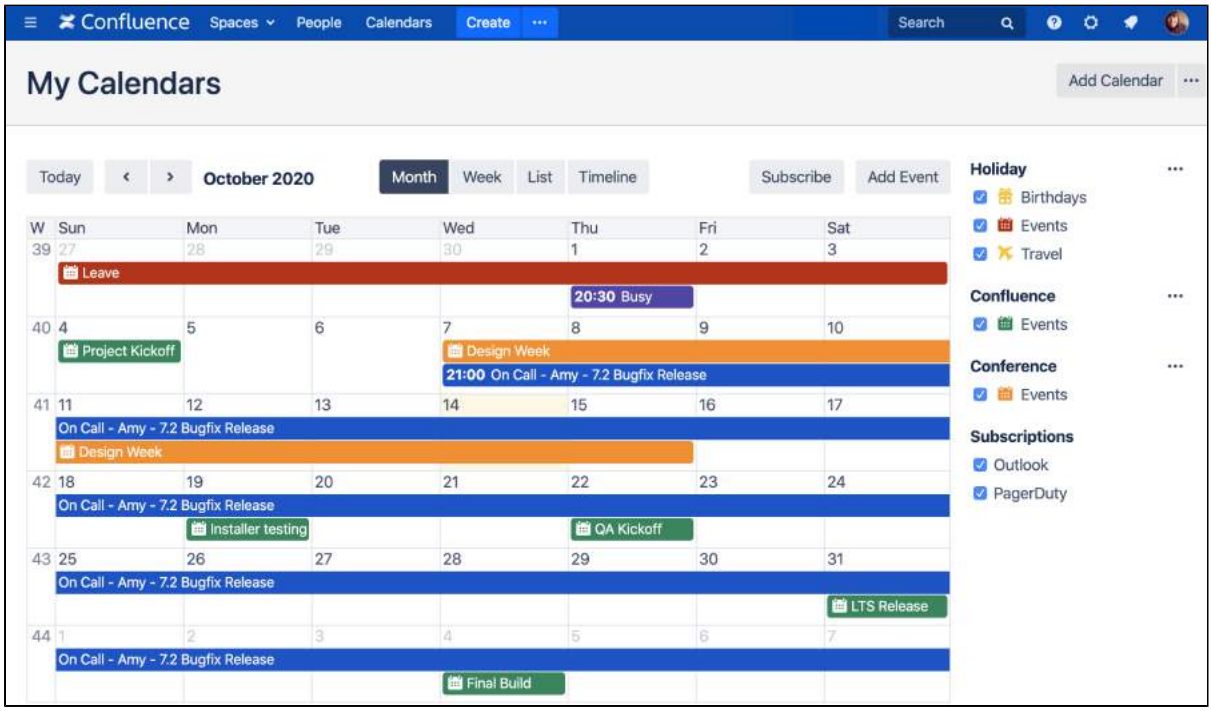

Calendars in a space sidebar

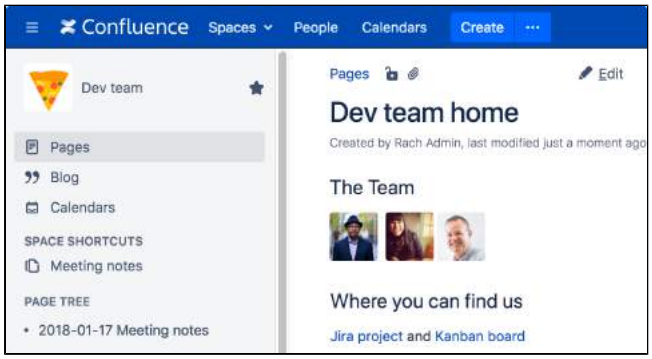

Create and add Calendars

To get started, you'll need to create a calendar or add an existing one. Select**Add Calendar** to quickly create a new calendar, or select••• next toAdd Calendarto add an existing calendar in the following ways:

- **Add Existing Calendar** to subscribe to another calendar in your Confluence instance.
- **Import Calendar** to import an ICS file.
- **Subscribe by URL** to Subscribe to Third-Party Calendars from Team Calendars.

SeeCreate, Add, and Edit Calendarsto learn more.

Add events to a calendar

To add an event to the calendar, you can do any of the following:

- Choose**Add Event**
- Click a date on the calendar in either the month or week view (or click and drag to choose a date range)
- Double-click on the Timeline view

You'll be prompted to enter a title, time and other details. Once you've added the details, choose **OK** to add the event to the calendar. If there are multiple calendars on your My Calendars page, you'll need to select the calendar you're adding the event to.

You can also click and drag on the calendar to select a range of days in the month and week views. Learn more by reading Add Events .

Choose your view

You can view your events in four different ways: month, week, list, and timeline. Team Calendars remembers your last selected view, so you don't need to choose it again next time you come back.

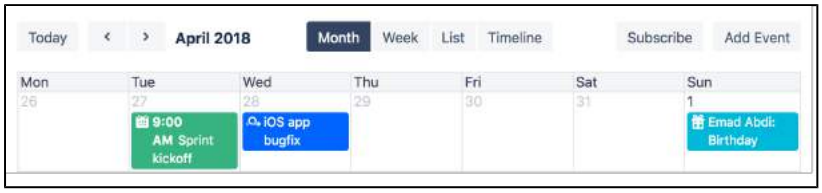

Embed calendars on Confluence pages

You can embed a calendar on any page in Confluence, ensuring you and your teams are always up-to-date. Take a look atEmbed Calendars on Confluence Pagesfor more information.

Team Calendars for Confluence is now part of Confluence Data Center
To get access to the features described on this page upgrade to Confluence Data Center 7.11 or later. Cant upgrade yet? Depending on your current Data Center version, you can access these features by installing the latest version of the app (at no cost). See our FAQ for all the details

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Create, Add, and Edit Calendars

You can create new calendars or add existing calendars in any Confluence space or your My Calendars page. Existing calendars can be other calendars from your Confluence site, or third-party calendars.

On this page:

- Create a calendar
- Add other calendars
- Import a calendar
- Edit a Calendar

Related pages:

- Restrict a Calendar
- Add Events
- Event Types

Create a calendar

To create a brand new calendar:

1. Do either of the following:
   - **Create a calendar in a space** Hit **Create from template** in the Confluence header, then choose **Calendars** in the create dialog (choose the right space too, if you're not already there)
   - **Create a calendar in My Calendars** Choose **Calendars** in the Confluence header to go to your My Calendars page, then hit **Add Calendar** at the top-right of the page
2. Choose **Add New Calendar**
3. Enter a **Name** for your calendar

   > If you're in My Calendars, you'll also need to enter a **Related space** . The calendar will appear in that space once you've created it.

4. Choose **OK**

Add other calendars

To add other calendars from your Confluence instance, or from third-party calendars, select ••• to the right of **Add Calendar**, then pick any of the following:

- **Add Existing Calendar** to subscribe to another calendar in your Confluence instance.
- **Subscribe by URL** to Subscribe to Third-Party Calendars from Team Calendars.

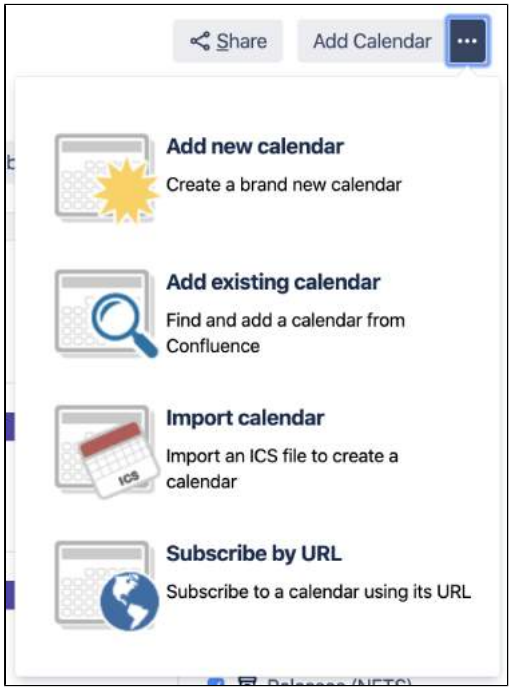

> Did you know you can add a calendar to a Confluence page? See Embed Calendars on Confluence Pages.

Import a calendar

Import a calendar from an an iCalendar (ICS) file. You can import the file to a new calendar, or add the events to an existing team calendar.

To import an ICS file as a new calendar:

1. Select ••• to the right of **Add Calendar**
2. Select **Import Calendar**
3. Enter a name for your new calendar
4. Locate and select the .ics file you saved and follow the prompts to import it.

To import an ICS file into an existing calendar:

1. Choose ••• to the right of the destination calendar and choose **Import iCalendar**
2. Locate and select the .ics file you saved and follow the prompts to import it.

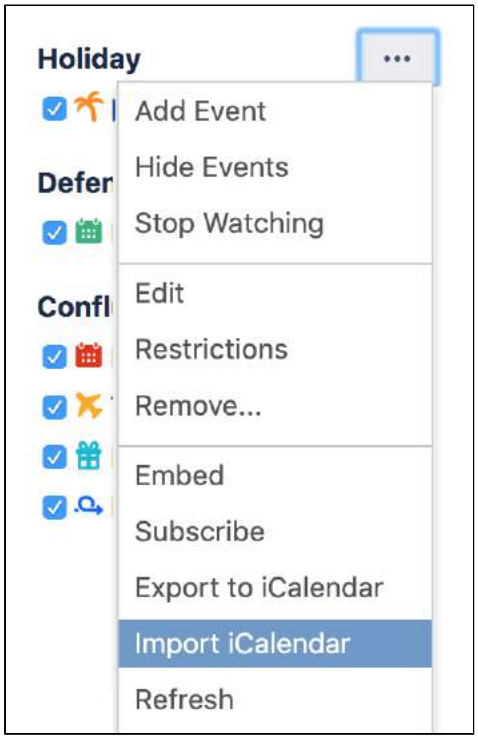

Edit a Calendar

If you want to change the name of a calendar, hide certain event types,or move the calendar to another space, you can edit the calendar provided you have sufficient permissions. You can also add custom event types and hide default event types to tailor the calendar to your needs.

1. Choose **Calendars** in the Confluence header or space sidebar
2. Choose ••• to the right of the calendar name, then choose **Edit**
3. Change the calendar's name or move it to another space in the **General** tab, or hide events in the **Event Types** tab

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

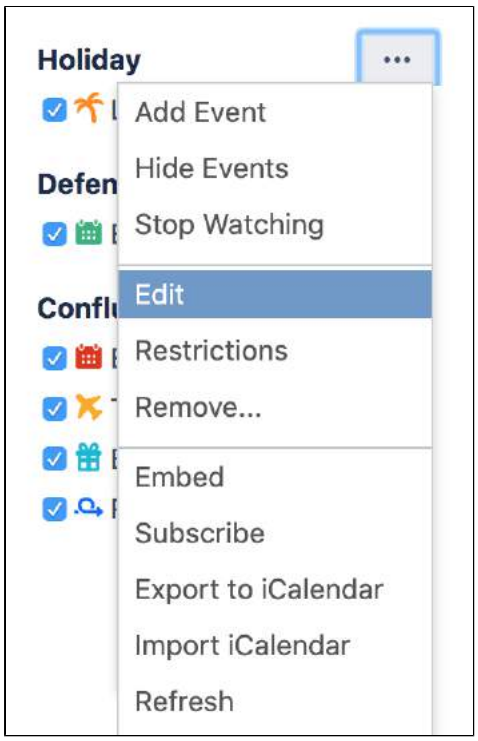

Team Calendars for Confluence is now part of Confluence Data Center
To get access to the features described on this page upgrade to Confluence Data Center 7.11 or later. Cant upgrade yet? Depending on your current Data Center version, you can access these features by installing the latest version of the app (at no cost). See our FAQ for all the details

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Restrict a Calendar

Calendars respect the view permissionsof their related space. When you create a calendar in a space, or add a related space to an existing calendar, only people with permission to view the space will be able to see the calendar.

If a calendar contains sensitive information, or shouldn't be edited by just anyone, you can further limit who can or edit it by adding restrictions. Restrictions in Team Calendars work in much the same way they do in Confluence, you can restrict a calendar to individual peopleorConfluence groups.

Related pages:

- Create, Add, and Edit Calendars
- Add Events

Change who can view a calendar

If no view restrictions are added, everyone who can view the space can view the calendar.

To restrict who can view a calendar:

1. Select **Calendars**in the Confluence header or space sidebar.
2. Select ••• to the right of a calendar's name, then choose**Restrictions**.
3. Select the **Add view restriction**radio button.
4. Search for and select a user, or choose the **Group**button and search for and select a Confluence group. Repeat this for all people and groups you want to be able to see the calendar.
5. Select **OK**.

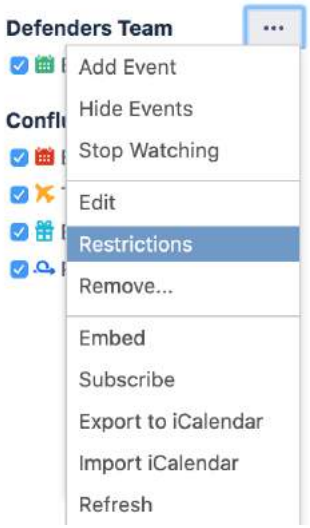

Only the people listed under **These people can view this calendar**in the restrictions dialog, and the calendar's creator, will be able to see the calendar. Remember, people also needview permissionsfor the space associated with this calendar to be able to see the calendar.

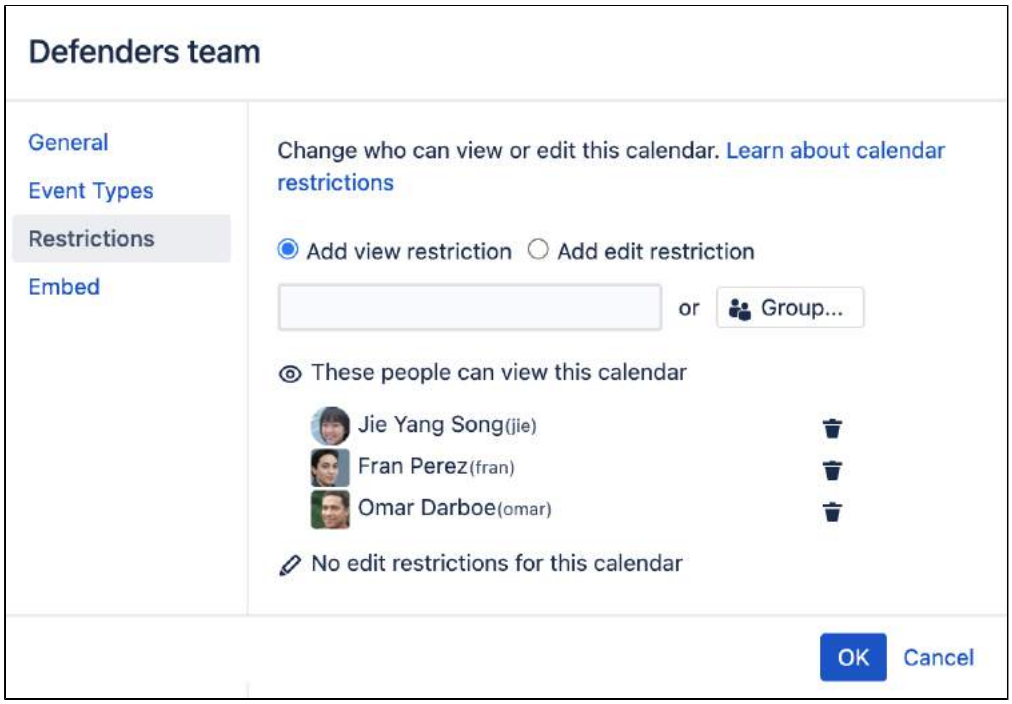

Change who can edit a calendar

If no edit restrictions are added, everyone who can view the calendar can edit the calendar.

To restrict who can edit a calendar:

1. Select **Calendars**in the Confluence header or space sidebar.
2. Select ••• to the right of a calendar's name, then choose**Restrictions**.
3. Select the **Add edit restriction**radio button
4. Search for and select a user, or choose the **Group**button and search for and select a Confluence group. Repeat this for all people and groups you want to be able to edit the calendar.
5. Select **OK**.

Only the people listed under **These people can edit this calendar**in the restrictions dialog, the calendar's creator, and any space administrators (who can see the calendar), will be able to edit the calendar, including adding events, editing existing events, and changing the restrictions.

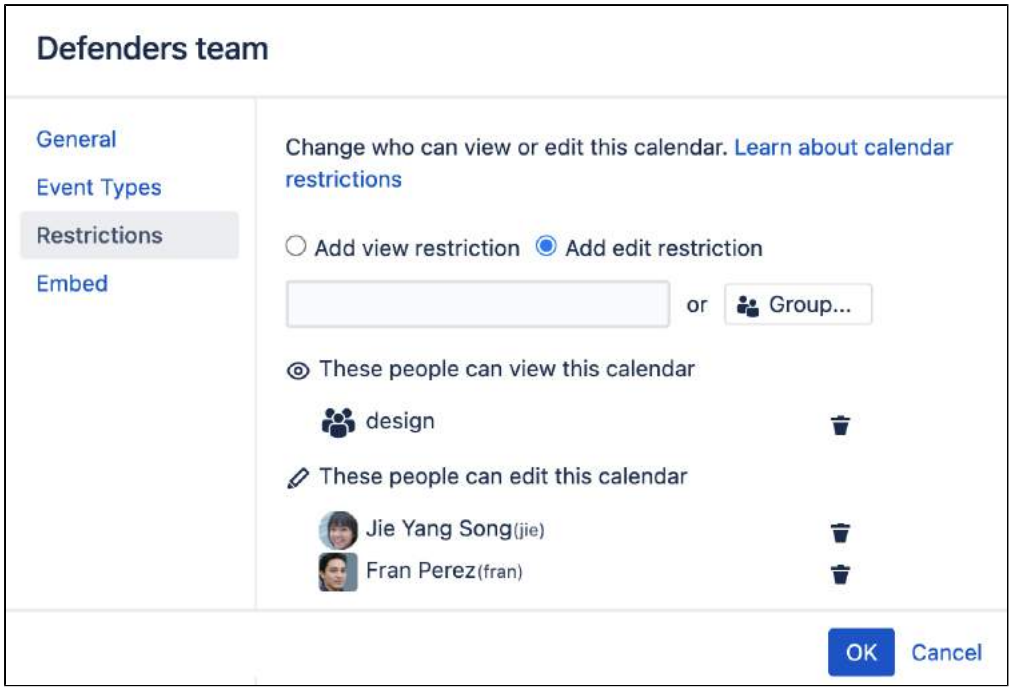

Team Calendars for Confluence is now part of Confluence Data Center
To get access to the features described on this page upgrade to Confluence Data Center 7.11 or later. Cant upgrade yet? Depending on your current Data Center version, you can access these features by installing the latest version of the app (at no cost). See our FAQ for all the details

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Add Events

You can quickly add events to your calendars andcategorize them into different event types for easier viewing and management.

On this page:

- Add an Event
- Event Types
- Repeat options
- Delete an Event

Related pages:

- Add Jira Events
- Event Types
- Custom Event Types

Add an Event

1. Do any of the following:
    - Choose**Add Event**
    - Click a date on the calendar in either the month or week view
    - Double-click on the Timeline view
2. Select the type of event you'd like to add from the**Event Type**drop-down
3. Complete the other details for your event and choose**Add**

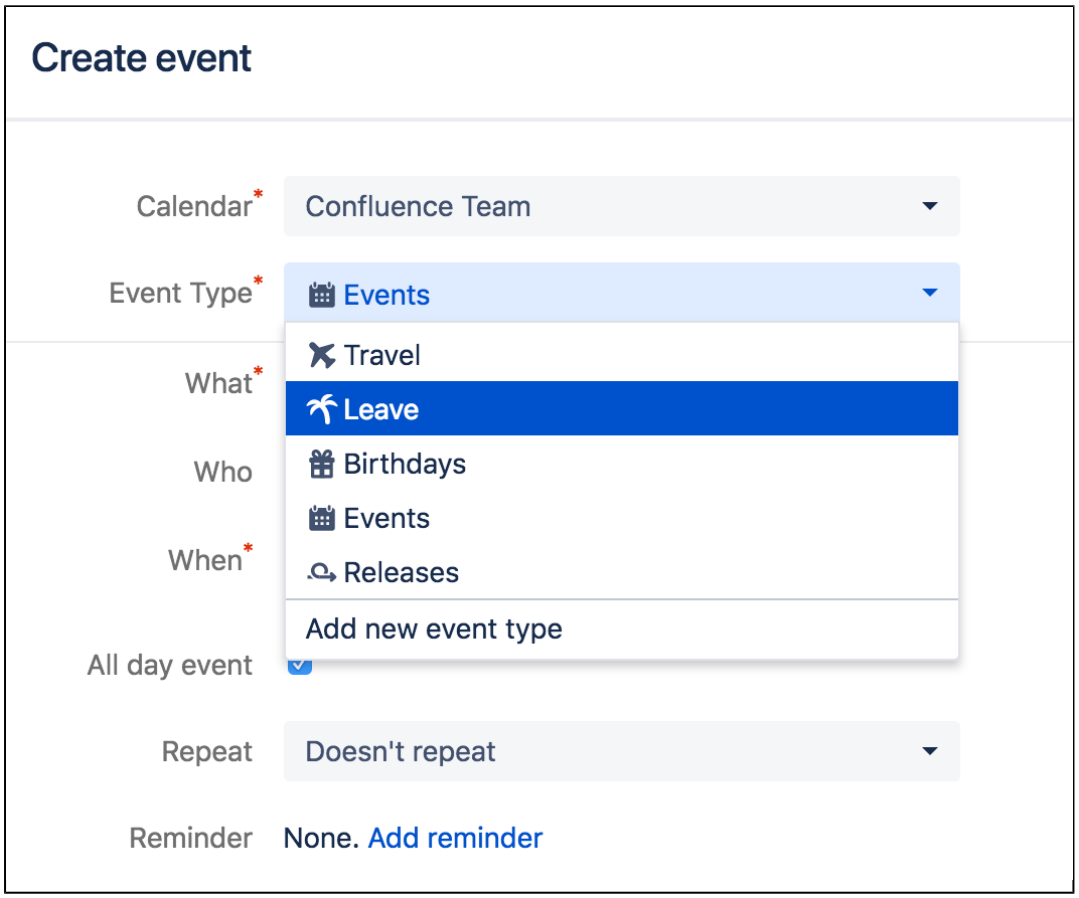

Event Types

Each calendar includes a standard set of event types, which you can use to classify the different events in the calendar.You can alsocreate custom event typesto capture events that don't fit into the standard event types.

Repeat options

You can choose to repeat your event daily, weekly, monthly or yearly.

Delete an Event

To delete an event, select the event and choose**Delete**from the inline dialog.

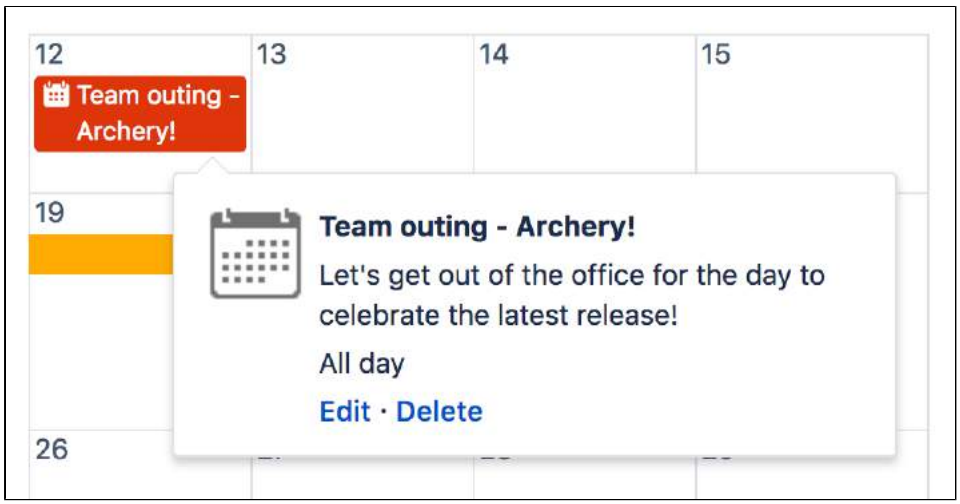

If you don't see the**Delete**option, the reason is usually:

- You don't have edit permissions on the calendar
- The event you're trying to edit is a JIRA event (sprint, version or issue).If this is the case, the change to the JIRA item has to happen in JIRA.

Team Calendars for Confluence is now part of Confluence Data Center
To get access to the features described on this page upgrade to Confluence Data Center 7.11 or later. Cant upgrade yet? Depending on your current Data Center version, you can access these features by installing the latest version of the app (at no cost). See our FAQ for all the details

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Add Jira Events

Jira events enable you to visualize versions, issues or sprint dates fromJira applications. To use this feature, Confluence and Jira need to be connected by an application link.If, after connecting Confluence to Jira you're still not able to add Jira events to a Team Calendar, make sure the**Jira iCalendar Plugin**is enabled in yourJIRA site. This is bundled with Jira, but may have been disabled by an administrator.

On this page:

- Add JIRA events
- Update JIRA events in Team Calendars
- Add JIRA issue dates

Related pages:

- Add Events
- Event Types
- Custom Event Types

Add JIRA events

When you create an event, there are multiple JIRA event types you can choose from. These are:

- **JIRA Issue Dates**Embed a single JIRA issue date or the duration between any two issue dates. For example, embed due dates for JIRA issues, or, if you havecustom fields, embed a scheduled deployment date.
  What if you want to see the duration for a particular issue? You may have fields like 'Deployment Start Date'and 'Deployment End Date'or 'Issue Start'and 'Issue End Date'. Team Calendars allows you to visualize the duration between any two dates.
- **JIRA Software Sprints**If your team is using JIRA Softwareto plan their work, you can also embed your Sprints inside your Team Calendar. This helps you see the duration of your sprint and how your team's availability or other team events could impact the sprint.
- **JIRA Project Releases** JIRA versions help you define a release for your projects. You can embed and visualize your JIRA project versions to see how your releases are progressing.

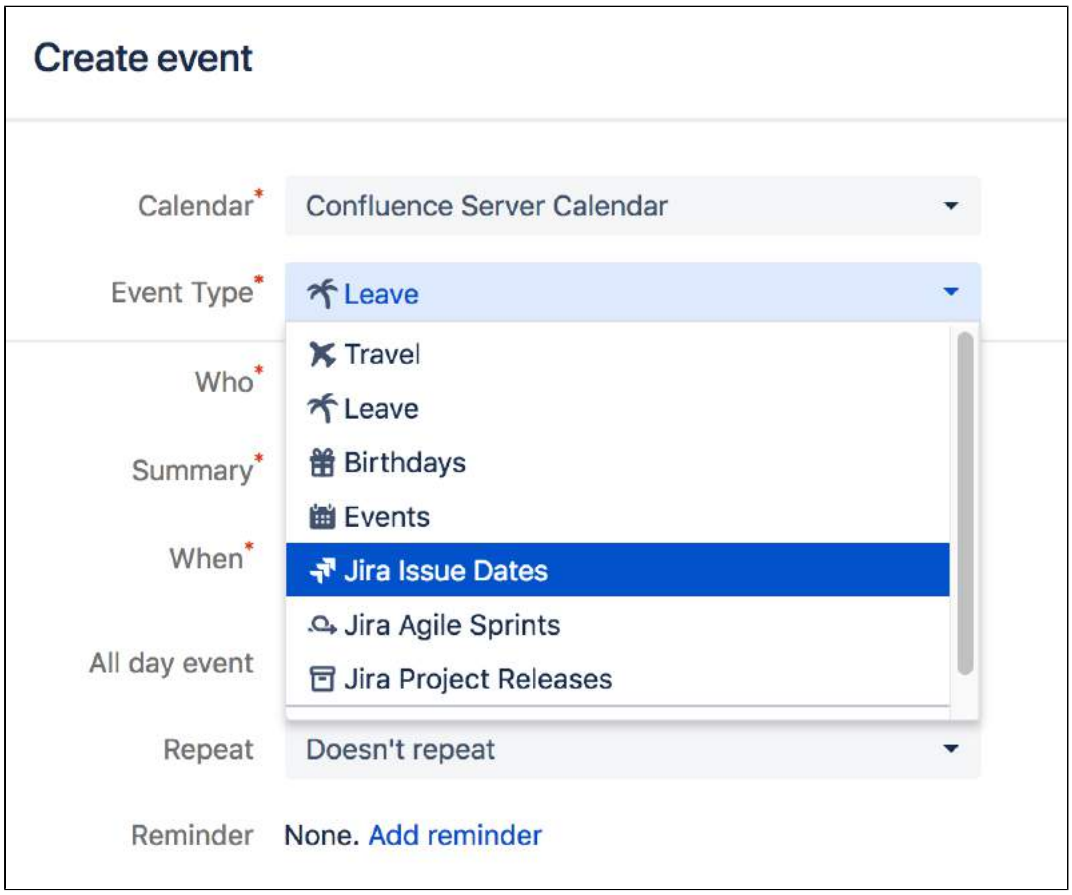

Update JIRA events in Team Calendars

From Team Calendars, you can:

- Update JIRA project version dates
- Update Issue dates

Drag-and-drop the event or update its date, and it will automatically update in your JIRA application. Currently, you can't update your JIRA Software sprint dates.

Add JIRA issue dates

You can embed virtually any issue date in Team Calendarsfrom your JIRA application. Embed JIRA issue dates as:

- A single date
- A duration between two specified dates

When adding JIRA issue dates, you'll be presented with the following display options:

- **Project** A quick and easy way to map out project issues on a calendar
- **Saved filter** For more advanced filtering as well as support for multiple projects.
- **JQL** Power users can useJQLto create a JIRA calendar.

> If you select**Saved filter**or**JQL**, Team Calendars will render all available date fields that are available for all projects selected in your filter. If there are dates that are only available to one project, and not another in the filter, they will not render in the list of possible fields to show.

Upon selecting a Project, Saved Filter or JQL, you'll be presented with all possible dates that can be rendered in Team Calendars. For those of you with Jira Software, Team Calendars can also embed your JIRA Software sprint durations right on your JIRA Calendars.

Once you've entered a JQL query, choose the dates you'd like to display in Team Calendars.

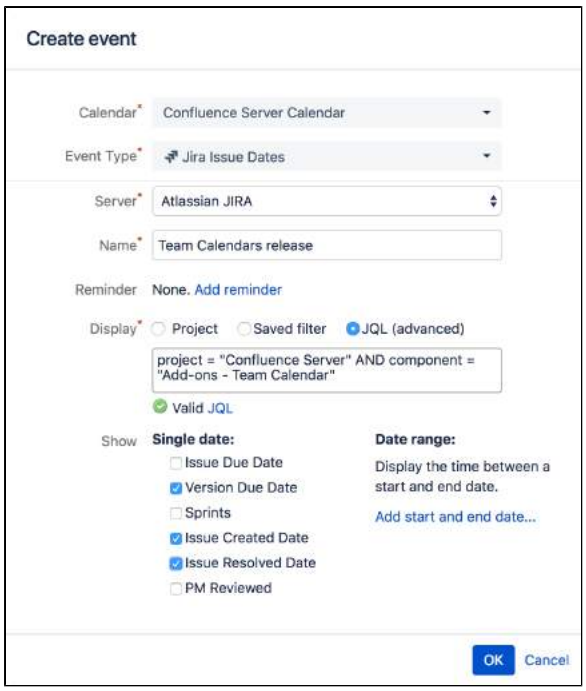

Once the calendar is created, depending on what you've selected, you can see:

- Issue dates
- JIRA project version due dates
- JIRA Software sprint durations

Examples:

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Issue dates

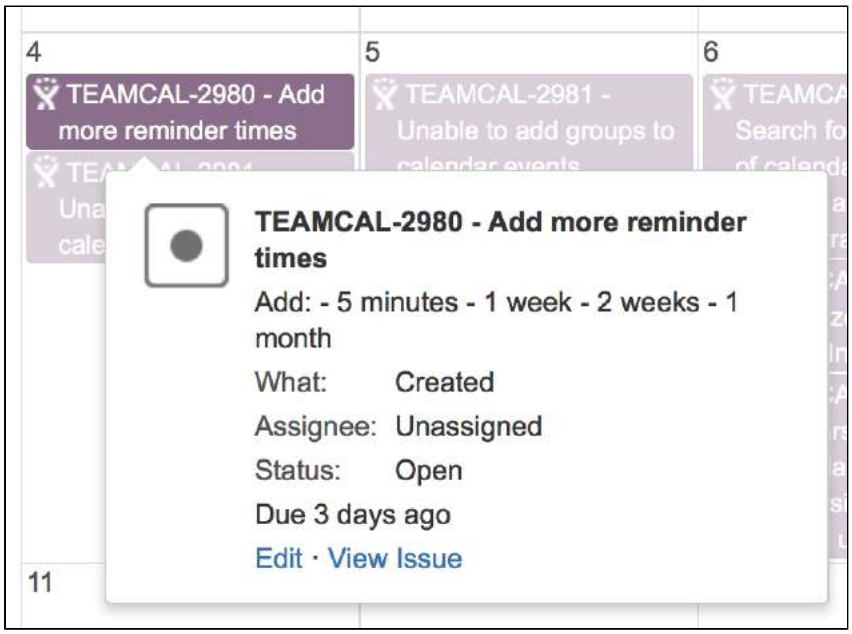

Project release (shows release date)

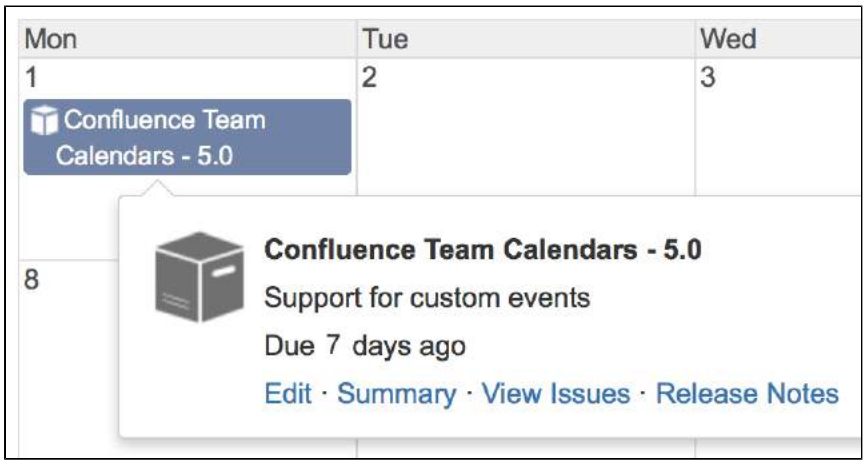

Software sprint (shows start and end dates)

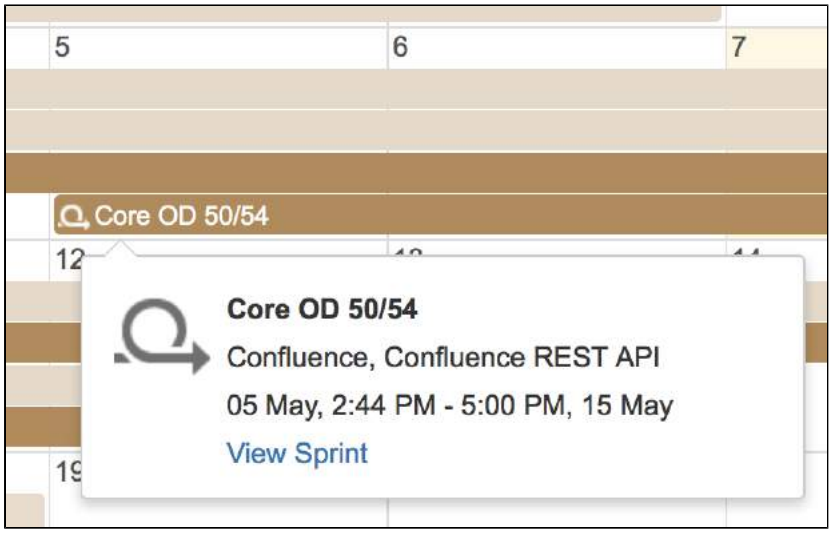

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Notes

- You can't add content to JIRA calendars as they source all of their date information from the JIRA system. To add issues and version deadlines, add them in the JIRA application's interface.
- You can drag-and-drop your JIRA issue dates and project versions to re-schedule them right from Team Calendars.
- When you create a JIRA calendar, all users in your system can see the calendar name in Confluence. The actual calendar events(issues, versions, sprints), however, automatically respect your JIRA application permissions, so it's possible for users to be subscribed to a JIRA calendar and not see all events.
- In order to make sure calendars load quickly, JIRA events are cached. You can configure the cache expiry (time to live) in **Team Calendars** > **Global settings** - decrease the cache expiry time for more accuracy.

Team Calendars for Confluence is now part of Confluence Data Center
To get access to the features described on this page upgrade to Confluence Data Center 7.11 or later. Cant upgrade yet? Depending on your current Data Center version, you can access these features by installing the latest version of the app (at no cost). See our FAQ for all the details

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Event Types

Each calendar includes a standard set of event types, which you can use to classify the different events in the calendar. You can alsocreate custom event typesto capture events that don't fit into the standard event types.

On this page:

- Standard event types
- Color-code event types
- Filter events

Related pages:

- Custom Event Types
- Add Events
- Add Jira Events

Standard event types

Plan your people

- **Leave** For tracking holidays and breaks. Who is out of the office and when.
- **Travel**For tracking your team's travel plans.
- **Birthdays**To help you celebrate those special moments with your team members.
- **Events** A generic event type where you can create any type of event.

Plan your projects

Team Calendars integrates withJIRA applicationsand allows you to plan your JIRA projects and see how your team's availability may impact them. You can create JIRA project events for:

- **JIRA Versions**To plan your releases.
- **JIRA Software Sprints**Visualize your teams Agile Sprints inside your calendar to see how your team's leave could impact your project.
- **JIRA Issue Dates**

Learn more about JIRA events

If you don't want to use any of the standard event types, hide them by choosing the calendar's drop-down menu, then select**Edit > Event Types**. Choose the minus icon to hide the event type.

Color-code event types

All events are color-coded according to the event type they belong to.To change the color of events in a particular event type:

1. Hover over the event type and click the••• that appears to the right
2. Choose the appropriate color for the event type

Changing event type colors changes the appearance of its events for you, but not for other viewers of the calendar. Each person viewing the calendar is free to set their preferred color scheme.

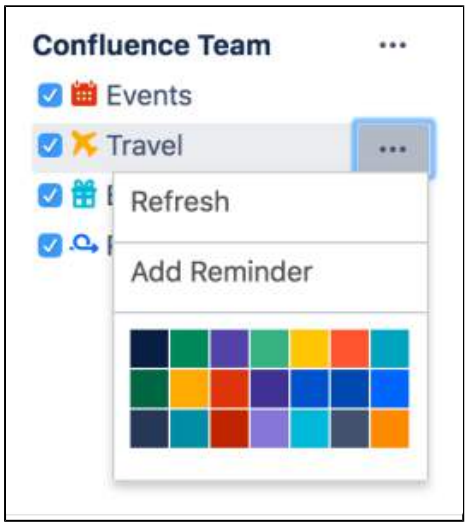

Filter events

If you've got a lot of events on your Team Calendar, you may want to show or hide certain event types or calendars to reduce clutter. Click on an event type to show or hide its events, or hit ••• next toa calendar and choose **Hide events**to hide all events in that calendar.

Team Calendars for Confluence is now part of Confluence Data Center
To get access to the features described on this page upgrade to Confluence Data Center 7.11 or later. Cant upgrade yet? Depending on your current Data Center version, you can access these features by installing the latest version of the app (at no cost). See our FAQ for all the details

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Custom Event Types

Event types allow you to categorize your events, making them easier to view and manage; custom event types take this a step further by giving you complete freedom to track any kind of event you can think of so if you need to create events fortraining,email campaigns, outages, or anything else, Team Calendars is up to it.

On this page:

- Add a custom event type
- Delete a custom event type

Related pages:

- Event Types
- Add Events

Add a custom event type

To add a custom event type:

1. Choose ••• to the right of thecalendar's name, then choose**Edit**
2. Choose**Event Types**on the left, then choose**Add new event type**
3. Enter the**Event Name**, select an**Event Icon,**then choose**Save**.

That's it! Your new event type won't appear in the sidebar immediately, but it'll show up as soon as youadd an eventof that type.

You'll need editing privileges for the calendar to add custom event types.

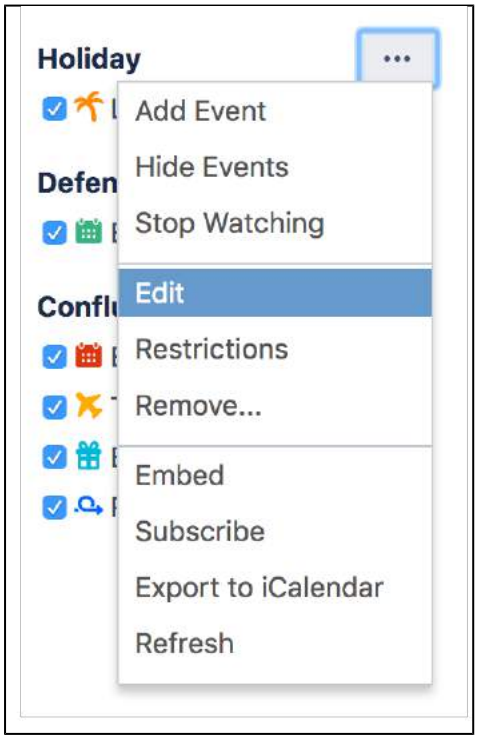

To quickly add a new custom event type, choose**Add Event**, then choose**Add new event type**from the**Event Type**menu.

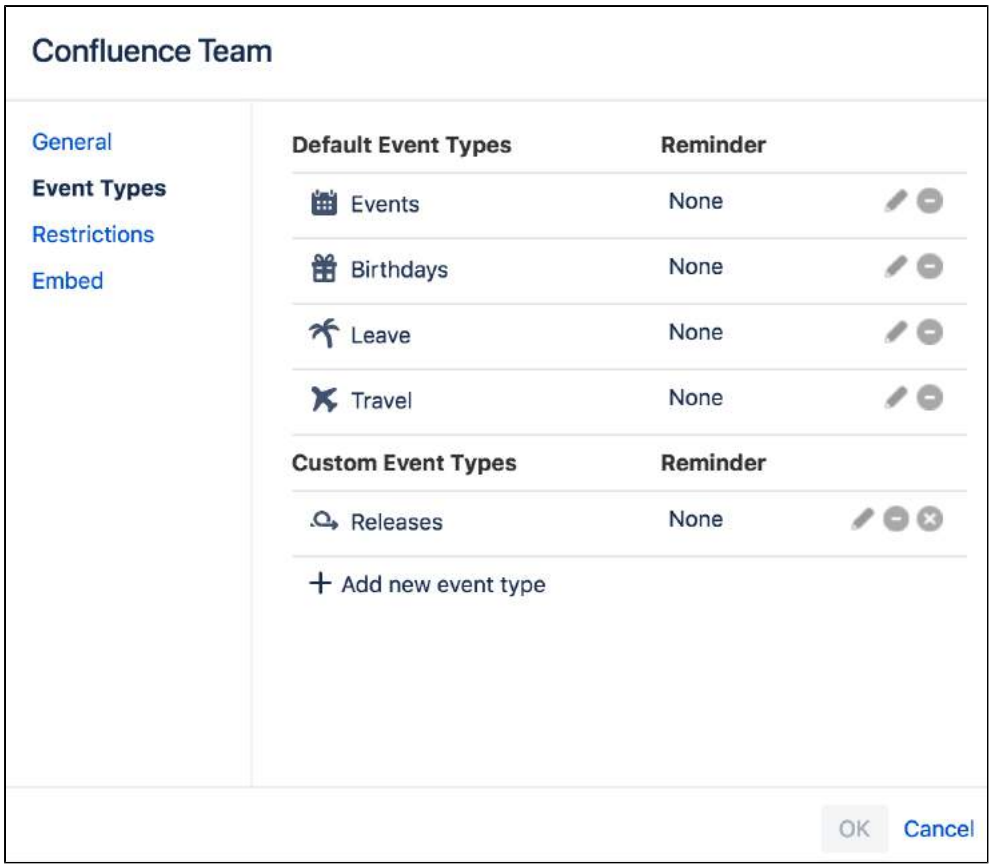

Delete a custom event type

If you no longer need a custom event type, you can delete it by choosing**Edit**>**Event Types**, then choose the icon to the right of the event type.

Deleting a custom event type will delete all events within that event type, so, if there are any events you want to keep, make sure you move them to another event type before deleting.

Team Calendars for Confluence is now part of Confluence Data Center
To get access to the features described on this page upgrade to Confluence Data Center 7.11 or later. Cant upgrade yet? Depending on your current Data Center version, you can access these features by installing the latest version of the app (at no cost). See our FAQ for all the details

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Reminders

There are some events that are just too important to miss, and you wouldn't want anyone in your team to miss them either. If you have edit permission for a calendar, you can set reminders for eachevent typeso that anyonewho's added the calendar to their 'My Calendars'pagewill receive an email at your chosen time.

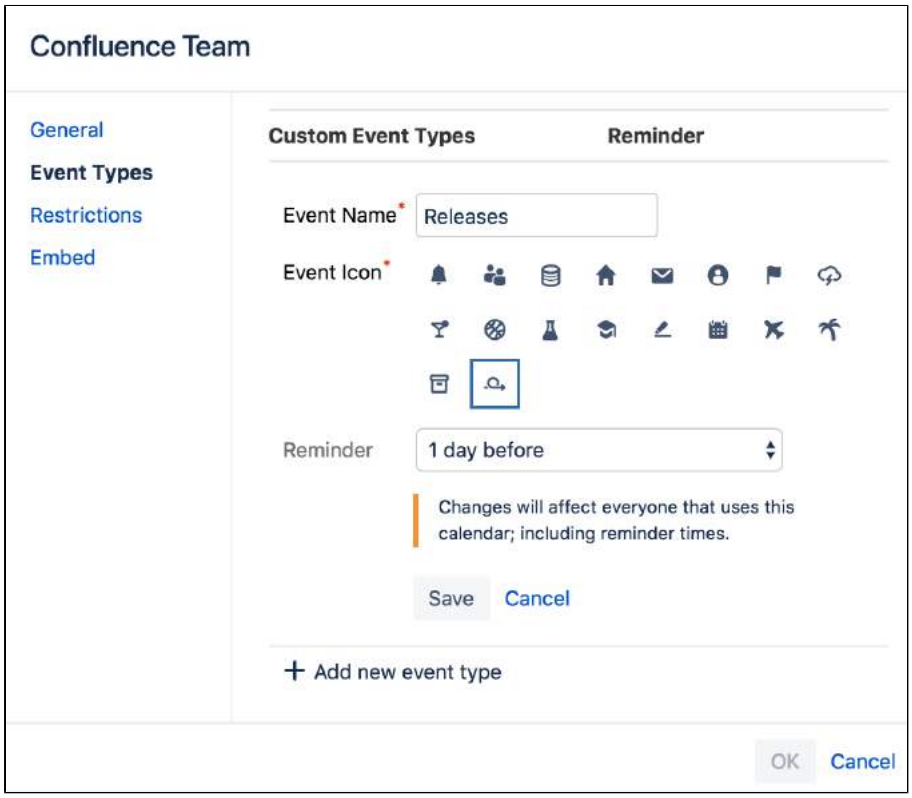

Add reminders to event types

To add a reminder to an event type:

1. Choose ••• >**Edit**
2. Choose**Event Types**
3. Choose the pencil icon to edit an event type
4. Add your reminder time and choose**Save**
   Repeat steps 3 and 4 for each event type you're adding a reminder to.
5. Choose **OK**

If you only want to add reminders to one event type, choose the ••• to the right of the event type, then choose**Add a reminder**.

From a reminder email you can click to view the event, or turn off reminders for the event type if you don't need them.To*remove*reminders for an event type, follow the above procedure and set the reminder time to 'None'.

Remember, when you change reminders for an event type, you're changing reminders for anyone who's added the calendar to their My Calendars page.

Turn reminder emails on or off

You can decide whether or not you want to receive reminder emails for each event type on your My Calendars page. This isn't about adding or removing reminders for an event type; it's essentially opting-in to or opting-out of the reminder emails.

To turn reminders on or off:

1. Choose **Calendars** in the Confluence header
2. Choose the ••• to the right of an event type
3. Select 'Turn on reminders' or 'Turn off reminders'

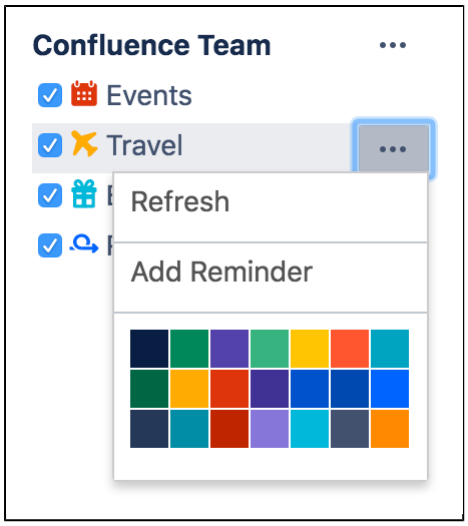

If it says 'Add a reminder' when you get to step 3, it means there's no reminder set for the event type. Only choose that option if you want to *set* reminders for the event type, which will be sent to everyone who's added the calendar to their 'My Calendars' page.

You can also opt-out of reminder emails for an event type by clicking the 'Stop reminding me' linkinthe reminder email

Team Calendars for Confluence is now part of Confluence Data Center
To get access to the features described on this page upgrade to Confluence Data Center 7.11 or later. Cant upgrade yet? Depending on your current Data Center version, you can access these features by installing the latest version of the app (at no cost). See our FAQ for all the details

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Embed Calendars on Confluence Pages

You can embed a calendar on any page in Confluence, making it easy to track and manage events. When you embed a calendar on a page, the calendar will also appear in the**Calendars** section of the space the page is in.

Below are a few ways of embedding a calendar on a Confluence page.

On this page:

- Paste the calendar's link
- Add the Team Calendar macro

Related pages:

- Subscribe to Team Calendars from Third-Party Calendars
- Create, Add, and Edit Calendars

Paste the calendar's link

1. Choose**Calendars** in the Confluence header or space sidebar
2. Hit the••• to the right of the calendar name
3. Choose**Embed**
4. Copy the calendar link provided
5. Paste the link on any page in Confluence

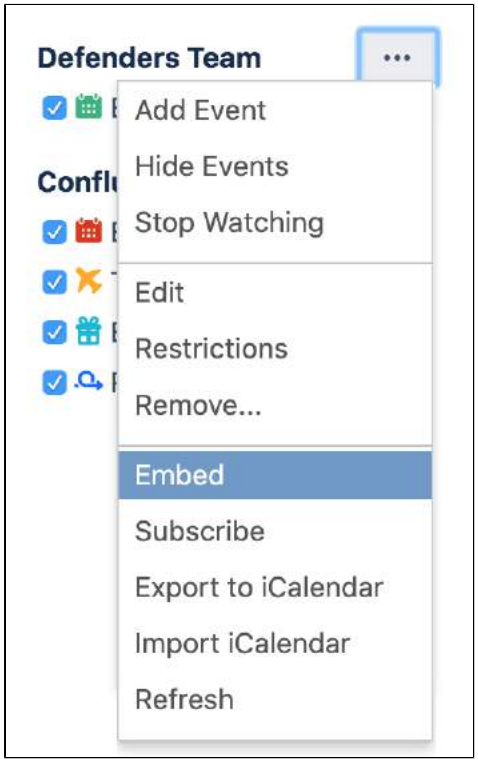

Add the Team Calendar macro

1. While editing a page, choose**Insert > Team Calendar**
2. Select**Add Existing Calendar**
3. Search for and select the calendar and choose**Add**

SeeTeam Calendar Macro for more information about using this macro.

If you'd like to edit properties for the Macro, choose the Team Calendars Macro that you added to the page and choose**Edit** .

Team Calendars for Confluence is now part of Confluence Data Center
To get access to the features described on this page upgrade to Confluence Data Center 7.11 or later. Cant upgrade yet? Depending on your current Data Center version, you can access these features by installing the latest version of the app (at no cost). See our FAQ for all the details

Watch a Calendar

Watching a calendar is a great way to stay up-to-date on changes made in the calendar, without necessarily adding it to your My Calendars page. When you watch a calendar, you'll be sent an email notification when any of the following events occur:

- A new event is created
- An existing event is edited
- An event is deleted
- The entire calendar is deleted

When you watch a Confluence space it includes watching all calendars in that space, so there's no need to manually watch calendars for your team or project.

On this page:

- Watch a calendar
- Stop watching a Calendar

Related pages:

- Create, Add, and Edit Calendars
- Subscribe to Team Calendars from Third-Party Calendars
- Embed Calendars on Confluence Pages

Watch a calendar

To watch a calendar, choose ••• to the right of a calendar name and select **Watch**. Notifications will be sent to your email address.

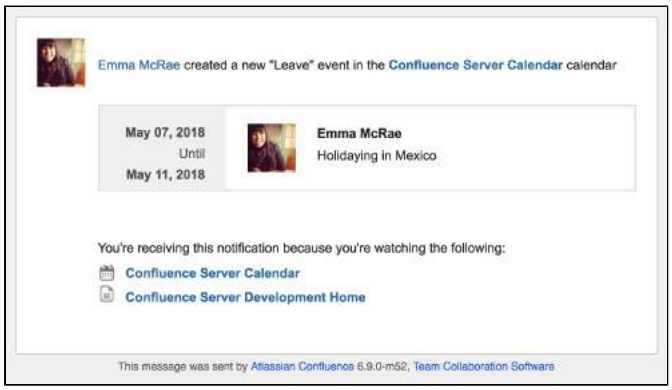

Stop watching a Calendar

You can also stop watching a calendar, if you no longer want to receive email notifications about changes to the calendar.

To stop watching a calendar, choose ••• to the right of the calendar name and select **Stop Watching**.

If the menu says 'Watching', you're watching the calendar because you're either watching a page the calendar is embedded on or you're watching the related space for the calendar. To stop watching in this case, you'd need to stop watching the relevant page or space containing the calendar.

Administrators can limit notifications to only people directly watching a calendar. Go to > **General Configuration** > **Team Calendars** and choose **Don't notify space or page watchers**.

Team Calendars for Confluence is now part of Confluence Data Center
To get access to the features described on this page upgrade to Confluence Data Center 7.11 or later. Cant upgrade yet? Depending on your current Data Center version, you can access these features by installing the latest version of the app (at no cost). See our FAQ for all the details

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Share Calendars

When viewing space calendars (all calendars in a particular Confluence space) or a single calendar from a space, use theSharebuttonto email anyone a link to space calendars or a single calendar. You can either copy the Share link from the share dialog, or enter a Confluence user, group or email address.The option to add people is only available if your site has a mail server configured.

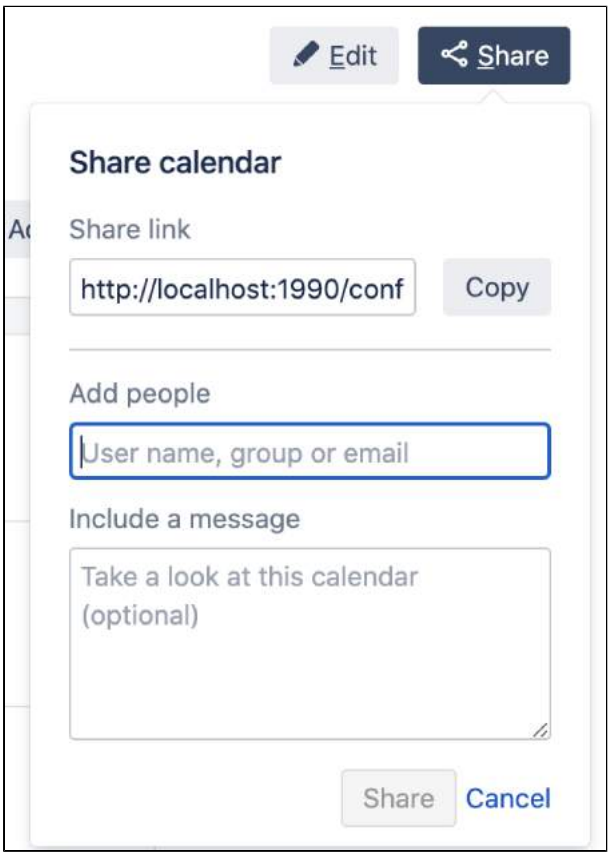

Share space calendars

1. Go to the space and select**Calendars**in the sidebar.
2. Select**Share**at the top right of the page.
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 (or use the trash icons to remove people from the list).
4. Enter an optional message.
5. Select**Share**to send the link via email.

Share a single calendar from a space

If a space has multiple calendars, but you only want to share one of them, follow the steps belowto get to the calendar's details page and share it.

1. Go to the space and choose**Calendars**in the sidebar.
2. Select the name of the calendar you want to share.You'll be taken to the detail page for the calendar it's just a page with that single calendar on it.
3. Select**Share**at the top right of the page.
4. 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 (or use the trash icons to remove people from the list).
5. Enter an optional message.
6. Select**Share**to send the link via email.

Team Calendars for Confluence is now part of Confluence Data Center
To get access to the features described on this page upgrade to Confluence Data Center 7.11 or later. Cant upgrade yet? Depending on your current Data Center version, you can access these features by installing the latest version of the app (at no cost). See our FAQ for all the details

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Subscribe to Team Calendars from Third-Party Calendars

Subscribing to Team Calendars from your favourite third-party calendar app is a great way to view your Team Calendars events when you're not in Confluence.

There are two ways to synchronize a Team Calendar with your calendar app:

- Two-way sync (CalDAV) - allows you to view and update Team Calendars events in your app.
- One-way sync (iCal) - allows you to view, but not update, Team Calendars events in your app.

Not all calendar apps support two-way sync, so check the table below to find out if your app supports it.

	Calendar app	Sync options
	Microsoft Outlook (Windows desktop)	One-way and two-way sync
	Microsoft Outlook (browser)	One way sync
	Apple Calendar (MacOS)	One-way and two-way sync
	Apple Calendar (iOS)	One-way and two-way sync
	Google Calendar (browser)	One way sync
	Google Calendar (Android)	One-way and two-way sync
	Thunderbird	One-way and two-way sync

You can also Subscribe to Third-Party Calendars from Team Calendars.

Other calendar apps

If the calendar app you use isn't listed above, you may still be able to subscribe to Team Calendars. Check your app's documentation to find out whether it supports the iCal or CalDAV standard.If it supports iCal, then you should be able to set up one-way syncronization. If it supports CalDAV, you should be able to set up two-way synchronization.

To subscribe to a calendar app that's not listed above:

In Confluence:

1. Choose the **Subscribe** button at the top of your calendar.
2. Select either**iCal** or **CalDAV**from the**Calendar app**dropdown (check your calendar app's documentation to find out which standard it supports).

3. Copy the**Calendar URL**.

Then, in your Calendar app, follow the steps in your app's documentation to add or subscribe to the calendar.

Timezones

Timezones can be confusing! Here's how we decide what time to display.

- We display the time of an event in Confluence using the timezone set in your Confluence profile -find out how to update this.
- We display the time of an event in your calendar app using the timezone set in your app or device.
- We store events using the timezone defined in the Team Calendar itself, and then convert them to the appropriate timezone when you view the event.

You should always see the correct time, but if you change locations, make sure you update your Confluence profile so you don't see a discrepancy between your calendar app and Confluence.

Confluence context path

CalDav uses Confluence's root URL as it's well-known URL. If you've configured Confluence with a context path, for examplewww.yoursite.com/confluence, your calendar app will not be able to connect to your Team Calendars.

There are two ways to resolve this issue.You can either:

- Configure your proxy so that any URL that **does not** end with `/confluence`(or whatever your context path is) or with `/synchrony`(if you're using collaborative editing) always redirects to `/confluence`. This option is only suitable if Confluence is the only application accessed via that domain.
- Configure your proxy to redirect the following calls to your URL, with context path. In the examples below our context path is `/confluence`.

Apache

```
RewriteRule ^/.well-known/(.*)$ /confluence/.well-known/$1 [R=301,NC,L]
RewriteRule ^/plugins/servlet/team-calendars/caldav/(.*)$ /confluence/plugins/servlet/team-calendars
/caldav/$1 [R=301,NC,L]
```

The order of directives is important in Apache. These rewrites should be placed*before* the Synchrony and Confluence location blocks. SeeUsing Apache with mod_proxyfor general information about configuring your proxy.

Nginx

```
rewrite ^/.well-known/(.*)$ /confluence/.well-known/$1;
rewrite ^/plugins/servlet/team-calendars/caldav/(.*)$ /confluence/plugins/servlet/team-calendars
/caldav/$1;
```

While order is not as important in Nginx, we tested this config with the rewrites before the Synchrony and Confluence location blocks. SeeRunning Confluence behind NGINX with SSLforgeneral information about configuring your proxy.

Reverse Proxies

If you are using NGINX or Apache HTTPD, don't forget to enable support for DAV methods on them, else, synchronisation may fail due to the following error:

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

```
405 Error
    </p>

    <h2 class="content-subhead">Method not allowed!</h2>
    <p>


    The REPORT
    method is not allowed for the requested URL.

</p>


<h2 class="content-subhead">Error 405</h2>
<p>
    Server is named <a href="/">confluence-atlas.atlassian.com</a>.

    Webserver program <code>Apache/2.4.29 (atlassian@atlassian.com - CEL 7.x) OpenSSL/1.1.0f mod_fcgid/2.3.9
</code> served
    the request with unique id <code>W1idbjAGhETV97vZN7bzyAAAAWU</code>
    from <code>64.103.79.190:3443</code> socket
    on <code>Wed Jul 25 16:55:26 2018</code>.

</p>
```

Modules:

For NGINX - Link

For HTTPD - Link

Troubleshooting

If you have problems adding a Team Calendar to your calendar app, you should check the documentation for your particular app. Here's some issues we know about:

- You can't subscribe to a Team Calendar on Outlook.com if your Confluence site is not accessible outside your network.
- If your Confluence site isn't accessible outside your network, you may need to be connected to your network when you add the subscription, and to sync event data.
- If your calendar app supports calendar discovery, your calendar must be listed on your **My Calendars** screen, so it is available to sync. To add a calendar to **My Calendars** choose **Calendars** on the header > **Add Calendar** > **Add existing calendar.** We'll also prompt you in the subscribe dialog, if the calendar is not listed on **My Calendars**.
- If you're unable to connect to your calendar, and your site has a context path (for example, you access it fromwww.mysite.com/confluence)your admin may need to make some changes to the proxy configuration. SeeConfluence context path above.
- You can always enable more logging from CalDavSynchronyzer as listed here for more logging and troubleshooting aid.

Team Calendars for Confluence is now part of Confluence Data Center
To get access to the features described on this page upgrade to Confluence Data Center 7.11 or later. Cant upgrade yet? Depending on your current Data Center version, you can access these features by installing the latest version of the app (at no cost). See our FAQ for all the details

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Subscribe to Team Calendars from Microsoft Outlook

Outlook support

Subscribing to Team Calendars from Microsoft Outlook has been tested and works with Outlook 2010, 2013, and 2016 on Windows.

The Outlook for Mac 2011/2016 desktop applications don't provide the ability to subscribe to an internet calendar, but there is a workaround that may work for you.

On this page:

- Outlook support
- Subscribe to a Team Calendar from Outlook (Windows)
  - Subscribe with two-way synchronization (CalDAV)
  - Subscribe with one-way synchronization (iCal)
- Subscribe to a Team Calendar from Outlook (Mac OS)

Subscribe to a Team Calendar from Outlook (Windows)

There are two ways to synchronize your calendar in Outlook:

- **Two-way sync** (CalDAV) - allows you to view and update Team Calendars events in Outlook. This option requires a free CalDAV Synchronization plugin for Outlook, and is only available for Windows users.
- **One-way sync** (iCal) - allows you to view, but not update, Team Calendars events in Outlook.

Subscribe with two-way synchronization (CalDAV)

Two-way synchronization is only available for people using the Microsoft Outlook desktop application on Windows, with the CalDAV Synchronization plugin installed.

1. Download and install the CalDav Synchronization plugin

Before you begin:

1. Download a CalDAV plugin for Outlook.
   We usedCalDAV Synchronizer, whichis free. Other plugins may be available.
2. Install the plugin. You may need to talk to your admin if you don't have permission to install applications on your PC.

2. Grab your Team Calendars URL

In Confluence:

1. Choose the **Subscribe** button at the top of your calendar.
2. Select **Outlook** from the **Calendar app** dropdown.

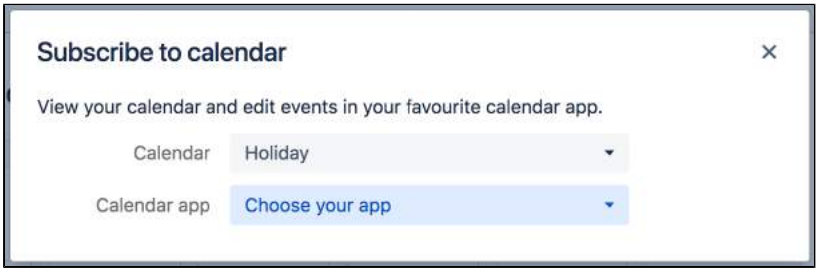

3. If prompted, add the calendar to**My Calendars**.
4. Copy the **Calendar URL**.
   Because the CalDAV Synchronization plugin supports calendar discovery, this will be your Confluence URL, not the URL of an individual calendar. You'll be able to choose which calendar (that you've added to My Calendars)to sync in Outlook.

3. Subscribe to the calendar in Outlook

In your Outlook desktop application:

1. Choose **CalDAV Synchronizer** from the ribbon / toolbar.
2. Choose **Synchronization profiles**.

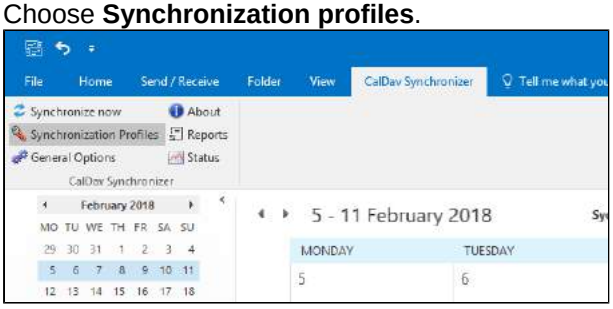

3. Click the **+** button to add a new profile.

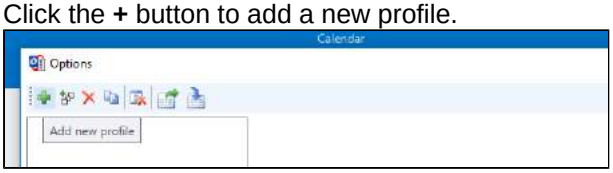

4. Choose **Generic CalDAV / CardDAV** from the **Profile type** screen.
5. Give your profile a name, for example **Confluence Team Calendars**
6. Select an existing **calendar folder**, or create a new one.

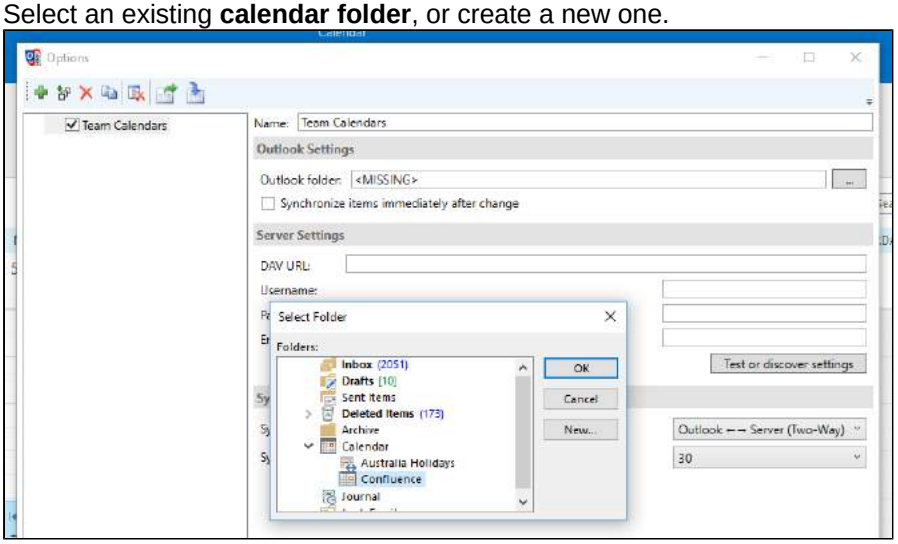

7. Enter the **Confluence URL** you copied earlier in the **DAV URL** field.
8. Enter your Confluence **username** and **password**.
9. Click the **Test or discover settings** button to connect to Team Calendars.
10. If the test is successful you'll see a list of resources. Choose the calendar you want to subscribe to.

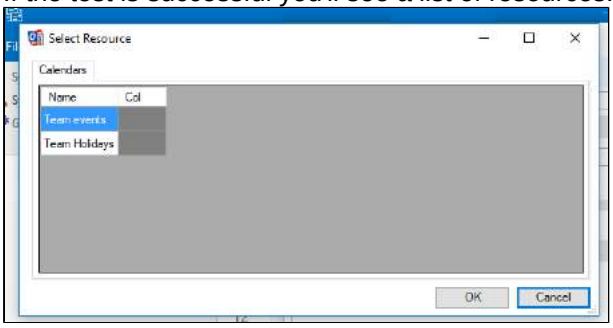

11. Choose how often you want Outlook to sync back to Team Calendars.
    The default is every 30 minutes, but you may want to sync your calendar more often.
12. Click **OK** to save your new profile.
13. Back on the **CalDAV Synchronizer** tab, choose **Synchronize now** to sync the calendar for the first time.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

14. If your calendar doesn't appear, you may need to select it from the**My Calendars** list.

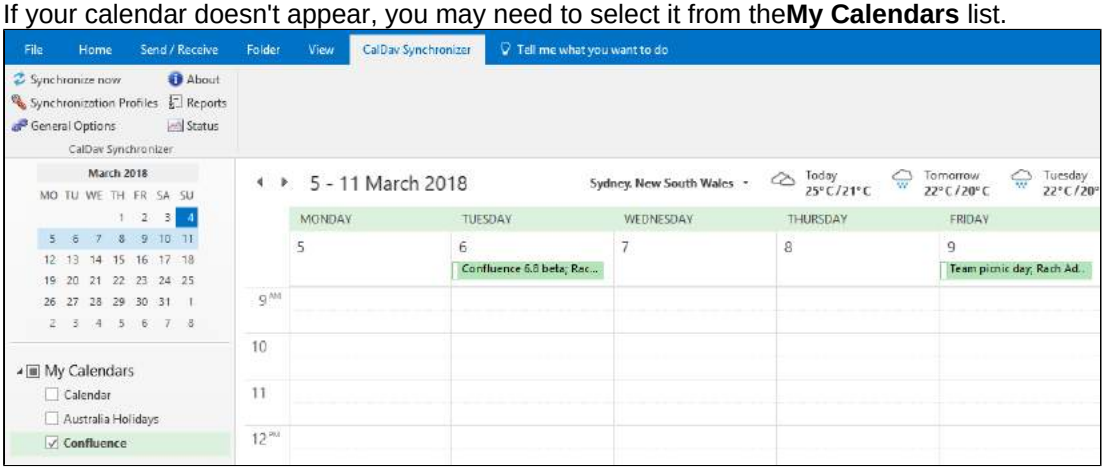

Seeing something different?These instructions are specific to the CalDAV Synchronizerplugin. If you're using a different plugin, check its documentation to find out how to do this step.

Subscribe with one-way synchronization (iCal)

One-way synchronization means that you can view, but not update, Team Calendars events in Outlook. This option is best for people who are not able to install the CalDAV Synchronization plugin.

1. Grab your Team Calendars URL

1. Choose ••• to the right of the calendar name, then choose**Subscribe**.
2. Select**iCal**from the**Calendar app**dropdown (don't choose Outlook for one-way sync).
3. Copy the**Calendar URL**.

2. Subscribe to the calendar in Outlook

1. Choose**Calendar** at the bottom left of the app.
2. Choose**Open Calendar**> **From Internet** in the ribbon.

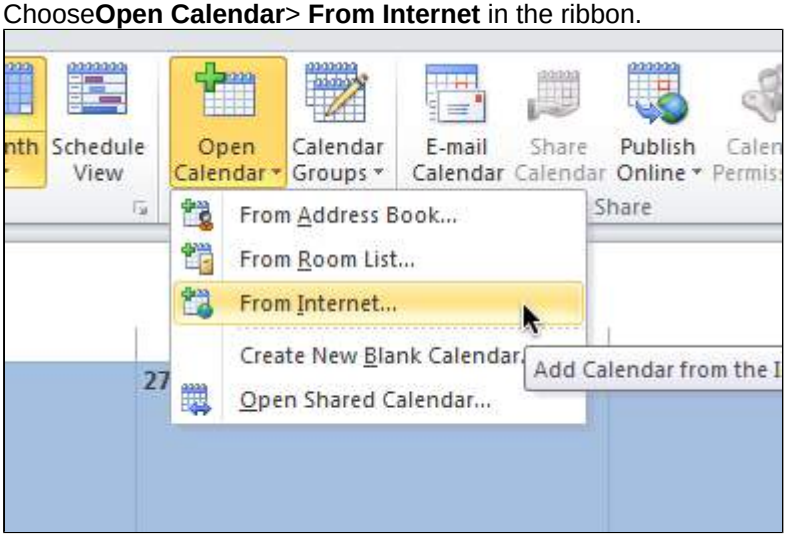

3. Paste the calendar address and choose**OK**.

Depending on your authentication setup in Confluence, you may be prompted to enter your Confluence username and password.

Subscribe to a Team Calendar from Outlook (Mac OS)

Outlook for Mac 2011/2016 desktop applications don't provide the ability to subscribe to an internet calendar.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

If you have access to Outlook Web Access (OWA), you may be able toopen an Internet Calendar, andhave it appear in your Outlook client.

Team Calendars for Confluence is now part of Confluence Data Center
To get access to the features described on this page upgrade to Confluence Data Center 7.11 or later. Cant upgrade yet? Depending on your current Data Center version, you can access these features by installing the latest version of the app (at no cost). See our FAQ for all the details

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Subscribe to Team Calendars from Outlook on the web

These instructions apply to Office 365, Outlook.com and Outlook on the web. If you are using the Outlook desktop app, head toSubscribe to Team Calendars from Microsoft Outlookto find out how to subscribe.

Subscribe to Team Calendars from Outlook on the web

Only one-way sync (iCal) is available for Outlook on the web. This means you can view, but not update, Team Calendars events in Outlook in your browser.

Subscribe with one-way synchronization (iCal)

These instructions are for Outlook.com (as at March 2018). Your version may differ slightly.

1. Grab your Team Calendars URL

In Confluence:

1. Choose the **Subscribe** button at the top of your calendar.
2. Select**Outlook (browser)**from the**Calendar app**dropdown.

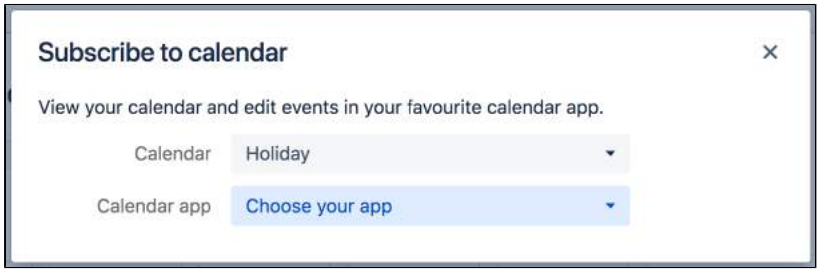

3. Copy the**Calendar URL**.

2. Subscribe to the calendar in Outlook on the web

In Outlook calendar in your browser:

1. Choose**Add Calendars**>**From Internet**.
2. Enter the Team Calendar URL you copied above.
3. Give your calendar a name. This is how it will appear in the calendar list.

You can now view, but not edit, Team Calendar events in Outlook on the web.

⚠ If your Confluence site is not accessible from outside your network, for example if you need to use a VPN to access it from outside your office, you won't be able to subscribe to your Team Calendars in Outlook on the web. Outlook will return a "Couldn't open calendar. Calendar address isn't formatted correctly error".

ⓘ **Team Calendars for Confluence is now part of Confluence Data Center**
To get access to the features described on this page upgrade to Confluence Data Center 7.11 or later. Cant upgrade yet? Depending on your current Data Center version, you can access these features by installing the latest version of the app (at no cost). See our FAQ for all the details

Subscribe to Team Calendars from Apple Calendar

If you want your Team Calendars events to appear in the built-in Calendar app on your Mac, you can subscribe to Team Calendars in a few easy steps.

On this page:

- Subscribe to Team Calendars from Apple Calendar
  - Subscribe with two-way synchronization (CalDAV)
  - Subscribe with one-way synchronization (iCal format)

Subscribe to Team Calendars from Apple Calendar

There are two ways to synchronize your calendar in Apple Calendar:

- **Two-way sync** (CalDAV) - allows you to view and update Team Calendars events in Apple Calendar.
- **One-way sync** (iCal) - allows you to view, but not update, Team Calendars events in Apple Calendar.

Subscribe with two-way synchronization (CalDAV)

Two-way synchronization allows you to view and update Team Calendars events in macOS.

These instructions are for Apple Calendar 10.0 on MacOS High Sierra. Your version may differ slightly.

1. Grab your Team Calendars URL

In Confluence:

1. Choose the **Subscribe** button at the top of your calendar.
2. Select**Apple Calendar (macOS)**from the**Calendar app**dropdown.

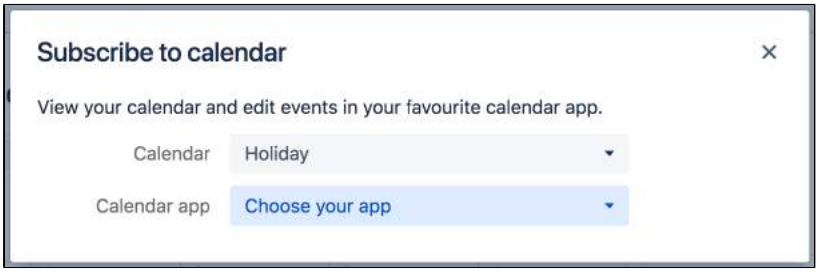

3. If prompted, add the calendar to**My Calendars**.
4. Make a note of the**Calendar URL**.
   Because Apple Calendar supports calendar discovery, this will be your Confluence URL, not the URL of an individual calendar. You'll be able to choose which calendars (that you've added to My Calendars) to sync in the app.

2. Add the calendar account in Apple Calendar

In the Calendar app on your Mac:

1. Choose **Calendar** > **Add Account**.
2. Choose **Other CalDAV account**.
3. Select **Manual** from the Account Type drop down.
4. Enter your Confluence **username** and **password**.
5. Enter your **Confluence URL**.
   Apple Calendar supports calendar discovery, so you don't need the URL of a specific calendar.
6. Click **Sign in**.

All available calendars will be listed under your Confluence URL. Deselect any calendars you don't want to see.

You can rename the calendar group. Head to **Calendar** > **Accounts** and edit the account listing.

Subscribe with one-way synchronization (iCal format)

These instructions are for Apple Calendar 10.0 on MacOS High Sierra. Your version may differ slightly.

1. Grab your Team Calendars URL

In Confluence:

1. Choose ••• to the right of a calendar name, then choose **Subscribe**.
2. Select **iCal** from the **Calendar app** dropdown (don't choose Apple Calendar for one-way sync).
3. Copy the calendar's URL.

2. Subscribe to the calendar in Apple Calendar

In the Calendar app on your Mac:

1. Choose **File** > **New Calendar Subscription...**
2. Paste the calendar URL and choose **Subscribe**.
3. Configure any other calendar information and choose **Save**.

You can now view, but not edit, Team Calendar events in Apple Calendar.

Team Calendars for Confluence is now part of Confluence Data Center
To get access to the features described on this page upgrade to Confluence Data Center 7.11 or later. Cant upgrade yet? Depending on your current Data Center version, you can access these features by installing the latest version of the app (at no cost). See our FAQ for all the details

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Subscribe to Team Calendars from Apple iOS Calendar

You can subscribe to Team Calendars from your iPhone or iPad device using the built-in iOS Calendar.

Subscribe to a Team Calendar from iOS

There are two ways to synchronize your calendar in iOS:

- **Two-way sync** (CalDAV) - allows you to view and update Team Calendars events in iOS.
- **One-way sync** (iCal) - allows you to view, but not update, Team Calendars events in iOS.

Subscribe with two-way synchronization (CalDAV)

Two-way synchronization allows you to view and update Team Calendars events in iOS.

These instructions are for iPhone, running iOS 11. Your version may differ slightly.

1. Grab your Team Calendars URL

In Confluence:

1. Choose the **Subscribe** button at the top of your calendar.
2. Select**Apple Calendar (iOS)**from the**Calendar app**dropdown.

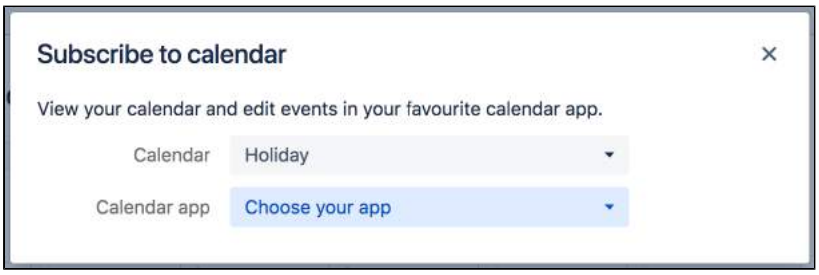

3. If prompted, add the calendar to**My Calendars**.
4. Make a note of the**Calendar URL**.
   Because the iOS app supports calendar discovery, this will be your Confluence URL, not the URL of an individual calendar. You'll be able to choose which calendar(that you've added to My Calendars) to sync in the app.

2. Subscribe to the calendar in the iOS Calendar app

On your iOS device:

1. Go to **Settings** > **Accounts and passwords** > **Add account**.
2. Choose **Other** >**Add CalDAV account**.
3. Enter the**Confluence URL**you copied earlier in the**Server**field.
4. Enter your Confluence **username** and **password** and follow the prompts to add the account.
5. In the Calendar app, choose **Calendars** and select the calendars you want to show.

Subscribe with one-way synchronization (iCal)

One-way synchronization means that you can view, but not update, Team Calendars events in your app.

1. Grab your Team Calendars URL

1. Choose ••• to the right of the calendar name, then choose**Subscribe**.
2. Select**iCal**from the**Calendar app**dropdown (don't choose Apple Calendar for one-way sync).
3. Copy the**Calendar URL**.

2. Subscribe to the calendar in the iOS Calendar app

On your iOS device:

1. Go to **Settings** > **Accounts and passwords** > **Add account**.
2. Choose **Other** > **Add Subscribed Calendar**.
3. Paste the calendar URL and choose **Next**.
4. Enter any other required information and choose **Save**.

You can now view, but not edit, Team Calendar events in Apple Calendar for iOS.

Troubleshooting

Firewall or VPN issues

If your Confluence site isn't accessible outside your network (for example, you need to use a VPN to access it when you're out of the office), you will need to be connected to your network when you create the account. Your calendar will also only successfully sync when you're connected to your network. This means you may not see updates to events until the next time you are connected to your network or VPN.

SSL error

We have seen that iOS occasionally throws an SSL error when you attempt to add the Team Calendars account. If this happens, we recommend removing the account, and trying again. This often fixes the problem.

Calendars don't appear

It can take a few minutes for your calendars to appear in the Calendar app, after you've added the account. Tap the **Calendars** link at the bottom of the Calendar app to see the list of calendars from each account, and pull down to refresh the list.

Team Calendars for Confluence is now part of Confluence Data Center
To get access to the features described on this page upgrade to Confluence Data Center 7.11 or later. Cant upgrade yet? Depending on your current Data Center version, you can access these features by installing the latest version of the app (at no cost). See our FAQ for all the details

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Subscribe to Team Calendars from Google Calendar

Integrating Team Calendars with your Google Calendar is a great way to keep track of your team's leave, travel, rosters, and projects, all in one place.

Subscribe to Team Calendars from Google Calendar

Only one-way sync (iCal) is available for Google Calendar. This means you can view, but not update, Team Calendars events in Google calendar in your browser.

Subscribe with one-way synchronization (iCal)

These instructions are for Google G Suite Calendar (as at March 2018). Your version may differ slightly.

1. Grab your Team Calendars URL

In Confluence:

1. Choose the **Subscribe** button at the top of your calendar.
2. Select**Google calendar**from the**Calendar app**dropdown.

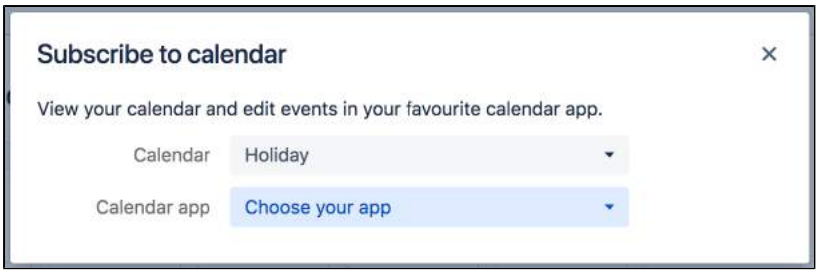

3. Copy the**Calendar URL**.

2. Subscribe to the calendar in Google Calendars

In Google calendar in your browser:

1. Choose **Add other calendars** > **From URL**.

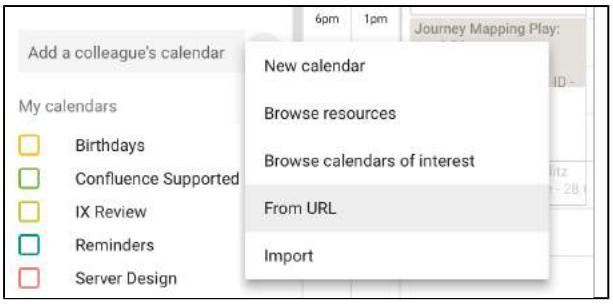

2. Paste your Team Calendars URL.
3. Choose **Add Calendar**.

The Calendar will be listed under Other Calendars, usually by URL.To rename the calendar Go to **Settings** > **Settings for other calendars**, select your Team Calendar and edit the calendar name.

Known Issues and Limitations

Known issues with Team Calendars and Google Calendars integration. These limitations are due to different issues in Google Calendar.

- There is currently a known issue with Google Calendar where a calendar subscription can't be added, or does not display events ~~CONFSERVER-53690~~ CLOSED
- Subscribed Team Calendars names are truncated if they have a space in the title (More information:TEAMCAL-455).
- Google Calendar refresh times are delayed. They range from 3 hours to 24 hours. (More information:TEAMCAL-458).
- Subscriptions are read-only (you cannot modify events from Google Calendar, as it does not provide a CalDAV option).

Troubleshooting

Adding a calendar to Google Calendars may give you this error: "Could not fetch the url because robots.txt prevents us from crawling the url.". This means your server administrator is disallowing search engines from indexing your servers. A server administrator will need to modify your `robots.txt` to allow Google to index Team Calendars. In your`robots.txt`file, add the following line:

```
User-agent: *
...
Allow: /rest/calendar-services/1.0/calendar/export/subcalendar/private/*
```

Team Calendars for Confluence is now part of Confluence Data Center
To get access to the features described on this page upgrade to Confluence Data Center 7.11 or later. Cant upgrade yet? Depending on your current Data Center version, you can access these features by installing the latest version of the app (at no cost). See our FAQ for all the details

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Subscribe to Team Calendars from Android Calendar

You can subscribe to Team Calendars from your phone or tablet using the built-in Google Calendar on Android devices.

Subscribe to a Team Calendar from Google Calendar on Android

There are two ways to synchronize your calendar in Android device:

- **Two-way sync** (CalDAV) - allows you to view and update Team Calendars events in Google Calendar. Requires a CalDav synchronization app.
- **One-way sync** (iCal) - allows you to view, but not update, Team Calendars events in Google Calendar.

Subscribe with two-way synchronization (CalDAV)

Two-way synchronization allows you to view and update Team Calendars events in Google Calendar on Android. It requires a CalDav synchronization app. We've tested Team Calendars with DAVdroid, an open source app which is available from the Play Store (paid) or via F-Droid (free).

These instructions are for Android 7.0 Nougat, with DAVdroid 1.10. Your version may differ slightly.

1. Install the DAVdroid app on your device

Before you begin install a CalDAV synchronization app for Android.

We used DAVdroid which you can install from thePlay Store(paid) or viaF-Droid(free). Other plugins may be available.

2. Grab your Team Calendars URL

In Confluence:

1. Choose the **Subscribe** button at the top of your calendar.
2. Select**Google Calendar (Android)**from the**Calendar app**dropdown.

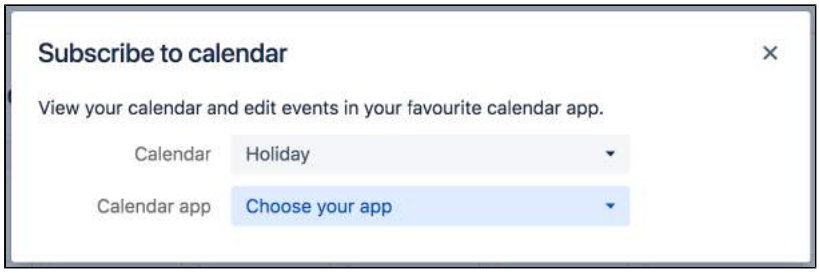

3. If prompted, add the calendar to **My Calendars**.
4. Make a note of the**Calendar URL**.
   Because DAVdroid supports calendar discovery, this will be your Confluence URL, not the URL of an individual calendar. You'll be able to choose which calendars (that you've added to My Calendars) to sync in the app.

3. Subscribe to the calendar in the DAVdroid app

On your Android device:

1. In DAVdroid, go to**Add account**>**Login with URL and user name**.
2. Enter the **Confluence URL** you copied earlier.

3. Enter your Confluence **username** and **password**.

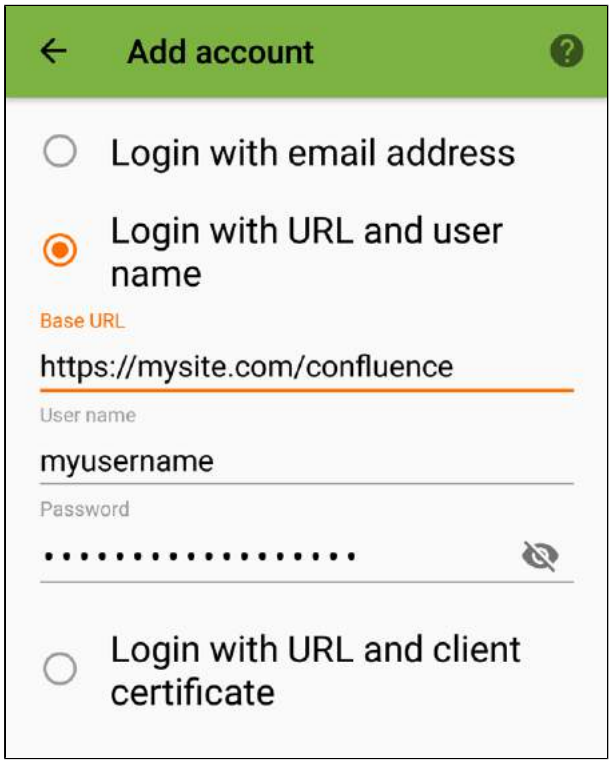

4. Follow the prompts to name your account.
5. Select the calendars you want to display in the Android Calendar app.

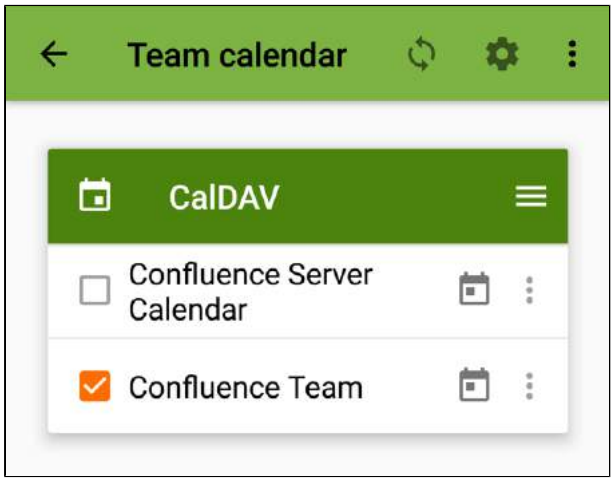

6. Open the Calendar app on your device.Your Team Calendars will be listed under the account name.

Subscribe with one-way synchronization (iCal format)

One-way synchronization means that you can view, but not update, Team Calendars events in your app.

You can't subscribe to an iCal calendar directly on your device. Instead you will need toSubscribe to Team Calendars from Google Calendar in your browser. Events from your Team Calendar will then be visible in the Android calendar app automatically.

Team Calendars for Confluence is now part of Confluence Data Center
To get access to the features described on this page upgrade to Confluence Data Center 7.11 or later. Cant upgrade yet? Depending on your current Data Center version, you can access these features by installing the latest version of the app (at no cost). See our FAQ for all the details

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Subscribe to Team Calendars from Thunderbird

If you want your Team Calendars events to appear in Thunderbird, you can subscribe to Team Calendars in a few easy steps.

On this page:

- Subscribe to Team Calendars from Thunderbird
  - Subscribe with two-way synchronization (CalDAV)
  - Subscribe with one-way synchronization (iCal)

Subscribe to Team Calendars from Thunderbird

There are two ways to synchronize your calendar in Thunderbird:

- **Two-way sync** (CalDAV) - allows you to view and update Team Calendars events in Thunderbird.
- **One-way sync** (iCal) - allows you to view, but not update, Team Calendars events in Thunderbird.

Subscribe with two-way synchronization (CalDAV)

Two-way synchronization allows you to view and update Team Calendars events in iOS.

These instructions are for Thunderbird 52 on MacOS High Sierra. Your version may differ slightly.

1. Grab your Team Calendars URL

In Confluence:

1. Choose the **Subscribe** button at the top of your calendar.
2. Select**Thunderbird**from the**Calendar app**dropdown.

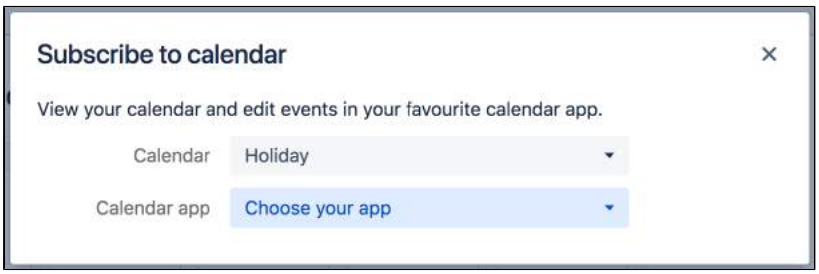

3. Copy the calendar URL.

2. Add the calendar in Thunderbird

In Thunderbird:

1. From the **Home** tab, choose **Create a new calendar**> **On the network**.
2. Choose**CalDAV**.

3. Paste the **Calendar URL** you copied earlier in the **Location** field.

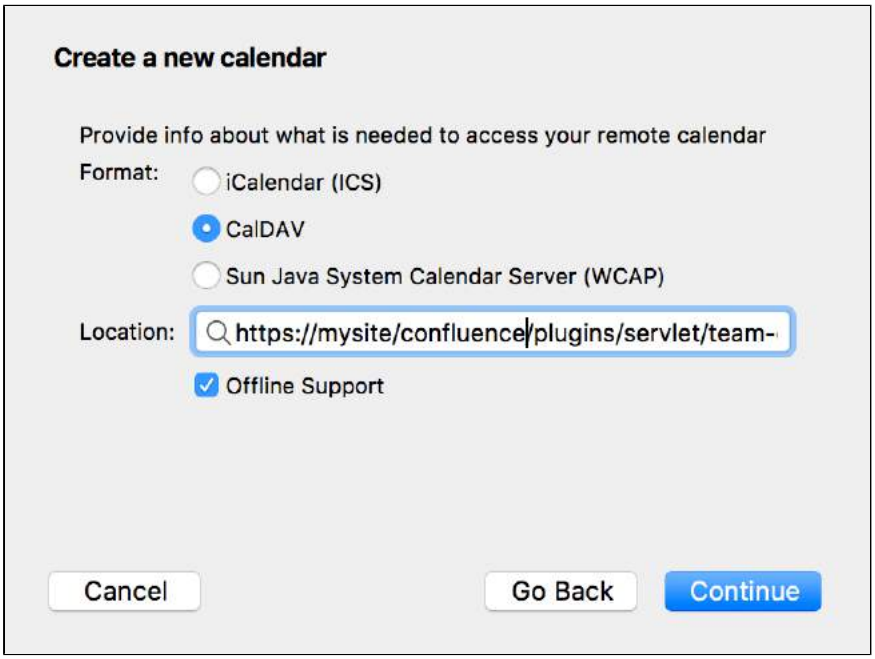

4. Name the calendar - this is the name that will appear in the Thunderbird calendar list.
5. Enter your Confluence **username** and **password**.

Repeat this process for any other Team Calendars you want to add.

Subscribe with one-way synchronization (iCal)

These instructions are for Thunderbird 52 on MacOS High Sierra. Your version may differ slightly.

1. Grab your Team Calendars URL

In Confluence:

1. Choose the **Subscribe** button at the top of your calendar.
2. Select**iCal**from the**Calendar app**dropdown (don't select Thunderbird for one-way sync).
3. Copy the calendar URL.

2. Add the calendar in Thunderbird

In Thunderbird:

1. From the**Home**tab, choose**Create a new calendar**>**On the network**.
2. Choose**iCalendar (ICS)**.
3. Paste the**Calendar URL**you copied earlier in the**Location**field.
4. Name the calendar.This is the name that will appear in the Thunderbird calendar list.
5. Enter your Confluence**username**and**password**.

You can now view, but not edit, Team Calendar events in Thunderbird.

Team Calendars for Confluence is now part of Confluence Data Center
To get access to the features described on this page upgrade to Confluence Data Center 7.11 or later. Cant upgrade yet? Depending on your current Data Center version, you can access these features by installing the latest version of the app (at no cost). See our FAQ for all the details

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Subscribe to Third-Party Calendars from Team Calendars

Subscribing to your favourite third-party calendar app from Team Calendars is a great way to view important events in Confluence.

There is only one way to synchronize a third-party calendar with Team Calendars:

- One-way sync (iCal) - allows you to view, but not update, third-party calendar events in Team Calendars.

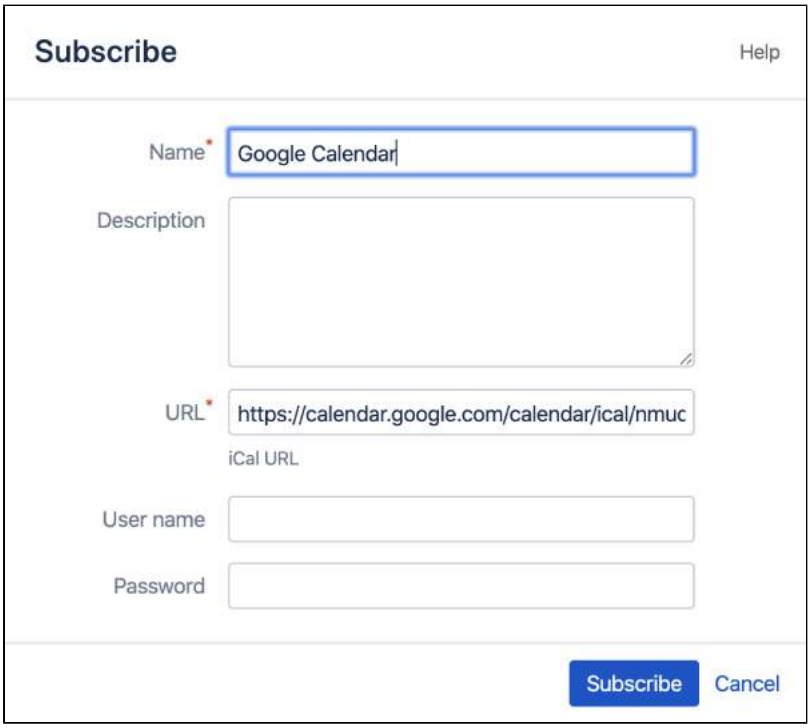

You can subscribe to these third-party calendars from Team Calendars:

- Outlook
- Apple
- Google
- Opsgenie
- PagerDuty
- Teamup

You can alsoSubscribe to Team Calendars from Third-Party Calendars.

Team Calendars for Confluence is now part of Confluence Data Center
To get access to the features described on this page upgrade to Confluence Data Center 7.11 or later. Cant upgrade yet? Depending on your current Data Center version, you can access these features by installing the latest version of the app (at no cost). See our FAQ for all the details

Subscribe to Outlook Calendar from Team Calendars

Subscribing to Outlook Calendar in Team Calendars is a great way to see events, appointments, and meetings in Confluence.

Subscribe to Outlook Calendar from Team Calendars

In your Outlook.com account:

1. Go to **Settings > View all Outlook settings.**
2. Select **Calendar > Shared Calendars.**
3. Under **Publish a Calendar, Select a calendar, Select permissions,** then **Publish.**
4. Copy the **ICS** link.

In Confluence:

1. Select **Calendars** from the Confluence header or space sidebar (to add the calendar in a space).
2. Select ••• next to **Add Calendar** and choose **Subscribe By URL.**
3. Enter the **Name** of the calendar and the ICS link you copied. If necessary, enter your username and password credentials for the calendar. Then select **Subscribe**.

Team Calendars for Confluence is now part of Confluence Data Center
To get access to the features described on this page upgrade to Confluence Data Center 7.11 or later. Cant upgrade yet? Depending on your current Data Center version, you can access these features by installing the latest version of the app (at no cost). See our FAQ for all the details

Subscribe to Apple Calendar from Team Calendars

You can view iCloud calendar events from Team Calendars in a few easy steps.

Subscribe to Apple Calendar from Team Calendars

Set up iCloud on your Mac to subscribe to iCloud calendars. In Apple Calendar on your Mac:

1. From the **Calendars** list, select a calendar under **iCloud.**
2. Select the **Share Calendar button**
3. Copy the **URL** (it will only appear if Public Calendar is selected).

In Confluence:

1. Select**Calendars**from the Confluence header or space sidebar (to add the calendar in a space).
2. Select••• next to **Add Calendar**and choose**Subscribe By URL.**
3. Enter the **Name** of the calendar and the URL you copied. You will need to change webcal to http or https at the beginning of the link. If necessary, enter your username and password credentials for the calendar. Then select **Subscribe**.

Team Calendars for Confluence is now part of Confluence Data Center
To get access to the features described on this page upgrade to Confluence Data Center 7.11 or later. Cant upgrade yet? Depending on your current Data Center version, you can access these features by installing the latest version of the app (at no cost). See our FAQ for all the details

Subscribe to Google Calendar from Team Calendars

Subscribing to Google Calendar in Team Calendars is a convenient way to see all your events in one place, meaning you don't need to leave Confluence just to check your appointments.

Subscribe to Google Calendar from Team Calendars

In Google Calendar:

1. From **My Calendars,** hover over a calendar and select **Options** (three vertical dots) > **Settings and sharing**.
2. From **Settings for my Calendars,** select **Integrate Calendar**.
3. **Copy the URL**listed under**Secret address in iCal format**. If your calendar is made available to the public, you can copy the URL listed under**Public address in iCal format**.

In Confluence:

1. Select**Calendars**from the Confluence header or space sidebar (to add the calendar in a space).
2. Select••• next to**Add Calendar**and choose**Subscribe By URL.**
3. Enter the**Name**of the calendar and the calendar URL you copied. If necessary, enter your username and password credentials for the calendar. Then select**Subscribe**.

Events from your Google Calendar will now appear in Team Calendars.

- Team Calendar reads Google's calendars and caches them for an hour by default. Currently, there's no way to configure the synchronization frequency or interval via the UI. The feature request is tracked here:
  CONFSERVER-50112 - Make Google Calendar to Team Calendars synchronisation frequency or interval to be configurable GATHERING INTEREST
  and there's a workaround available in the report.
- Subscriptions are read-only (you can't modify events from Team Calendars)
  ~~CONFSERVER-51340~~ - CalDAV support CLOSED

Team Calendars for Confluence is now part of Confluence Data Center
To get access to the features described on this page upgrade to Confluence Data Center 7.11 or later. Cant upgrade yet? Depending on your current Data Center version, you can access these features by installing the latest version of the app (at no cost). See our FAQ for all the details

Subscribe to Opsgenie Calendars from Team Calendars

Subscribing to Opsgenie Calendars in Team Calendarsis a great way to see Your on-call schedules and Who is on-call schedules in Confluence.

Subscribe to Your on-call schedules from Team Calendars

In Opsgenie:

1. Go to **Settings** > **Your on-call schedule.**
2. Select the**Open calendar**link icon (it will appear after you hover over the calendar).
3. Select the**Copy calendar link to clipboard**link icon (it will appear to the left of the Open calendar icon).

In Confluence:

1. Select**Calendars**from the Confluence header or space sidebar (to add the calendar in a space).
2. Select••• next to**Add Calendar**and choose**Subscribe By URL.**
3. Enter the**Name**of the calendar and the calendar link you copied. If necessary, enter your username and password credentials for the calendar. Then select**Subscribe**.

Subscribe to Who is on-call schedules from Team Calendars

In Opsgenie:

1. Select **Who is on-call** from the header.
2. Select a schedule name to view the schedule.
3. Select the**Open calendar**link icon (it will appear after you hover over the calendar).
4. Select the**Copy calendar link to clipboard**link icon (it will appear to the left of the Open calendar icon).

In Confluence:

1. Select**Calendars**from the Confluence header or space sidebar (to add the calendar in a space).
2. Select••• next to **Add Calendar**and choose**Subscribe By URL.**
3. Enter the **Name** of the calendar and the calendar link you copied. You will need to change webcal to http or https at the beginning of the link. If necessary, enter your username and password credentials for the calendar. Then select **Subscribe**.

Team Calendars for Confluence is now part of Confluence Data Center
To get access to the features described on this page upgrade to Confluence Data Center 7.11 or later. Cant upgrade yet? Depending on your current Data Center version, you can access these features by installing the latest version of the app (at no cost). See our FAQ for all the details

Subscribe to PagerDuty Schedules from Team Calendars

Subscribing to PagerDuty schedules in Team Calendars is an easy way to see all your on-call shifts in Confluence.

Subscribe to All Schedulesfrom Team Calendars

In PagerDuty:

1. Go to**My Profile**in the drop-down menu from your avatar.
2. Select the**User Settings**tab.
3. Under**Calendar**settings, select**WebCal feed**and copy the link.

In Confluence:

1. Select**Calendars**from the Confluence header or space sidebar (to add the calendar in a space).
2. Select••• next to**Add Calendar**and choose**Subscribe By URL.**
3. Enter the**Name**of the calendar and the WebCal feed link you copied. You will need to change webcal to http or https at the beginning of the link.If necessary, enter your username and password credentials for the calendar. Then select**Subscribe**.

All of your personal on-call schedules from PagerDutywill now appear in Team Calendars.

Subscribe to a Single Schedule from Team Calendars

To subscribe to only your on-call shifts, in PagerDuty:

1. Go to**Configuration > Schedules**.
2. Next to a calendar select**Export > Just My Calendar**and copy the calendar URL.

In Confluence:

1. Select**Calendars**from the Confluence header or space sidebar (to add the calendar in a space).
2. Select••• next to**Add Calendar**and choose**Subscribe By URL.**
3. Enter the**Name**of the calendar and the calendar URL you copied. You will need to change webcal to http or https at the beginning of the URL.If necessary, enter your username and password credentials for the calendar.Then select**Subscribe**.

To subscribe to an entire Schedule (everyone's shifts), in PagerDuty:

1. Go to**Configuration > Schedules**.
2. Next to a calendar select**Export > Everyone**and copy the calendar URL.

In Confluence: same as instructions above.

Team Calendars for Confluence is now part of Confluence Data Center
To get access to the features described on this page upgrade to Confluence Data Center 7.11 or later. Cant upgrade yet? Depending on your current Data Center version, you can access these features by installing the latest version of the app (at no cost). See our FAQ for all the details

Subscribe to Teamup Calendars from Team Calendars

Subscribing to Teamup Calendars in Team Calendars is an efficient way tohelp groups manage shared time and resources.

Subscribe to Teamup Calendars from Team Calendars

In Teamup:

1. Select in the top right and go to**Preferences > iCalendar feeds**.
2. Find your desired calendar on the list or scroll to the bottom of the list for**All Calendars**. Select**Copy**to the right of the feed URL.

In Confluence:

1. Select**Calendars**from the Confluence header or space sidebar (to add the calendar in a space).
2. Select next to**Add Calendar**and choose**Subscribe By URL.**
3. Enter the**Name**of the calendar and the feed URL you copied.If necessary, enter your username and password credentials for the calendar. Then select**Subscribe**.

Team Calendars for Confluence is now part of Confluence Data Center
To get access to the features described on this page upgrade to Confluence Data Center 7.11 or later. Cant upgrade yet? Depending on your current Data Center version, you can access these features by installing the latest version of the app (at no cost). See our FAQ for all the details

Export Team Calendars Content to Other Calendars

Team Calendars allows you to quickly export calendar data so you can open it in other calendar apps like Google Calendar, Apple Calendar, Microsoft Outlook, or any application that supports importing .ics files.

On this page:

- Google Calendar
- Outlook
- Apple Calendar
- Another team calendar
- Using another calendar app?

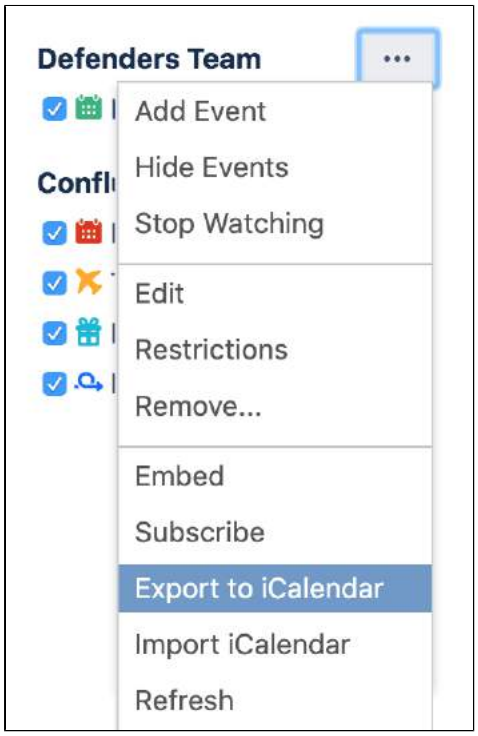

Google Calendar

These instructions are for Google G Suite Calendar (as at March 2018). Your version may differ slightly.

1. Choose ••• to the right of the calendar and choose**Export to iCalendar**
2. Save the .ics file to your computer
3. Open Google Calendar, and choose**Add other calendars**>**Import**.
4. Locate and select the .ics file you saved and follow the prompts to import it.

Outlook

These instructions are for Outlook 2016, your version may differ.

1. Choose ••• to the right of the calendar and choose**Export to iCalendar**
2. Save the .ics file to your computer
3. Open Outlook and choose**File**>**Open & Export** > **Import / Export**
4. Choose **Import an iCalendar (.ics) or vCalendar file (.vcs)**
5. Locate and select the .ics file you saved and follow the prompts to import it.

Apple Calendar

These instructions are for Apple Calendar 10, your version may differ.

1. Choose ••• to the right of the calendar and choose**Export to iCalendar**
2. Save the .ics file to your computer
3. Open the Calendar app on your Mac and choose**File**>**Import**
4. Locate and select the .ics file you saved and choose**Import**
5. Choose the calendar you want to import the events into or choose **New Calendar**, then choose**OK**

Another team calendar

You can also choose to import events from one team calendar into another team calendar.

1. Choose ••• to the right of the calendar and choose **Export to iCalendar**
2. Save the .ics file to your computer
3. Choose ••• to the right of the destination calendar and choose **Import iCalendar**
4. Locate and select the .ics file you saved and follow the prompts to import it.

This process can be used to merge multiple team calendars into one calendar. Once you've confirmed the events were imported successfully, you can delete the original calendar.

Using another calendar app?

Many calendar apps support importing iCal files. Check the documentation for your app to find out how to import.

The information you import into your chosen calendar app is no longer connected to Team Calendars and won't be updated when you update Team Calendars. If you'd like your calendar app to be synchronised with Team Calendars, see Subscribe to Team Calendars from Third-Party Calendars.

Team Calendars for Confluence is now part of Confluence Data Center
To get access to the features described on this page upgrade to Confluence Data Center 7.11 or later. Cant upgrade yet? Depending on your current Data Center version, you can access these features by installing the latest version of the app (at no cost). See our FAQ for all the details

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Delete or Remove a Calendar

If you don't want to see a calendar in a space or on your My Calendars page any more, there are two ways you can remove it. How you proceed depends upon whether you're just hiding the calendar from your view, or deleting it completely.

On this page:

- Remove a Calendar
- Delete a Calendar

Remove a Calendar

To remove a Calendar from a space or your My Calendars page:

1. Find the calendar in a space or on your My Calendars page
2. Choose ••• to the right of the calendar name and select **Remove...**
   You'll be asked if you want to delete the calendar or remove it from your list.
3. Choose the appropriate option:
   - In a space **Remove this calendar from space calendar view**
   - In your My Calendars page **Remove this calendar from my calendar view**

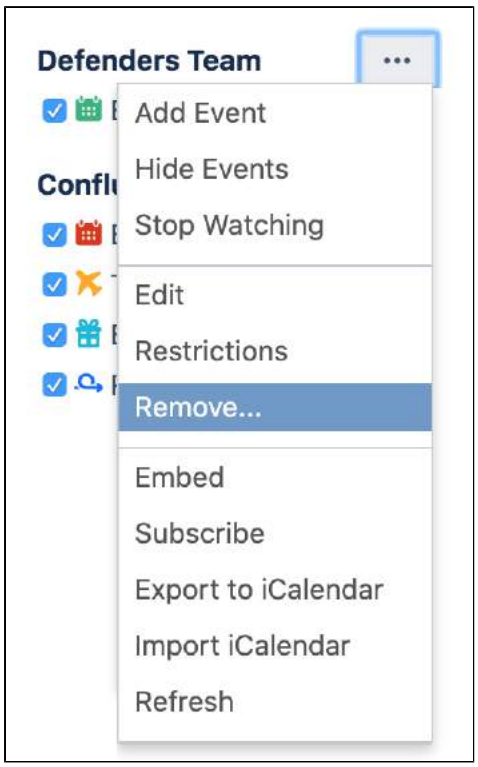

Delete a Calendar

⚠ Be Careful! This action is permanent, and will delete the calendar for all users.

To permanently delete a Calendar for all users:

1. Find the calendar in a space or on your My Calendars page
2. Choose ••• to the right of the calendar name and select **Remove...**
3. You'll be asked if you want to delete the calendar or remove it from your list.
4. Select **Permanently delete this calendar for all users**

If you're using Confluence Cloud, and you'd like Atlassian Cloud site admins to be able to delete calendars, in addition to the calendar's owner, you can enable that option in your Team Calendars settings. Type **Team Calendars** in the Confluence quick search to get to your Team Calendars settings, then tick **Allow site administrators to manage calendars**.

Team Calendars for Confluence is now part of Confluence Data Center
To get access to the features described on this page upgrade to Confluence Data Center 7.11 or later. Cant upgrade yet? Depending on your current Data Center version, you can access these features by installing the latest version of the app (at no cost). See our FAQ for all the details

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Team Calendars FAQ

Get answers to your Team Calendars for Confluence questions. Click any of the questions below to view the answer.

Team Calendars is part of Confluence Data Center. Upgrade to Confluence Data Center 7.11 or later to use Team Calendars.

Yes you can. Learn more about this here:

- Subscribe to Team Calendars from Google Calendar
- Subscribe to Google Calendar from Team Calendars

Yes you can! SeeAdd Jira Events to find out more.

Maybe! It depends whether you have a calendar app that supports iCal or CalDAV. See Subscribe to Team Calendars from Third-Party Calendars to find out how to get set up.

Yes you can. Just add multiple names in the people field.

Team Calendars is available in English and many of the the bundled Confluence languages. See Choosing a Default Language for a list of languages available in Confluence.

The date format isdependenton your language settings and automatically adjusts accordingly - this is a Confluence user and global setting. An admin can set the site default language (seeChoosing a Default Language) and you can also set your preference (seeEdit Your User Settings).

Team Calendars uses the language set in each user's profile to determine which day of the week should be the first day in the calendar. For example, if a user has selected US English, Sunday will be the first day of the week. If they select UK English, Monday will be the first day of the week.

If you want to override the default for all users, regardless of their profile setting, go to > **General Configuration** > **Team Calendars** and choose a specific day of the week.

Currently, it's not possible to subscribe to Google Calendar from Team Calendars using **Public URL to this Calendar** from Google. This is a known issue with Google Calendar; it appears that Google Calendar Public URL doesn't contain the data/event that can be read by Team Calendars to display events. You can subscribe to Google Calendar from Team Calendars using these links:

- Secret address in iCal format
- Public address in iCal format (if your calendar is made available to the public)

Events from JIRA and external calendars are cached to help your calendars load more quickly. The cache expiry (time to live) defaults to 10 minutes, and is configurable. Go to **Team Calendars** > **General Settings** and decrease the Cache expiry time if you want these events to update more frequently.

Team Calendars for Confluence is now part of Confluence Data Center
To get access to the features described on this page upgrade to Confluence Data Center 7.11 or later. Cant upgrade yet? Depending on your current Data Center version, you can access these features by installing the latest version of the app (at no cost). See our FAQ for all the details

Add-ons and integrations

Confluence has a wide range of features on its own, but you can also extend those features with Marketplace apps, and by integrating Confluence with other applications. Integrating with Jiraapplicationscan really take your Confluence experience to the next level by improving the way your teams 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 apps. Whether you need to create diagrams, like the ones you can create withGliffy, or you want to make awesome mockups and wireframes with Balsamiq, there are heaps of great apps in the marketplace. You may even find a really useful app you never knew you needed, but now can't live without.

In this section:

- Use Jira applications and Confluence together
- Request Marketplace Apps
- Use a WebDAV Client to Work with Pages
- Mail Archives
- Gadgets

Related pages:

- Use Jira applications and Confluence together
- Use a WebDAV Client to Work with Pages
- Gadgets
- Request Marketplace Apps

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 usingApplication Links,youcan display and create Jira issues and more from within Confluence.

What you can do with Confluence and Jira depends on the Jira application, version, and how it is hosted. Find out about the required applications and versions later in this page.

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
    - Link to pages from Jira
- For software teams
  - Define your requirements
  - Manage your sprints
- For service teams
  - Provide self help resources for your customers
  - Create knowledge base articles
  - Allow any active user to see knowledge base spaces
- Jira applications required

Want to learn how to integrate Confluence Cloud and Jira Software Cloud?Check out this guide.

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.

- `<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 macroand display a list of issues matching the saved filter.
- `<yourjirasite.com>/issues/?jql=project%20%3D%20CONF` will insert the Jira Issues macroand display a list of issues matching the Jira search.

Alternatively, you can add theJira Issues Macroto the page and search for issues directly:

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 customize 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 theJIRA Report blueprintto create a Change Log or Status report.
- Use theJira Chart Macroto display data as a chart, including pie charts, created vs resolved, and two dimensional charts.
- Use JIRAGadgetsto 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**iconthat 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 apopup when you highlight text, check thatText Selectis 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 Macropage.

Move between Jira and Confluence

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

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 itreally easy to jump fromConfluence 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.

The button doesn't detect links from issues displayed in the Jira Issues macro in table format.

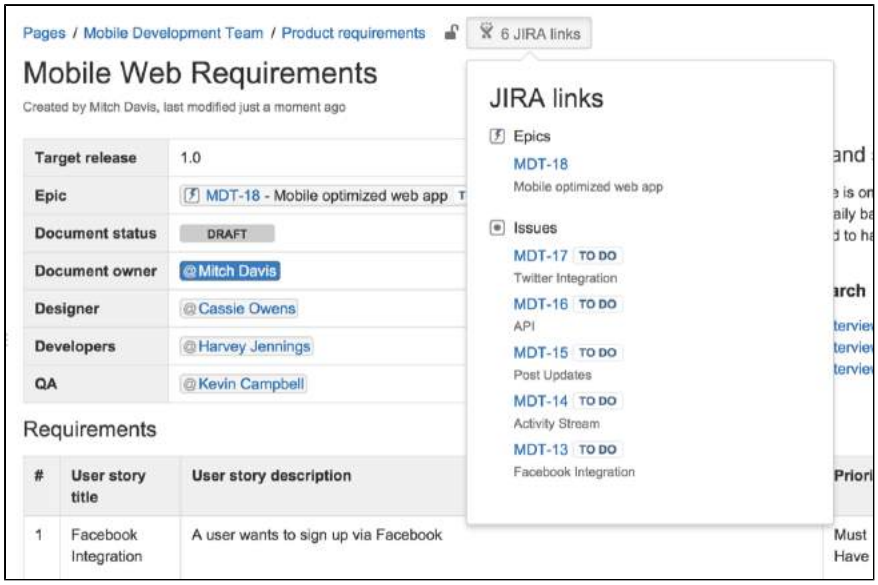

Link to pages from Jira

When viewing an issue in Jira, you can link it to a relevant Confluence page.

How you do this depends on your Jira application, but there will be a Link option that will allow you to search for a Confluence page, or enter a URL.

If you use a Jira cloud application, the option to search for a Confluence Server or Data Center page is only available in sites that do not have Confluence cloud. This is because Jira defaults to search Confluence cloud, over any other linked Jira applications.

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 theProduct Requirements Blueprintto 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 theProduct 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 tocreate 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.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

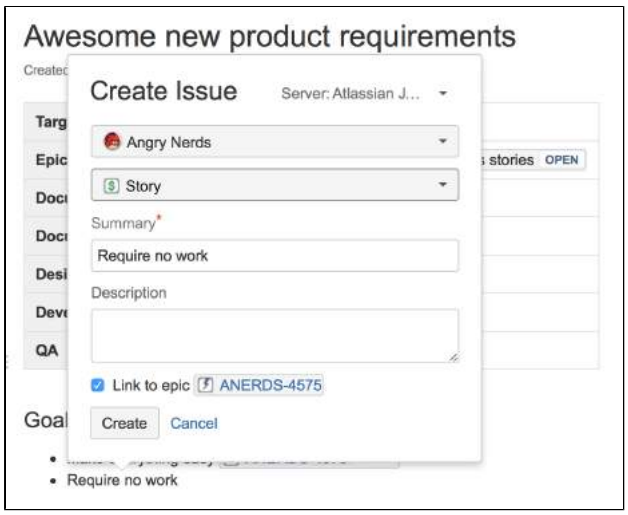

The tight integration between Confluence and Jira Software means you can easily access issues from the Confluence page and see their status at a glance, and from within Jira Software you can see links to related Confluence pages. All the information you need is right there.

If you have Jira Software Cloud, the option to automatically link issues created in Confluence to your epic isonly available in Classic projects.

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 theMeeting Notes Blueprint a handy template that helps capture the details you need andis 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 theJIRA Report blueprintin Confluence.
- Use the Retrospective Blueprintin 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.

If you have Jira Software Cloud, the option to create pages from your sprint is only available in Classic projects.

For service teams

Available for Jira Service Management (formerly Jira Service Desk) Server and Data Center only.If you use Jira Service Management Cloud, you cannot connect your project to a knowledge base space on a Confluence Server or Data Center site. You can only connect to a Confluence Cloud site. Read about how to migrate from Confluence Server to Cloud.

Provide self help resources for your customers

If you use Jira Service Management, you can help your customers resolve their issues without creating a request by connecting your service project to a knowledge base in Confluence. You'll need Confluence 5.10 or later.

In Jira Service Management, head to **Project settings**>**Knowledge base**to connect or create a Confluence space.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

When customers search in the Customer Portal, pages in the linked knowledge base space will be suggested, allowing customers to help themselves.

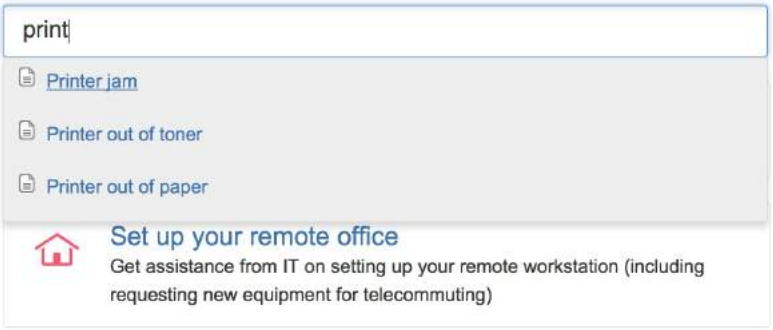

Create knowledge base articles

The Knowledge Base space blueprint, along with templates for how-to and troubleshooting articles make creating new knowledge base articles super simple for your agents.

The templates used in the how-to and troubleshooting blueprints are completely customizabletoo. 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 service 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 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 Management, but you can revoke access to the whole site or to particular spaces via Confluence's global permissions or space 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 needif you're not using the latest Jira Server or Data Center applications.
- The minimum Cloud plan required, and whether there are any limitationson the project type or apps in your site (for example, some integrations won't be available if you also have a ConfluenceCloud site).

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Feature	Jira Core	Jira Software	Jira Service Management (formerly Jira Service Desk)	Minimum legacy version	Jira Cloud plan
Display issues using Jira issues macro	✓	✓	✓	Jira 4.3	✓ Any
Display issue and project information using Jira chart macro	✓	✓	✓	Jira 5.0	✓ Any
Display issue and project information using the Jira Report blueprint	✓	✓	✓	Jira 5.0	✓ Any
Create an issue from Jira issues macro	✓	✓	✓	Jira 4.3	✓ Any
Create issues by highlighting text on a Confluence page	✓	✓	✓	Jira 6.3.1	✓ Any
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	⚠ Classic projects only
Link and create Confluence pages from epics and sprints	✗	✓	✗	Jira 6.3.1 and Jira Agile 6.3.5	⚠ Jira-only sites ⚠ Classic projects only
View linked issues with the Jira links button in Confluence	✓	✓	✓	Jira 6.3.1	✓ Any
Create a space using the Software Project space blueprint	✗	✓	✗	Jira Software 7.0	✓ Any
Use a Confluence space as a knowledge base in the customer portal	✗	✗	✓	Jira 5.2 and Jira Service Desk 1.0	✗
Allow Service Management customers to view knowledge base articles without a Confluence license	✗	✗	✓	Jira Service Desk 3.1	✗
Search for an existing Confluence page within the Jira link dialog	✓	✓	✓	Jira 4.3	⚠ Jira-only sites
Delegate user management to Jira	✓	✓	✓	Jira 4.3	✗
Add a Jira gadget to a Confluence page	✓	✓	✓	Jira 4.3	✓ Any

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 outIntegrating Jira and Confluence.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Request Marketplace Apps

The Atlassian Marketplace website offers hundreds of apps that the administrator of your Atlassian application can install to enhance and extend Confluence. If the app request feature is enabled for your Confluence instance, you can submit requests for appsfrom 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 apps that can help them work. When you find an appof interest, you can submit a request with just a few clicks.

Submit an apprequest

To browse for add-ons in the Atlassian Marketplace, follow these steps:

1. Choose**yourprofile picture**at top right of the screen, then choose**Atlassian Marketplace**.
2. In the Atlassian Marketplace page, use the search box to find apps or use the category menus to browse or filter 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 app 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 app.

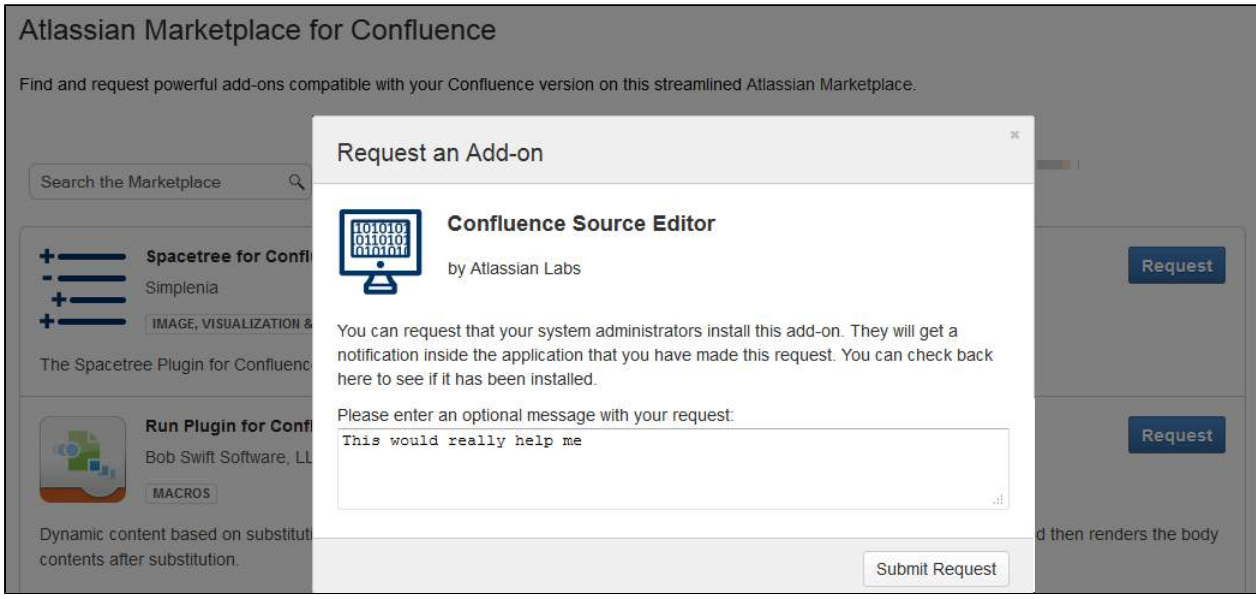

5. When ready, click **Submit Request**.
6. Click **Close** to dismiss the 'Success!' message dialog box.

At this point, a notification appears in the interface your administrators use to administer apps. Also your request message will appear in the app details view, visible from the administrator's 'Find New Apps' page. From there, your administrator can purchase the app, try it out or dismiss requests.

Update an app 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 app immediately.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

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. SeeConfiguring a WebDAV client for Confluencefor 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 organized.

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.

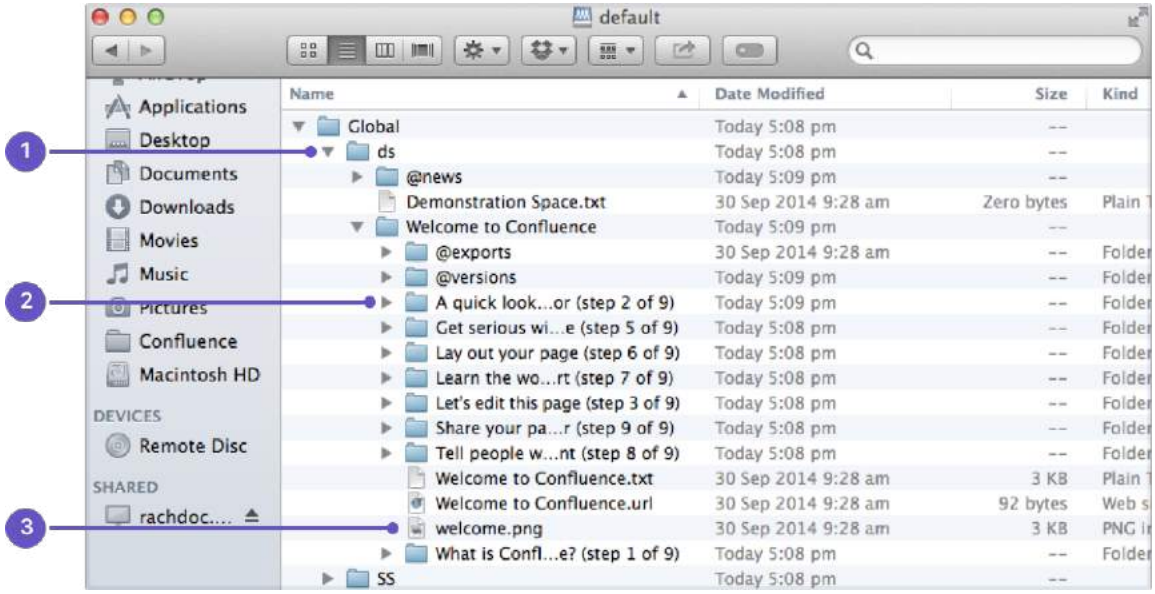

1. **Space key**
2. **Page title**
3. **Attached file**

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).

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

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 system app. To activate mail archiving features in Confluence, enable the app go to > **Manage apps** 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**
- 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
- Create content with email

Screenshot: Viewing a message in the mail archive

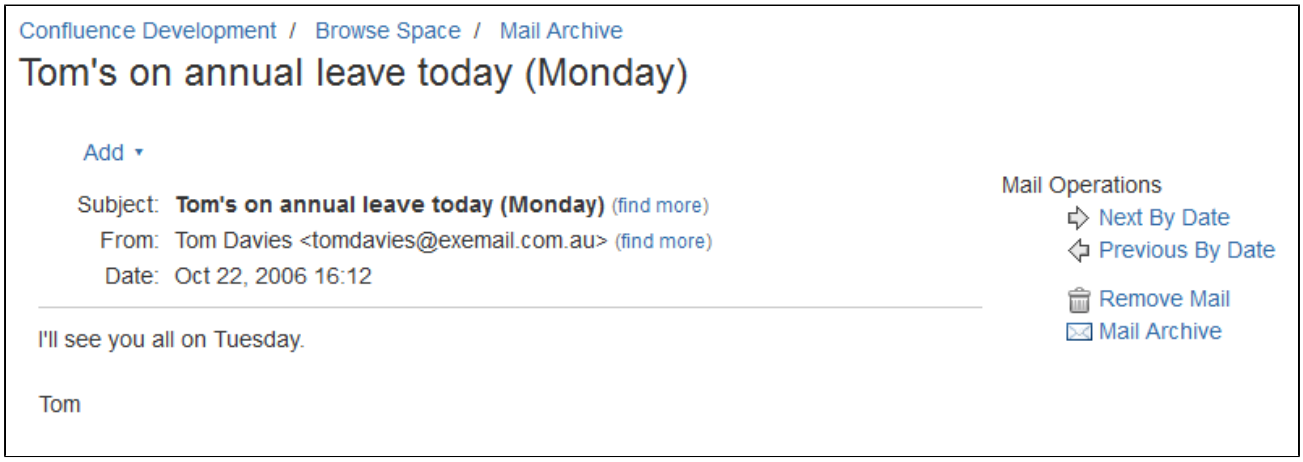

Notes

- 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 system app. To activate mail archiving features in Confluence, enable the app go to > **Manage apps** 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 Space 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 `conf-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 to 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 receivedemail messages to the clone account.

Step 2. Configure Confluence to archive the clone account

1. Go to the space and choose**Space tools**>**Integrations**from the bottom of the sidebar.
2. Choose **Mail Accounts** > **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.
    - **Authentication**:By default, this is set to basic authentication. To select a different authentication method like OAuth 2.0, you'll need to first configure an outgoing linkwith your service provider.
    - **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

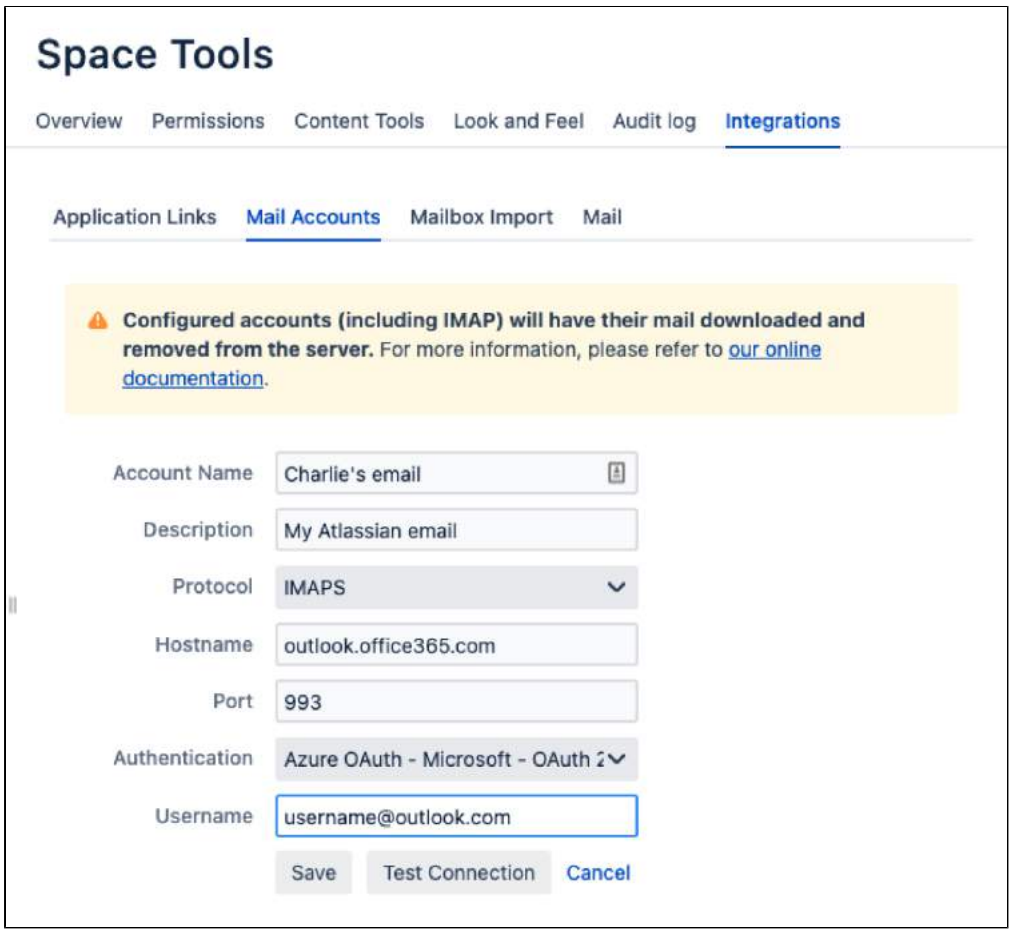

Screenshot: Adding a mail account.

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**.

You need to be aspace administratorto manually retrieve mail. SeeSpace 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 system app. To activate mail archiving features in Confluence, enable the app go to > **Manage apps** 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.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Delete and Restore Mail

To delete mail from a space, you need 'Delete Mail' permission.

Only a space administrator can delete all email messages in the space simultaneously.

To delete mail from a space:

1. Go to a space and choose **Space tools** > **Integrations** > **Mail**
   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 **Delete 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.

To restore deleted mail:

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.

Related pages:

- Mail Archives
- Add a Mail Account

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.

Related pages:

- Mail Archives
- Add a Mail Account

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**.

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 system app. To activate mail archiving features in Confluence, enable the app go to > **Manage apps** 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. SeeConfiguring System Properties.

Create content with email

Use email to create new pages and reply to comments even when you can't access your Confluence instance. Whether you're offsite, on mobile, or have lost network connection to your instance, these apps allow you to continue collaborating and creating content with flexibility and less disruption:

- Confluence Email Gateway
- Confluence Email to Page
- Confluence Reply to Email

These apps are considered experimental and only allow you to create content within your personal space. They are disabled by default with your Confluence installation.

On this page, youll learn how to configure each of these apps and what they do.

Before you begin, you will need to:

- configure an outgoing link to allow Confluence to be an OAuth client
- configure a Server for Outgoing Mail to allow Confluence to send email messages
- configure a Server for Incoming Mail to allow Confluence to receive email messages

Email Gateway

The Email Gateway app allows users to create content by sending an email to Confluence.

To enable the app:

You need to be a Confluence admin to perform these steps.

1. Go to > **Manage apps**
2. Change the filter to **All apps**
3. Search for Email
4. In the results, expand **Confluence Email Gateway**
5. Select **Enable**

Screenshot: Confluence admin email gateway configuration

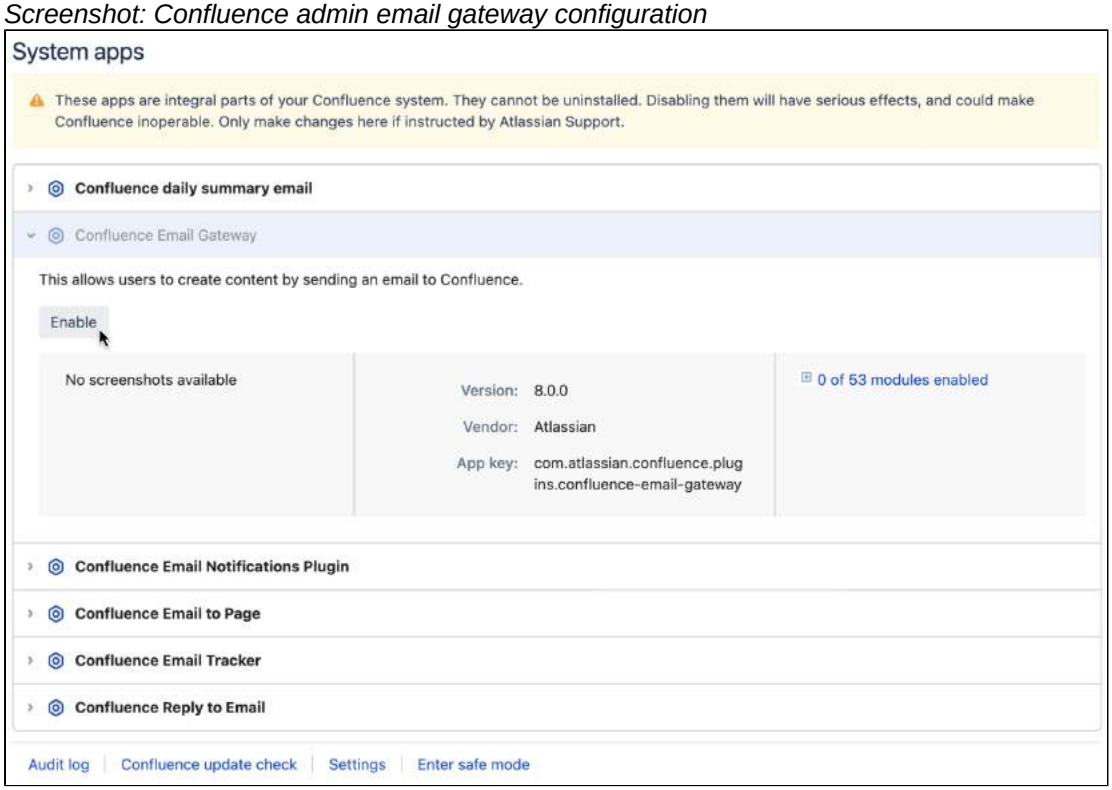

Screenshot: Create by email configuration screen. Once the Email Gateway app is enabled, you can access "Create by email" in the Confluence admin side panel.

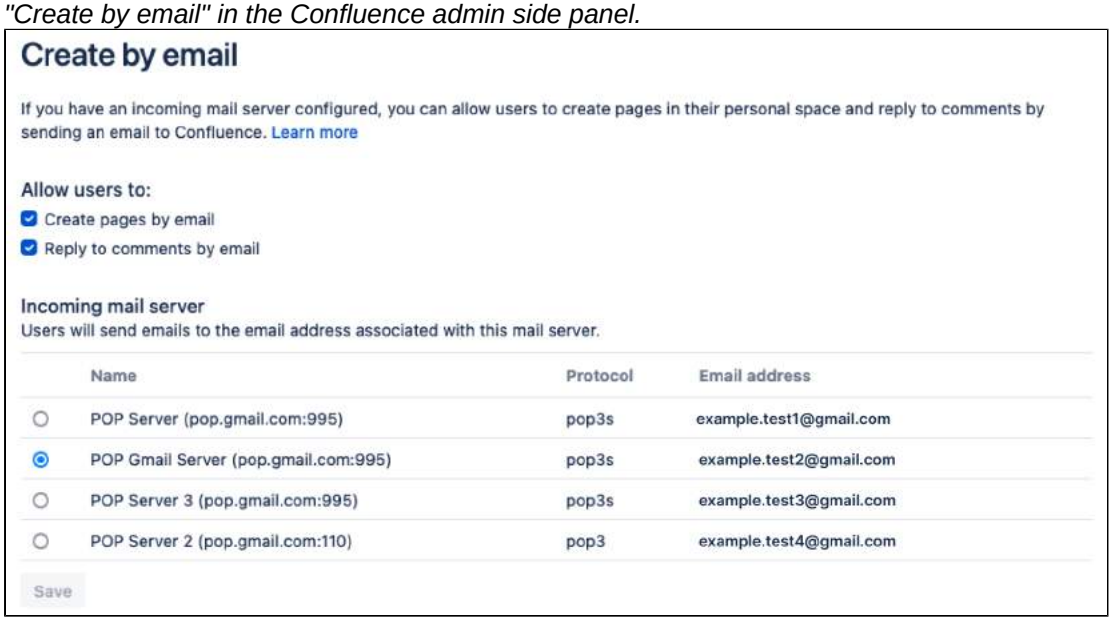

Email to page

The Email to Page app allows users to create pages in their personal space by sending emails to a configured email address.

To enable the app:

You need to be a Confluence admin to perform these steps.

Ensure you have enabled the Email Gateway for this to work. See above

1. Go to ⚙ > **Manage apps**
2. Set the filter to **All apps**
3. Search for 'Email'
4. In the results, expand **Confluence Email to Page**
5. Select **Enable**
6. Go to ⚙> **General Configuration**> **Create by page**
7. Check the box next to **Create pages by email**
8. Select the incoming mail server (this is the mail account users will email to create a page)
9. Select **Save**

Check yourspace:

Users will be able to see the app has been enabled in their personal space (see screenshot below).

1. Go to your personal space
2. Select **Space tools** > **Integrations**
3. Select the tab **Create By Email**
4. You should see **Create pages by email** listed with details of the inbox thats been configured. Email this address with the content of a new page, and youll receive an email reply with a Create page link to create the page in your personal space.

Reply to Email

The Reply to Email app allows users to publish comments in their personal space by replying to email notifications.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

To enable the app:

You need to be a Confluence admin to perform these steps.

Ensure you have enabled the Email Gateway app for this to work. See above

Go to > **Manage apps**

1. Set the filter to **All apps**
2. Search for Email
3. In the results, expand **Confluence Reply to Email**
4. Select **Enable**
5. Go to **> General Configuration**> **Create by page**
6. Check the box next to **Reply to comments by email**
7. Select the incoming mail server (this is the mail account users will reply to to create a comment)
8. Select **Save**

Check your space:

Users will be able to see the app has been enabled in their personal space (see screenshot below).

1. Go to your personal space
2. Select **Space tools** > **Integrations**
3. Select the tab **Create By Email**
4. You should see **Reply to comments by email** listed

Screenshot: When (1) Email to Page and (2) Reply to Email apps are enabled, the user will see this in their space integrations menu, and can email admin@companyname.com to create content.

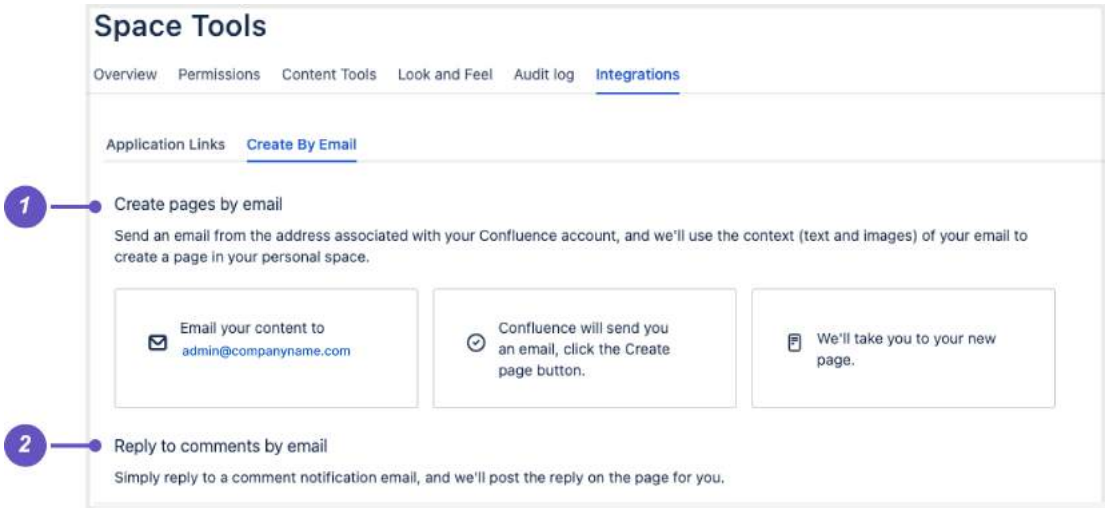

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

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 theintroduction 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**.

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

Add a Confluence gadget to a page

We ended support for the following Confluence gadgets in Confluence 7.0:

- Confluence Page Gadget
- Activity Stream Gadget
- Confluence QuickNav Gadget

These gadgets no longer appear in the macro browser and can't be added to a page. Any gadget already on a page, or used in another application like Jira, will still work.

The Confluence News gadget was removed entirely in Confluence 7.0. This gadget displayed news from Atlassian and hadn't been working for some time.

If you're wanting to display information within Confluence, we recommend using the following macros as an alternative:

- Include Page macro
- Recently Updated macro
- Livesearch macro

Add a Jira gadget to a page

For displaying basic Jira information, such as issues and charts, we recommend using theJira Issues Macroand Jira Chart Macroas 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:

1. In the editor go to**Insert**>**Other Macros**.
2. Select the gadget you wish to add, and use the preview area to configure the gadget.
3. Choose**Insert**.

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 authorizedexternal gadgetsin 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 dashboard:

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.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Activity Stream Gadget

> **We ended support for this gadget in Confluence 7.0**
> The gadget no longer appears in the macro browser and can't be added to a page.
> Any gadget already on a page, or used in another application like Jira, will still work.

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 Page Gadget

We ended support for this gadget in Confluence 7.0
The gadget no longer appears in the macro browser and can't be added to a page.
Any gadget already on a page, or used in another application like Jira, will still work.

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

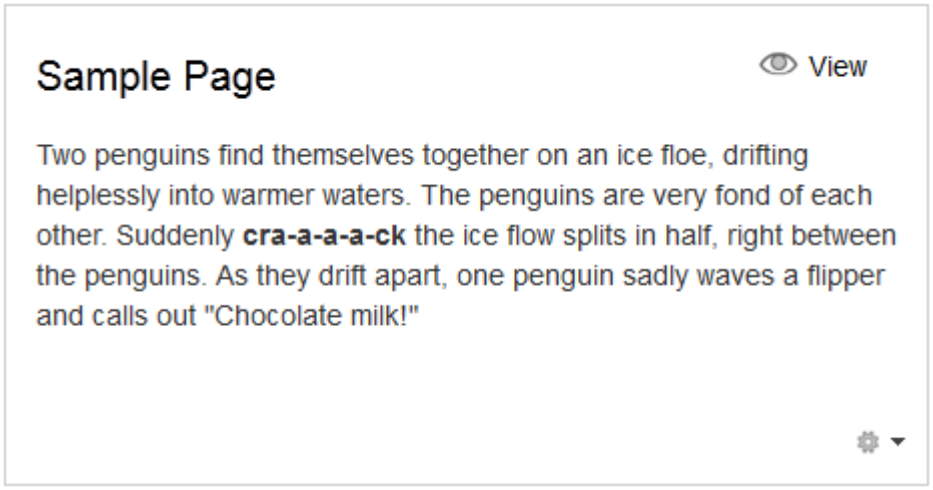

Confluence Page Gadget Properties

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

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

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 Server	✗	See here for the Improvement Request: CONFSERVER-51407 - Make Team Calendars display in a Confluence Page Gadget, for use on a JIRA Dashboard GATHERING INTEREST
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

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Confluence QuickNav Gadget

We ended support for this gadget in Confluence 7.0
The gadget no longer appears in the macro browser and can't be added to a page.
Any gadget already on a page, or used in another application like Jira, will still work.

The QuickNav Gadget allows you to add the quick navigation functionality to search your Confluence site.

To search Confluence using a QuickNav Gadget, type the name of a page, blog post, person, file, or space into the search box, and choose from the list of results displayed.

If you don't immediately see what you need, hit **Enter** or choose the '**Search for'** option at the bottom of the search results to go to the advanced search page. Learn more about searching Confluence.

Screenshot: Using the QuickNav Gadget

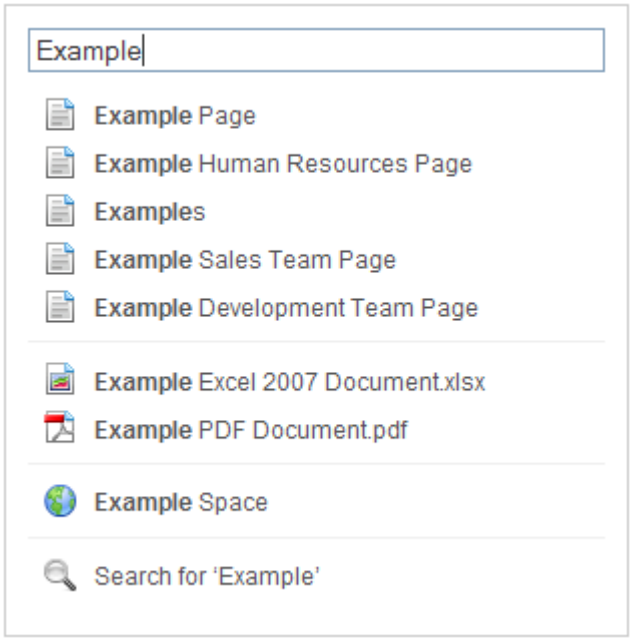

You can also search for administrative options in the QuickNav Gadget. For example, type 'general' into the search field to go to the General Configuration screen.

More information about the QuickNav Gadget:

- The QuickNav Gadget returns matches based on the title only, not the content of the page or file.
- Matching items are grouped by type, so that you can quickly find the type you want. Confluence shows a maximum of 3 admin options, 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.

Confluence QuickNav Gadget Properties

This gadget has no properties and cannot be customized.

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 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. SeeConfluence for Software Teams.

Develop Technical Documentation in Confluence

Confluence is a flexible platform with a range of features and Marketplace apps that can help you capture, distribute, and update yourtechnical 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.

For another great overview of how you can use Confluence for documentation check out Rock the Docs from our solution partner, K15t.

On this page:

- Create your Documentation Space
- Save time by re-using content
- Use page templates
- Draft your work
- Use links and anchors
- Useful macros
- Keep track of page updates
- Customize PDF export
- Other useful tools and apps
- Documentation theme (discontinued)

Related pages:

- Spaces

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 customize the home page at any time; what it looks like is completely up to you!

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.

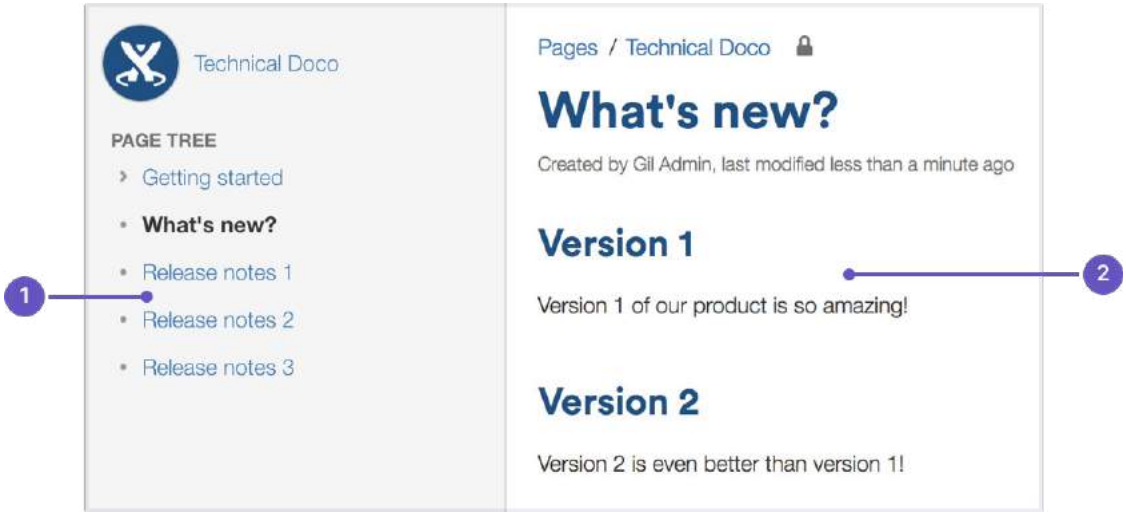

1. **Excerpts:** the intros to these pages are in excerpt macros.
2. **Excerpt include:** these are excerpt include macros.

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 pagewith 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 contentand**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** 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 normalConfluence pages.

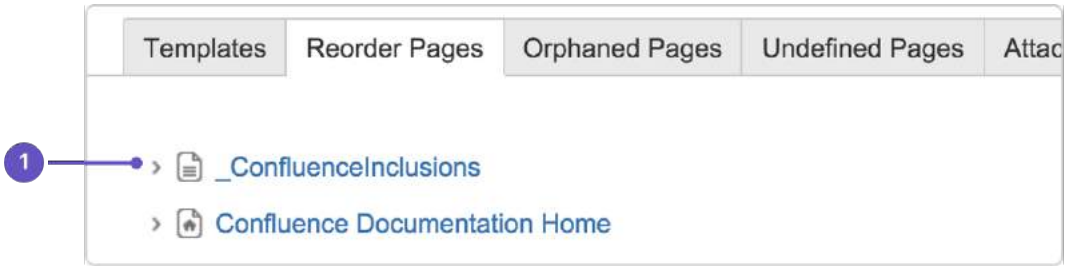

1. **Inclusions library location:** drag your inclusions library here, above the rest of your documentation.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

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 maximize 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 globalpage template** button at the top-right
4. Create your template page and choose**Save**

For detailed info on page templates, seeCreate a Template.

To get to**GlobalTemplates 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 reviewit 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 finalized, 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 itand turn it into a link.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

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 summarizing 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.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

View pagehistory

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 atPage History and Page Comparison Viewsfor 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 customizable list of people who've contributed by creating, editing, or, optionally, commenting on the page.

Customize 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 customize the look of the PDF by adding a title page, header, and footer.

The process you take depends on whether you're trying to customize 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 onCustomize Exports to PDF for more detailed instructions.

Other useful tools and apps

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 apps available on theAtlassian Marketplace,most of which we use ourselves, which canextend the functionality of Confluence.

New apps are hitting the marketplace all the time. This is by no means an exhaustive list!

Scroll Versions (supported by vendor)

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)

TheCopy Space add-on, by Atlassian Labs,does what its name suggests; itallows 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 app is also useful when you need toarchive 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.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

At this point this plugin won't copy page history, blog posts and email.

Scroll PDF Exporter(supported by vendor)

If you're going to produce a PDF of your documentation space, wouldn't you like it to be professionally formatted? TheScroll PDF Exporter, by K15t, lets you style single pages or whole spaces for export, using handy PDF templates.

Gliffy (supported by vendor)

Create diagrams, wireframes, flowcharts and more withGliffy. Gliffy featuresa 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 by vendor)

Lucidchart is available in versions forCloudandServer, and allows you tocreate 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 alsocomes with a free Visio viewer, so you can viewMicrosoft Visio(.vsd) files, Visio stencils (.vss) and it also supports exporting back to Visio.

Documentation theme (discontinued)

Confluence historically provided a theme specifically for documentation. This theme was removed inConfluence 6.0.

A number of the original Documentation theme features, such as headers, footers and the ability to add custom content to the sidebar, are available in Confluence's more modern default theme, which make it a great choice for your documentation space.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

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 organization or team.

Learn more about how a knowledge base helps your team work smarter

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 spaceto be self curating, and do not want to spend a lot of time manually organizing 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 homepagewith 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-organizing 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.

Customize 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 customizing the how-to and troubleshooting article templates to make them relevant for your organization. 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 customizing 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 color scheme (choose **Space Tools** > **Look and Feel** > **ColorScheme**).

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 Management** -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 thecustomer portal. You can connect Jira Service Management with Confluence 5.10 or later.
- **If you use Questions for Confluence Server** - 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 Marketplace apps

The Atlassian Marketplace has a large number of apps 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 app that's right for your knowledge base.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Use Confluence as your Intranet

Your intranet is the hub of your organization. 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 organization.

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 organize 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.Allowpeople toadd themselves as usersof the site; invite people to sign up bysending them an invitation link;add new usersmanually; oruse an existing directory like an LDAP directory for authentication and to manage users and groups.

Whichever way you choose, you canquickly build a community of Confluence users and give them access to your intranet; you'll also have a ready-madepeople directory.

Match your company branding

Upload your company logo, and Confluence's auto look and feel will change the color 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

AConfluence spaceis 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 organize things is to create a space for each team or department within your organization. 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.

Just choose**Spaces**>**Create space**from the header, and Confluence provides a list of space blueprints to help get you started.

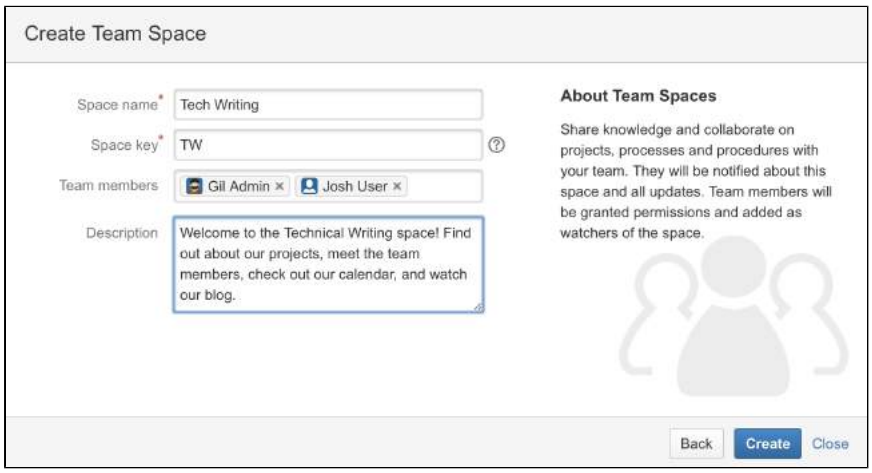

Each space can have its own color scheme and has a customizable home page, whichyou can edit to suit your purpose like displaying and tracking team goals and displaying a list of team members. Use the built-in 'Team Space' template to automaticallyadd all members of the team to the homepage, to help everyone get to know each other.

You canset 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 alsocreate 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 organization (or just those that you want to see them).

Create pages, meeting notes and more

You cancreate pagesfor 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**Cr eate**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-inpage blueprints, which help you with the content and formatting of the page. Themeeting notesanddecisions blueprintsare 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 ownblog, where you and your teams can share news and events, discuss important projects and developments, or congratulate a teammate for aspecial effort; blogging is a great way to foster company culture and celebrate achievements across your organization.

You can watch any blog to make sure you get updated when there's a new post. Blog posts are automatically organized 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 tomentionthem 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 blogsis 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.

TheConfluence dashboardalso has a recent activity feed, which allows you and your team to see what's trending throughout the company or in yournetwork.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

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

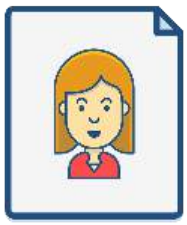

Creating insightful customer interview pages

Create sprint retrospective and demo pages (like a BOSS)

How to make better decisions as a development team

How to document releases 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, SeeSpaces.

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

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

Downloads

Download the Confluence documentation in PDF format.

Other resources

Confluence installation and upgrade guide

Confluence Knowledge Base

Atlassian Answers

    - Working with Confluence Logs
    - Scheduled Jobs
    - Configuring the Allowlist
    - Configuring the Time Interval at which Drafts are Saved
- Configuring Confluence Security
    - Confluence Security Overview and Advisories
    - Proxy and HTTPS setup for Confluence
    - Configuring Secure Administrator Sessions
    - Confluence Cookies
    - Using Fail2Ban to limit login attempts
    - Securing Confluence with Apache
    - Best Practices for Configuring Confluence Security
    - Hiding the People Directory
    - Configuring Captcha for Spam Prevention
    - Hiding External Links From Search Engines
    - Configuring Captcha for Failed Logins
    - Configuring XSRF Protection
    - User Email Visibility
    - Anonymous Access to Remote API
    - Configuring RSS Feeds
    - Preventing and Cleaning Up Spam
- Configuring a Confluence Environment
    - Confluence Home and other important directories
    - Application Server Configuration
    - Starting Confluence Automatically on System Startup
- Performance Tuning
    - Cache Performance Tuning
    - Memory Usage and Requirements
    - Requesting Performance Support
    - Compressing an HTTP Response within Confluence
    - Garbage Collector Performance Issues
    - Troubleshooting Slow Performance Using Page Request Profiling
    - Confluence Diagnostics
    - Faster permissions service
    - Confluence guardrails
- Data Collection Policy
- Managing emojis
- Administering Collaborative Editing
    - Possible Confluence and Synchrony Configurations
    - Configuring Synchrony
    - Set up a Synchrony cluster for Confluence Data Center
    - Migrate from a standalone Synchrony cluster to managed Synchrony
    - Troubleshooting Collaborative Editing
- Using read-only mode for site maintenance
- Administering the Atlassian Companion App
- Notifications from Atlassian
- Administer analytics
- Monitor application performance
    - Monitor Confluence with Prometheus and Grafana

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

- App metrics reference
- Live Monitoring Using the JMX Interface

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Getting Started as Confluence Administrator

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

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

On this page:

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

Diagram: a Confluence installation

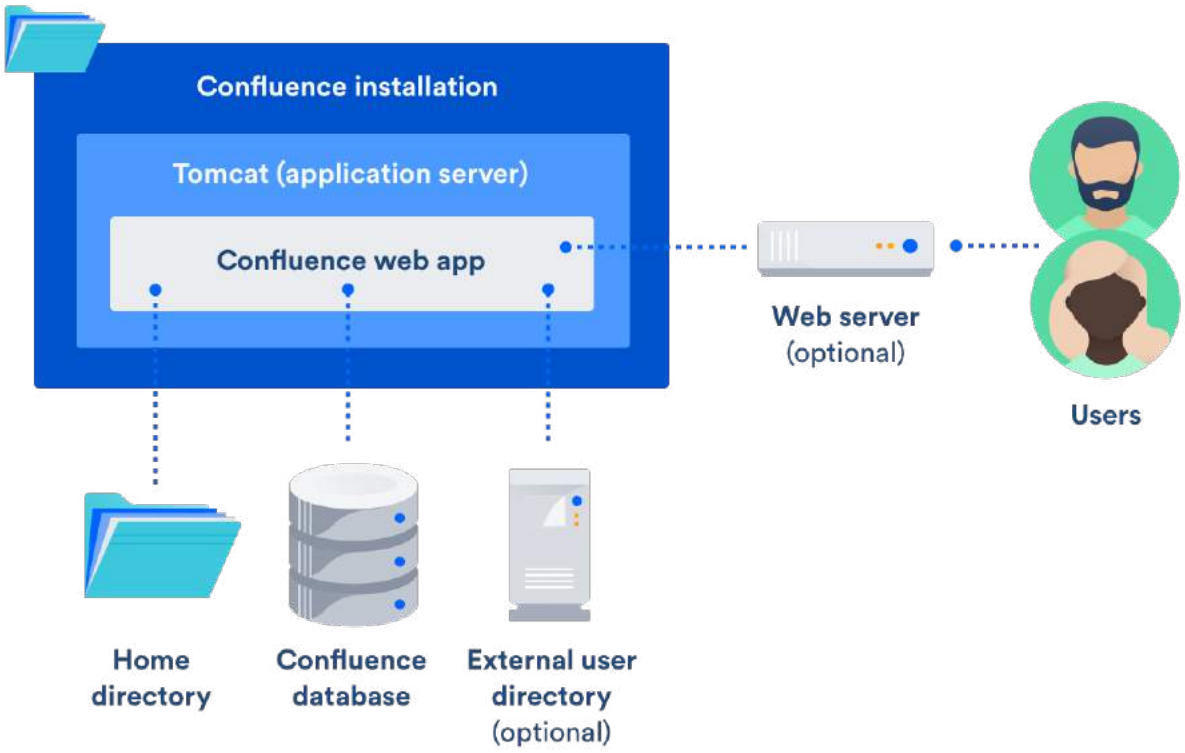

Quick access to admin functions via search

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

Screenshot: searching for admin options

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

Notes about finding admin functions via search:

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

How to administer and configure Confluence

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

To access the Confluence Administration Console:

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

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

Getting started on a new Confluence site

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

- Decide whether you want to allow public (anonymous) access to your site.SeeSetting Up Public Access.
- Add a space and some content. SeeCreate a SpacethenPages and blogs.
- Invite some users to your site. SeeAdd and Invite Users.

- Decide whether you will manage your users in Confluence or hook up an external LDAP directory. See Configuring User Directories.
- Make sure you have set up an email server. The above task list will include this step, but it is worth mentioning it here again. Email notifications are an important part of collaborating on Confluence. See Configuring a Server for Outgoing Mail.

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

Getting to know an existing Confluence site

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

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

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

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Manage Users

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

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

On this page:

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

Related pages:

- Configuring Confluence Security

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

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

Authentication

Seraph

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

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

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

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

XML-RPC and SOAP authentication

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

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

Password authentication

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

SAML single sign-on

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

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

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Add and Invite Users

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

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

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

On this page:

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

Related pages:

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

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 organization can add their own usernames.

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

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

Manageuser signupnotifications

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

To disable this notification:

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

4. Choose**Save**

Screenshot: User signup options

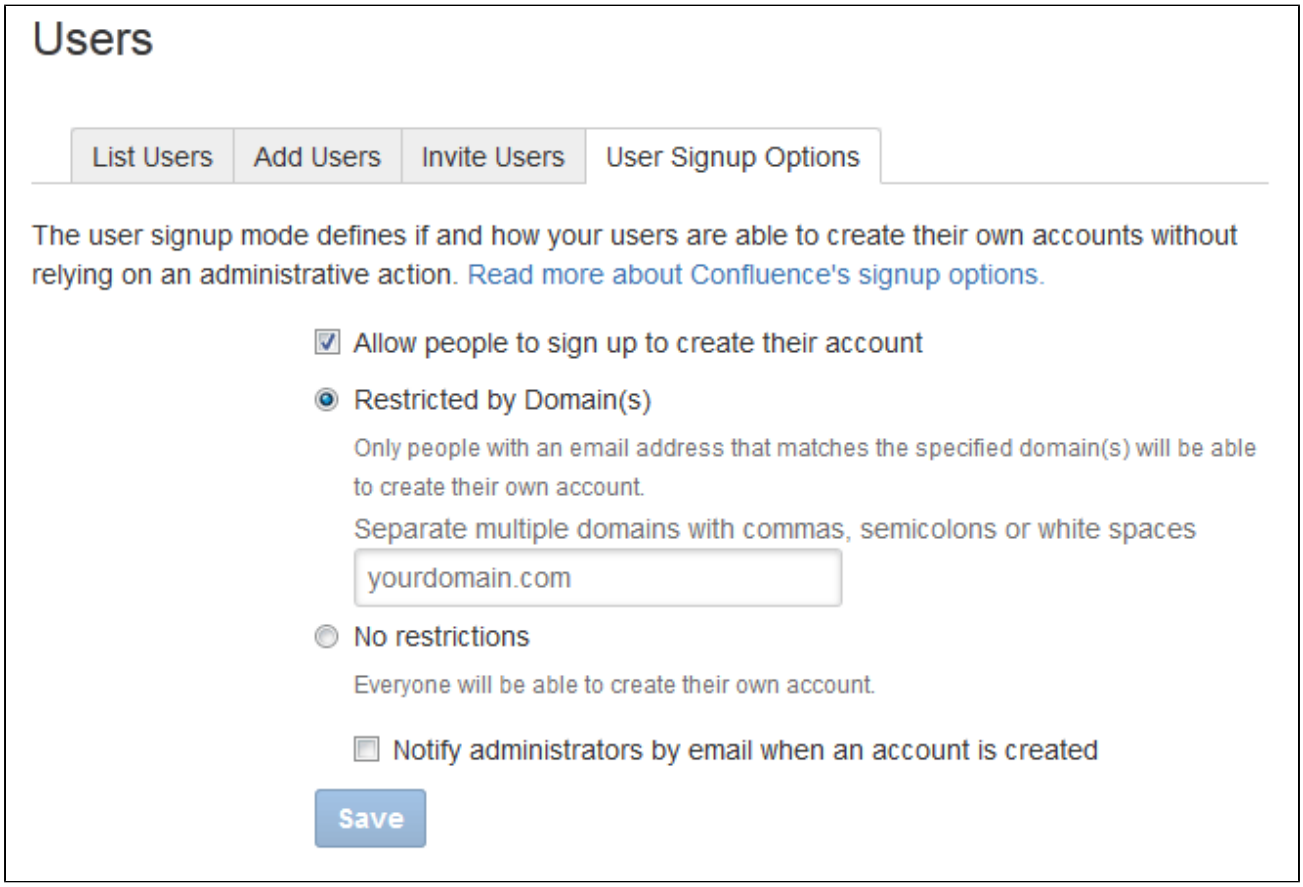

Invite people to sign up

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**InviteUsers**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 labeled**Email To**
     Separate the addresses with commas. For example:`john@example.com, sarah@example.com`
     b. Change the**Message**if you want to
     c. Choose**Send**

Reset the invitation link

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

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

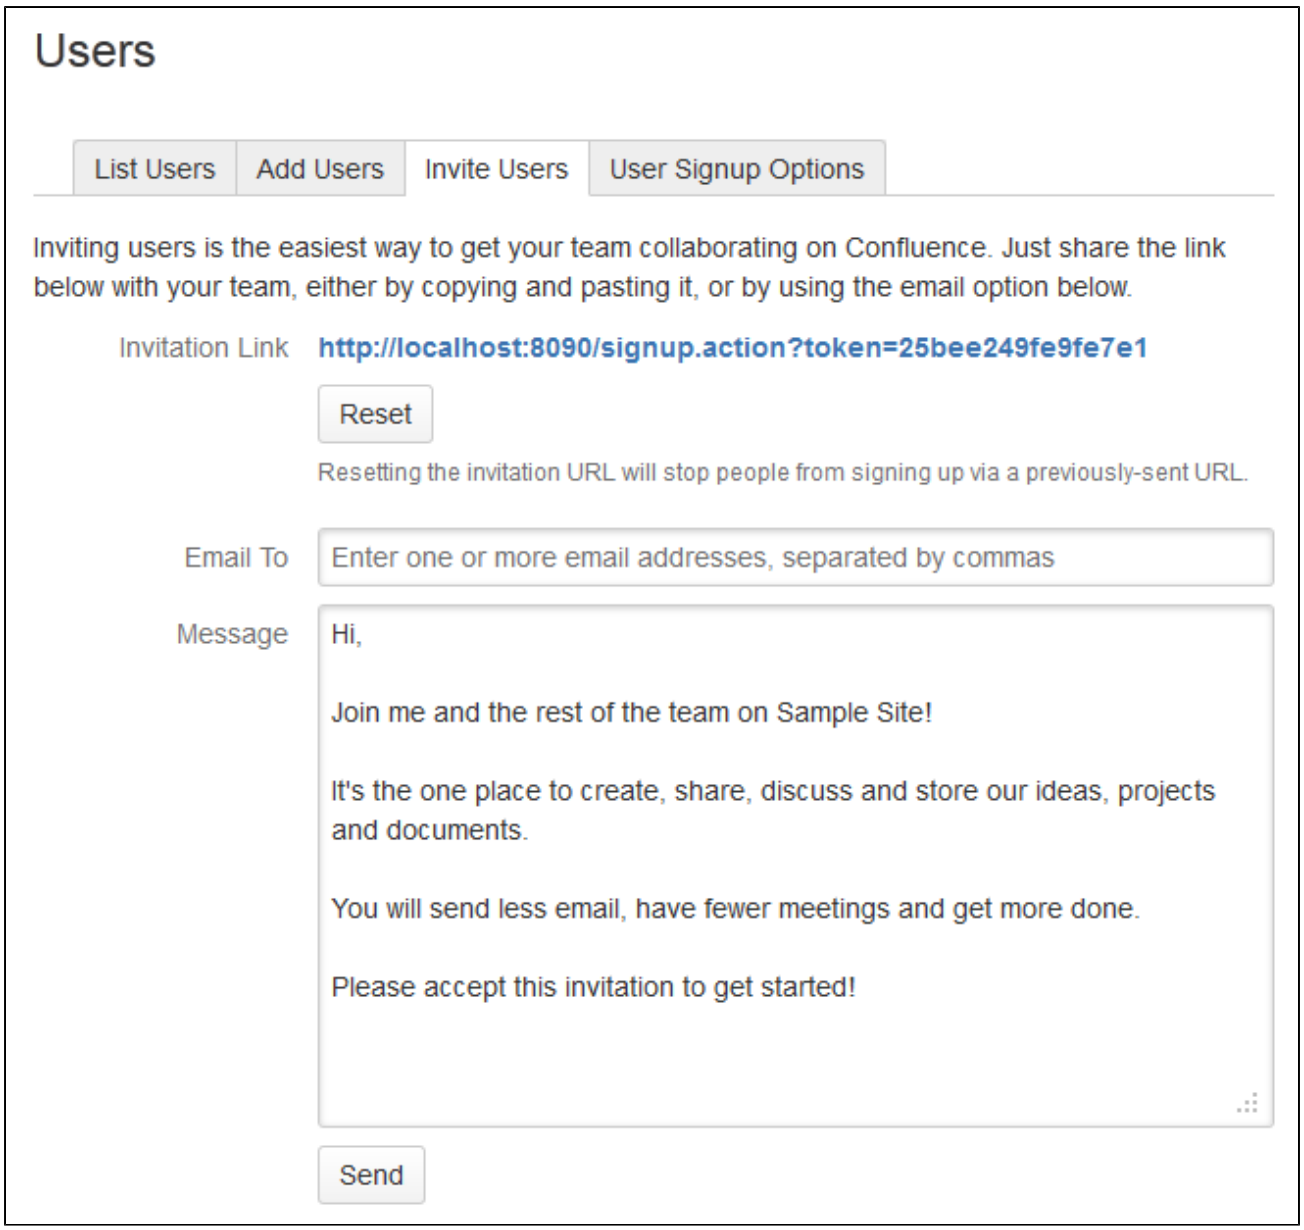

Add users manually

To add a new user:

1. Choose >**User management**
2. Select the**Add Users**tab
3. Enter the user's details

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

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

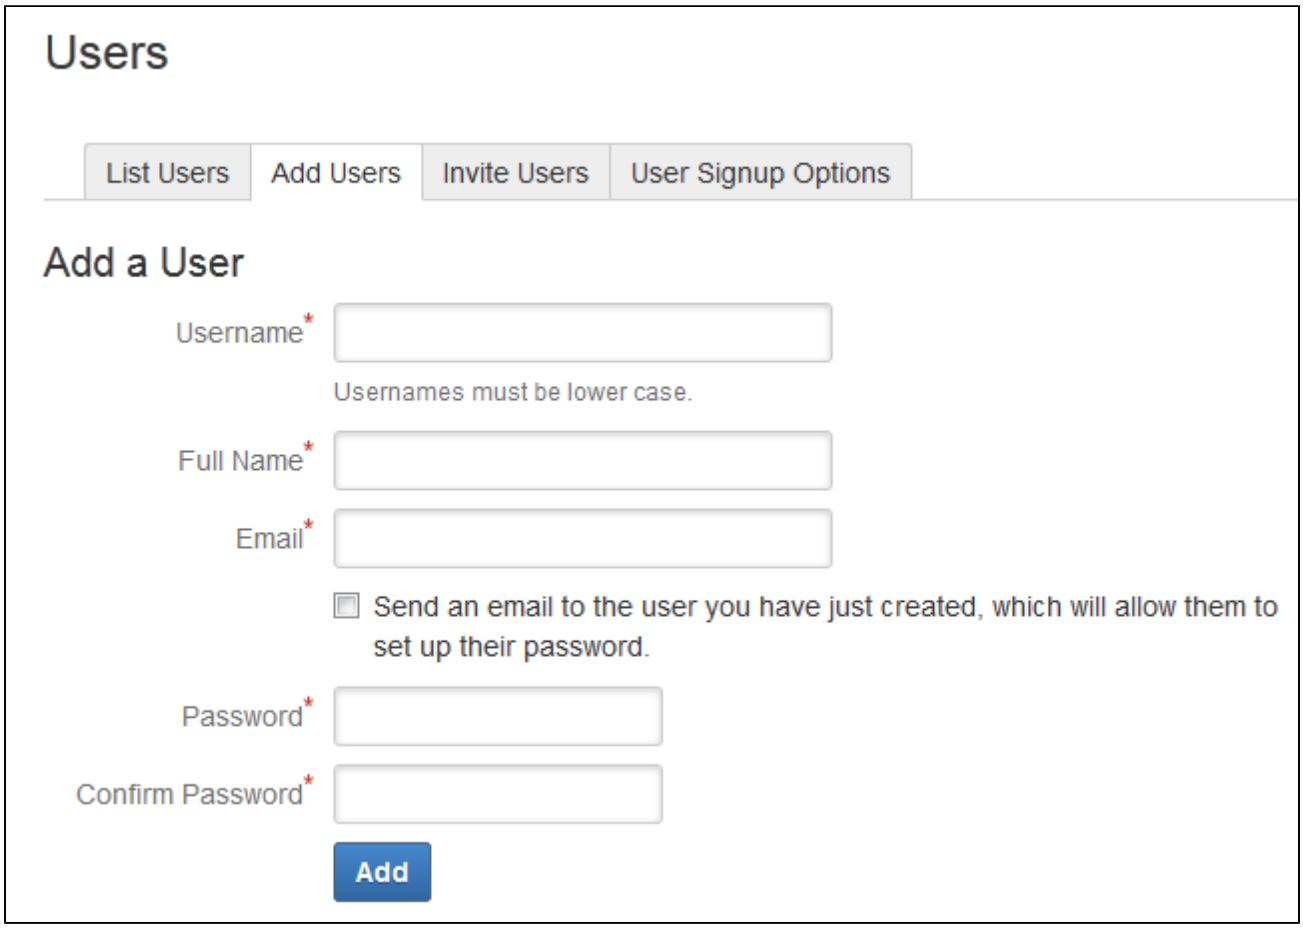

Notes

- **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.

  SeeManaging Multiple Directories.

- **Email server required for domain restricted signup and for invitations** You need to set up amail serverfor Confluence, before you can configure domain restricted signup or send email invitations to users.
- **Are the user management options not visible?**If you have external user management turned on, internal user management is disabled. To configure external user management, go to > **General Configuration**>**Security Configuration**. SeeDisabling the Built-In User Management.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

- **Avoid hash, slash and question characters in usernames** - there is a known issue where users with #, ? or / in their username cannot create spaces. See CONFSERVER-43494 GATHERING IMPACT and CONFSERVER-13479 GATHERING IMPACT for more information.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Delete or Disable Users

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

On this page:

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

Related pages:

- Manage Users
- Configuring User Directories

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

Delete, disable, or unsync?

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

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

When an user account is**disabled**:

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

When a user account is**unsynced**from all external directories:

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

When a user account is**deleted**:

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

Disable a user account

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

You need the Confluence Administrator global permission to do this.

To disable a user account:

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

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

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

Unsync a user account

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

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

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

Delete a user account

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

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

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

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

You need the Confluence Administrator global permissionto do this.

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

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

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

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

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

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

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

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

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

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

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

How deleted users appear to other people

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

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

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

Only remove access to Confluence

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

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

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

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

Limitations when deleting a user account

> The ability to delete and anonymize a user account was added in **Confluence 6.13**.
>
> For earlier Confluence versions there's a workaround you can use to permanently delete a user account via the database. See Right to erasure in Confluence Server and Data Center.
>
> You can also head to Confluence Server and Data Center GDPR support guidesto read more about Confluence and GDPR generally.

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

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Free text is not anonymised

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

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

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

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

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

Data stored in Synchrony is not deleted immediately

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

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

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

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

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

Personal spaces are not deleted

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

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

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

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

Workbox notifications don't disappear immediately

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

Data stored by third-party apps is not deleted

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

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

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

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Restore Passwords To Recover Admin User Rights

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

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

On this page:

- Use recovery mode to restore access

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

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

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

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

Use recovery mode to restore access

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

To recover administrator user rights:

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

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

   The way you do this depends on how you run Confluence. SeeConfiguring System Properties for more information on how to apply system properties.
3. Start Confluenceusing your usual method.
4. Log in to Confluence with the username**recovery_admin**and the temporary password you specified in the system property.
5. Reset the password for your existing admin account, or create a new account and add it to the appropriate administrator group.
6. Confirm that you can successfully log in with your new account.
7. Stop Confluence.
8. Remove the system property you added earlier.
9. Restart Confluence using your usual method (manually or by starting the service).

Good to know:

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

Edit User Details

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

On this page:

- Edit a user's details
- Reset login count

Related pages:

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

Edit a user's details

1. Choose >**User management**
2. Do either of the following:
    - Choose**Show all users**to list everyone inthe 'confluence-users' or 'users'group
    - Enter ausername, 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 removethis user from a group.
- **Edit Details** Change details such as the user's name, email address, contact details and team or department information. In some instances you may be able to change usernames as well. SeeChange a Usernamefor information.
- **Delete Profile Picture** - remove current and all previous profile pictures uploaded by the user.
- **Set Password** Edit the user's password details.
- **Disable** You candisable(i.e. deactivate) access for a user who no longer needs access to Confluence.
- **Delete** You can permanentlydeletea user, and replace their full name and username with an anonymous alias.

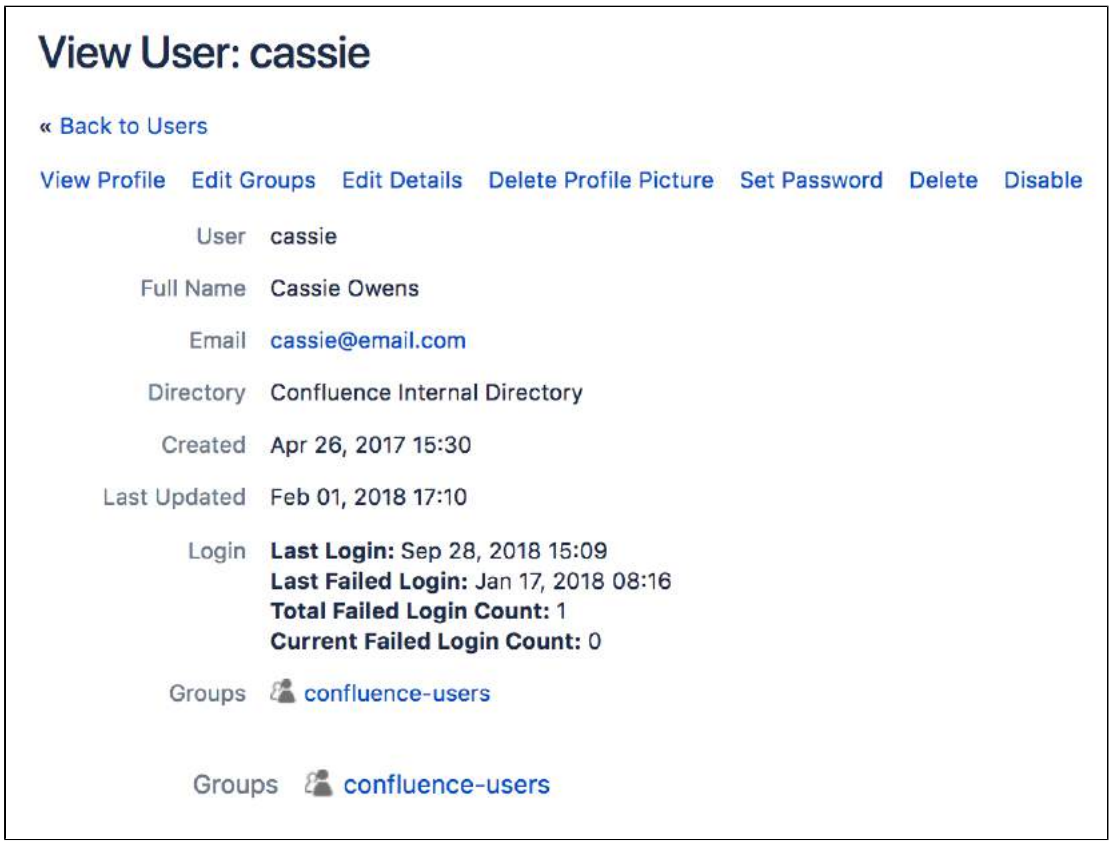

Reset login count

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

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

Multiple user directories

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

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

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

SeeManaging Multiple Directories.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

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 updateexternalusers, but it will detect changes in usernames coming from *some* external directories.

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

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

Notes

Some important things to note about changing usernames:

- **Mentions and page history** Any user mentions in current pages will automatically reflect the user's new username, but any mentions in page versions created prior to Confluence 5.3 will include the user's old username.

- **Personal Spaces** If a Confluence Administrator renames a user who has a personal space, the space key for that space will remain as the original username. For example, if jsmith's username is changed to jbrown, their personal space key will remain ~jsmith.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Managing Site-Wide Permissions and Groups

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

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

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

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

Related pages:

- Confluence Security Overview and Advisories
- Global Permissions Overview

Confluence Groups for Administrators

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

Groups can be used when setting:

- global permissions
- space permissions
- page restrictions.

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

On this page:

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

Related pages:

- Manage Users
- Global Permissions Overview

Default groups

The two default groups in Confluence are:

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

Create a new group

To add a new group:

1. Go to > **General Configuration**> **Groups**.
2. Choose**Add Group**.
3. Enter a name for your group and choose**Save**.Group names must be lower case.

You're now ready to startadding usersto the group.

Delete a group

To delete a group:

1. Go to > **General Configuration**> **Groups**.
2. Choose**Delete**next to the group you want to remove.

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

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

Confluence-administrators super group

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

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

About multiple user directories

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

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

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

SeeManaging Multiple Directories.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Adding or Removing Users in Groups

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

You can edit group membership in two places:

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

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

On this page:

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

Related pages:

- Manage Users
- Confluence Groups
- Global Permissions Overview

Add people to a group

To add members to a group from the Groups screen:

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

Screenshot: Adding members

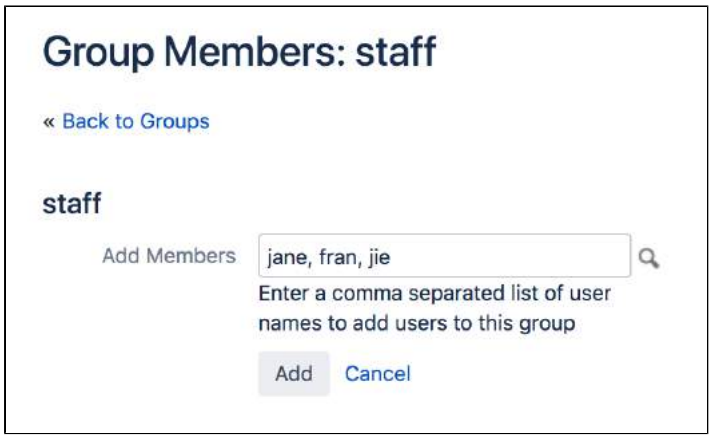

You can also change a user's group membership in the user management screen. Navigate to the user, then choose **Edit groups**, and select the groups the person should be a member of.

Remove people from a group

To remove members from a group:

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

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

About multiple directories

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

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

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

SeeManaging Multiple Directories.

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

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Global Permissions Overview

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

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

On this page:

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

Overview of global permissions

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

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

Grant global permissions

To grant global permissions to a user or group:

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

Screenshot: Editing global permissions

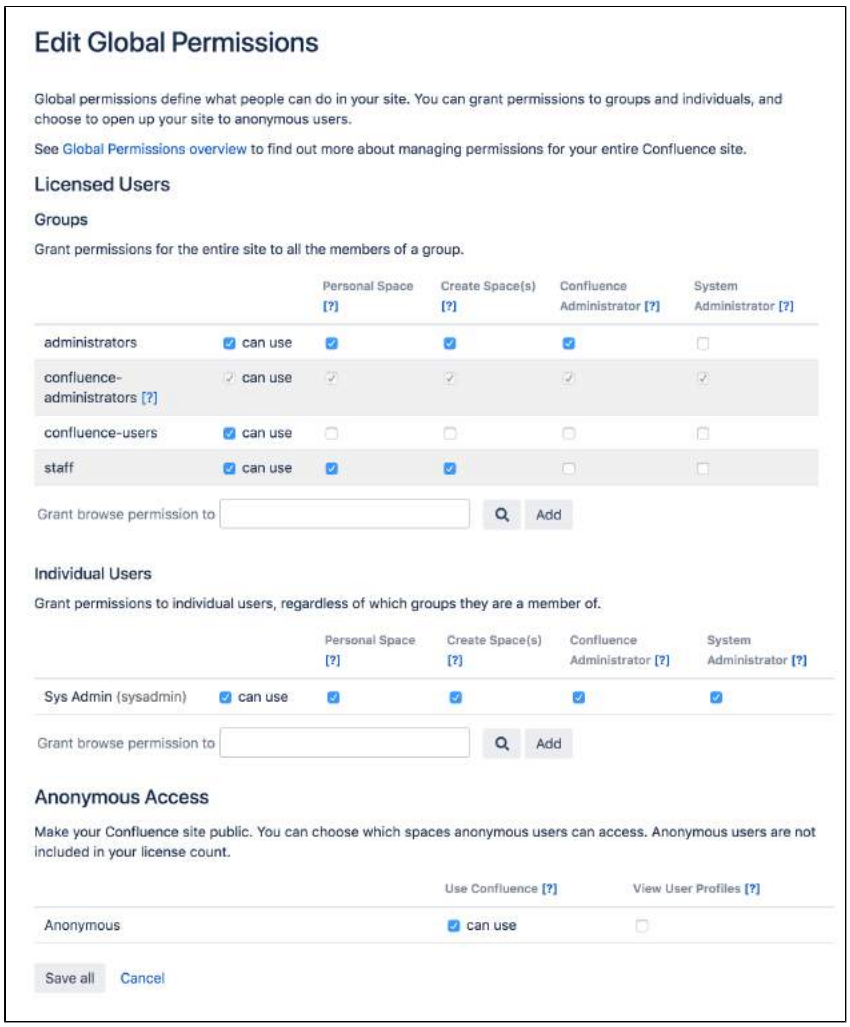

Revoke global permissions

To revoke the global permissions for a user or group:

1. Go to > **General Configuration**>**Global Permissions**
2. Choose**Edit Permissions**.
3. Locate the user or group you want to edit, and deselect all checkboxes.
4. Choose**Save all**.

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

System Administrator and Confluence Administrator permissions compared

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

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

Admin console	Confluence administrator	System administrator

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Configuration	• General Configuration, except:
	  • Base URL
	  • Connection timeout
	• Further configuration, except:
	  • Remote API
	• Languages
	• Shortcut Links
	• Global Templates and Blueprints
	• Recommended Updates Email
	• Configure Code Macro, except
	  • add new languages
	• WebDAV Configuration
		• General Configuration
		• Further configuration
		• Backup Administration
		• Languages
		• Shortcut Links
		• External Gadgets
		• Global Templates and Blueprints
		• Recommended Updates Email
		• Mail Servers
		• User Macros
		• In-app Notifications
		• Spam Prevention
		• PDF Export Language Support
		• Configure Code Macro
		• Office Connector
		• WebDAV Configuration

Marketplace	• Find new apps • Manage apps, except: • Upload an app.	• Find new apps • Manage apps

Users & security	• Users
	• Groups
	• SSO 2.0
	• Security Configuration, except:
	  • External user management
	  • Append wildcards to user and group searches
	  • Enable Custom Stylesheets for Spaces
	  • Show system information on the 500 page
	  • RSS settings
	  • XSRF Protection
	  • Attachment download security
	• Global Permissions
	• Space Permissions
	• Inspect Permissions (Data Center)
		• Users
		• Groups
		• SSO 2.0
		• Security Configuration
		• Global Permissions
		• Space Permissions
		• Inspect Permissions (Data Center)
		• User Directories
		• Whitelist

Look and feel	• Themes • Color Scheme • Site Logo and Favicon • PDF Layout • PDF Stylesheet • Sidebar, header and footer • Default Space Logo	• Themes • Color Scheme • Layouts • Stylesheet • Site Logo and Favicon • PDF Layout • PDF Stylesheet • Sidebar, header and footer • Default Space Logo • Custom HTML
Upgrade	• Latest upgrade report	• Latest upgrade report • Plan your upgrade

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

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

Confluence-administrators super group

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

When you install Confluence you'll be prompted to create a system administrator account. This user will be a member of the`confluence-administrators`super group.

What can members of this group do?

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

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

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

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

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

Should I use the confluence-administrators group?

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

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

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

Troubleshooting

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

- 'User/Group not found' - This message may appear if your LDAP repository is unavailable, or if the user/group has been deleted after the permission was created.
- If you're unable to log in to Confluence as an administrator (for example, you've lost the administrator password) you can start Confluence in recovery mode to recover your admin user rights. SeeRestore Passwords To Recover Admin User Rights.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Setting Up Public Access

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

On this page:

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

Related pages:

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

Allow anonymous access to the site

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

To enable anonymous access to your site:

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

Disable anonymous access to the site

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

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

Allow anonymous access to a space

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

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

Alternatives to making your site public

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

Revoke access for unlicensed users from Jira Service Management

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

To revoke access for unlicensed users:

1. Go to > **General Configuration**>**Global Permissions**.
2. 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 Management.

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

Screenshot: Unlicensed access section of the Global Permissions page.

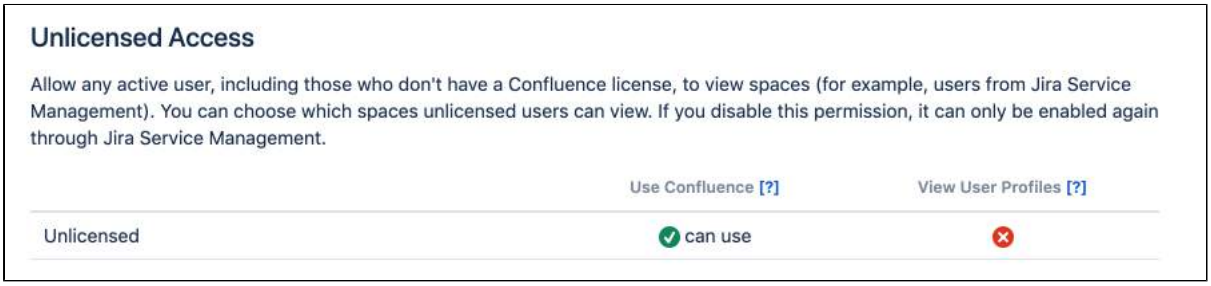

This section only appears on the Global Permissions page in Confluence if you have linked a space to your service project (as a Knowledge base), and chosen to allow all active users and customers to access without a Confluence license. SeeSet up a knowledge base for self-servicein the Service Management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

Configuring User Directories

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

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

On this page:

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

Related pages:

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

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

Configuring User Directories in Confluence

To configure your Confluence user directories:

1. Choose the**cog icon** , then choose**General Configuration**
2. Click '**User Directories**' in the left-hand panel.

Connecting to a Directory

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

- Confluence's internal directory. See Configuring the Internal Directory.
- Microsoft Active Directory. See Connecting to an LDAP Directory.
- Various other LDAP directory servers. See Connecting to an LDAP Directory.
- An LDAP directory for delegated authentication. See Connecting to an Internal Directory with LDAP Authentication.
- Atlassian Crowd or Jira 4.3 or later. See Connecting to Crowd or Jira for User Management.

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

Updating Directories

Limitations when Editing Directories

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

This limitation applies to all directory types. For example:

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

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

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

Recommendations

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

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

Enabling, Disabling and Removing Directories

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

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

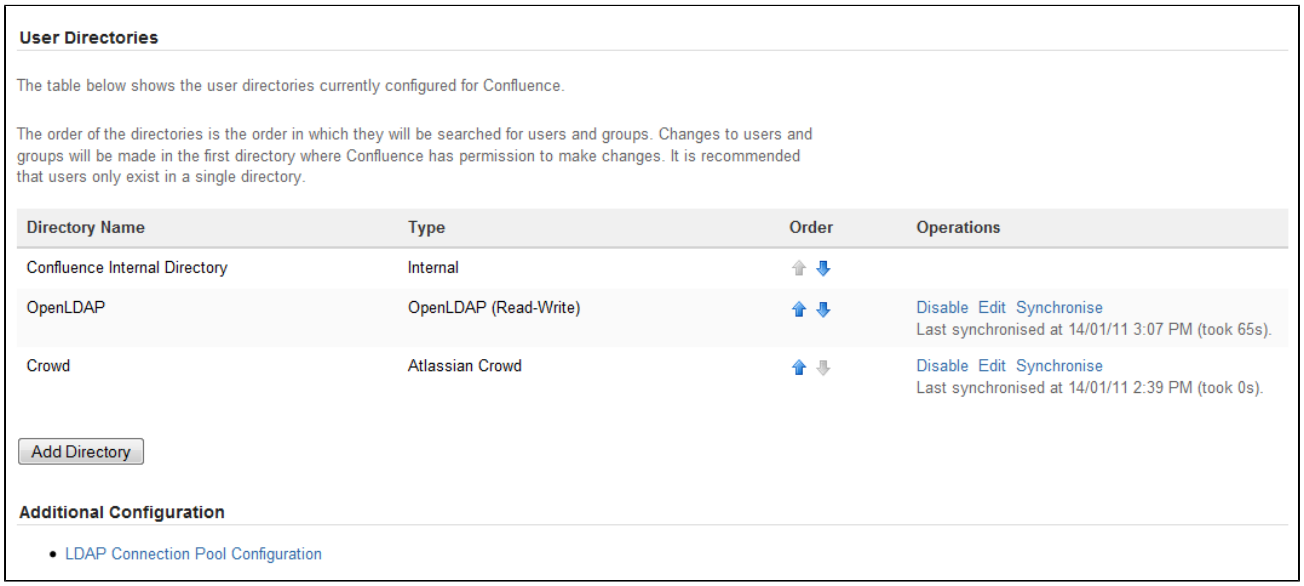

Screenshot above: Configuring user directories

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Configuring the Internal Directory

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

On this page:

- Overview
- Diagram of Possible Configuration

Related pages:

- Configuring User Directories
- How to Reenable the Internal Directory (Knowledge base article)

Overview

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

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

Diagram of Possible Configuration

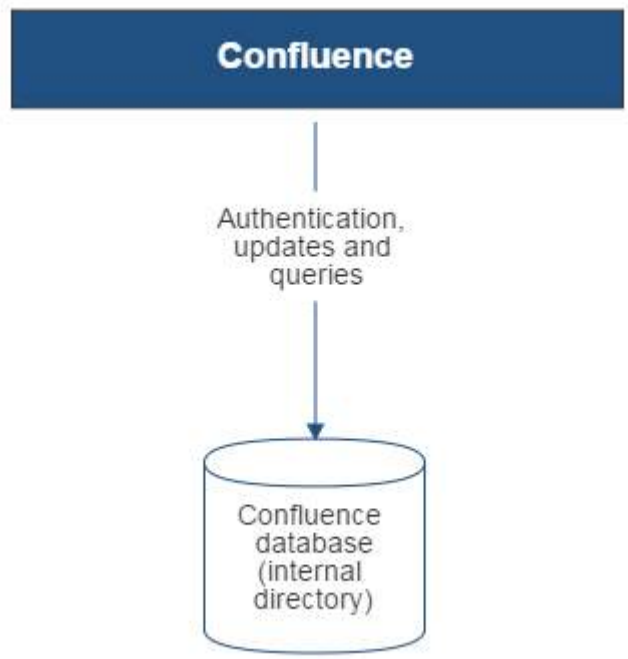

Diagram above: Confluence using its internal directory for user management.

Connecting to an LDAP Directory

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

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

On this page:

- Overview
- 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

Overview

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

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

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

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

Connecting to an LDAP Directory in Confluence

To connect Confluence to an LDAP directory:

1. Choose the**cog icon** , then choose**General Configuration**
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 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` By default, all users can read the uSNChanged attribute; however, only administrators or users with relevant permissions can access the Deleted Objects container. The specific privileges required by the user to connect to LDAP are "Bind" and "Read" (user info, group info, group membership, update sequence number, deleted objects), which the user can obtain by being a member of the Active Directory's built-in administrators group. Note that the incremental sync will fail silently if the Active Directory is accessed by a user without these privileges. This has been reported asCWD-3093.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

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`

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

Permission Settings

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

Setting	Description
Read Only	LDAP users, groups and memberships are retrieved from your directory server and can only be modified via your directory server. You cannot modify LDAP users, groups or memberships via the application administration screens.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

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 ~~CONFSERVER-28621~~ - User Loses all Local Group Memberships If LDAP Sync is Unable to find the User, but the User appears again in subsequent syncs CLOSED
Read /Write	LDAP users, groups and memberships are retrieved from your directory server. When you modify a user, group or membership via the application administration screens, the changes will be applied directly to your LDAP directory server. 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 behavior allows users to be removed from automatically-added groups. In Confluence 3.5 and 3.5.1, they would be re-added upon next login. Please note that there is no validation of the group names. If you mis-type the group name, authorization failures will result users will not be able to access the applications or functionality based on the intended group name. Examples: • `confluence-users` • `confluence-users,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*. Nested groups simplify permissions by allowing sub-groups to inherit permissions from a parent group.
Manage User Status Locally	If true, you can activate and deactivate users in Crowd independent of their status in the directory server.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Filter out expired users	If true, user accounts marked as expired in Active Directory will be automatically removed. For cached directories, the removal of a user will occur during the first synchronization after the account's expiration date. **Note**: This is available in Embedded Crowd 2.0.0 and above, but not available in the 2.0.0 m04 release.
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 Synchronization	Enable incremental synchronization if you only want changes since the last synchronization to be queried when synchronizing a directory. ⚠ Be aware that when using this option, the user account configured for synchronization must have read access to: • The`uSNChanged` attribute of all users and groups in the directory that need to be synchronized. • The objects and attributes in the Active Directory deleted objects container. If at least one of these conditions is not met, you may end up with users who are added to (or deleted from) the Active Directory not being respectively added (or deleted) in the application. This setting is only available if the directory type is set to "Microsoft Active Directory".

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Update group memberships when logging in	This setting enables updating group memberships during authentication and can be set to the following options:
	• **Every time the user logs in**: during the authentication, the users **direct** group memberships will be updated to match whats in the remote directory:
	  • Remove the user from all groups that the user no longer belongs to in the remote directory.
	  • Add the user to all the groups that the user belongs to in the remote directory. New groups with matching names and descriptions will be created locally if needed. The group will only contain the current user and other memberships will be populated when users who belong to the same group log in or when the synchronization happens.
	• **For newly added users only**: when a new user logs in for the first time, the users **direct** group memberships will be updated to match whats in the remote directory.
	  ⓘ Consider that the user's group memberships will be updated only if the user was created during the authentication.
	• **Never**: during the authentication, the user's group memberships wont change, even if the local state doesnt match whats in the remote.

Synchronization Interval (minutes)	Synchronization is the process by which the application updates its internal store of user data to agree with the data on the directory server. The application will send a request to your directory server every x minutes, where 'x' is the number specified here. The default value is 60 minutes.
Read Timeout (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 10. • The time to wait when getting a connection from the connection pool. A value of 0 (zero) means there is no limit, so wait indefinitely. • The time, in seconds, to wait when opening new server connections. A value of 0 (zero) means that the TCP network timeout will be used, which may be several minutes.

User Schema Settings

Setting	Description
User Object Class	This is the name of the class used for the LDAP user object. Example: • `user`
User Object Filter	The filter to use when searching user objects. Example: • `(&(objectCategory=Person)(sAMAccountName=*))` More examples can be found in our knowledge base. See How to write LDAP search filters.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

User Name Attribute	The attribute field to use when loading the username. Examples: • `cn` • `sAMAccountName` NB: In Active Directory, the 'sAMAccountName' is the 'User Logon Name (pre-Windows 2000)' field. The User Logon Name field is referenced by 'cn'.
User Name RDN Attribute	The RDN (relative distinguished name) to use when loading the username. The DN for each LDAP entry is composed of two parts: the RDN and the location within the LDAP directory where the record resides. The RDN is the portion of your DN that is not related to the directory tree structure. Example: • `cn`
User First Name Attribute	The attribute field to use when loading the user's first name. Example: • `givenName`
User Last Name Attribute	The attribute field to use when loading the user's last name. Example: • `sn`
User Display Name Attribute	The attribute field to use when loading the user's full name. Example: • `displayName`
User Email Attribute	The attribute field to use when loading the user's email address. Example: • `mail`
User Password Attribute	The attribute field to use when loading a user's password. Example: • `unicodePwd`
User Unique ID Attribute	The attribute used as a unique immutable identifier for user objects. This is used to track username changes and is optional. If this attribute is not set (or is set to an invalid value), user renames will not be detected they will be interpreted as a user deletion then a new user addition. This should normally point to a UUID value. Standards-compliant LDAP servers will implement this as 'entryUUID' according toRFC 4530. This setting exists because it is known under different names on some servers, e.g. 'objectGUID' in Microsoft Active Directory.

Group Schema Settings

Setting	Description

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

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

Membership Schema Settings

Setting	Description
Group Members Attribute	The attribute field to use when loading the group's members. Example: • `member`
User Membership Attribute	The attribute field to use when loading the user's groups. Example: • `memberOf`
Use the User Membership Attribute, when finding the user's group membership	Check this if your directory server supports the group membership attribute on the user. (By default, this is the '`memberOf`' attribute.) • If this checkbox is selected, your application will use the group membership attribute on the user when **retrieving the list of groups to which a given user belongs**. This will result in a more efficient retrieval. • If this checkbox is not selected, your application will use the members attribute on the group ('`member`' by default) for the search. • If the **Enable Nested Groups** checkbox is selected, your application will ignore the **Use the User Membership Attribute** option and will use the members attribute on the group for the search.
Use the User Membership Attribute, when finding the members of a group	Check this if your directory server supports the user membership attribute on the group. (By default, this is the '`member`' attribute.) • If this checkbox is selected, your application will use the group membership attribute on the user when **retrieving the members of a given group**. This will result in a more efficient search. • If this checkbox is not selected, your application will use the members attribute on the group ('`member`' by default) for the search.

Diagrams of Some Possible Configurations

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

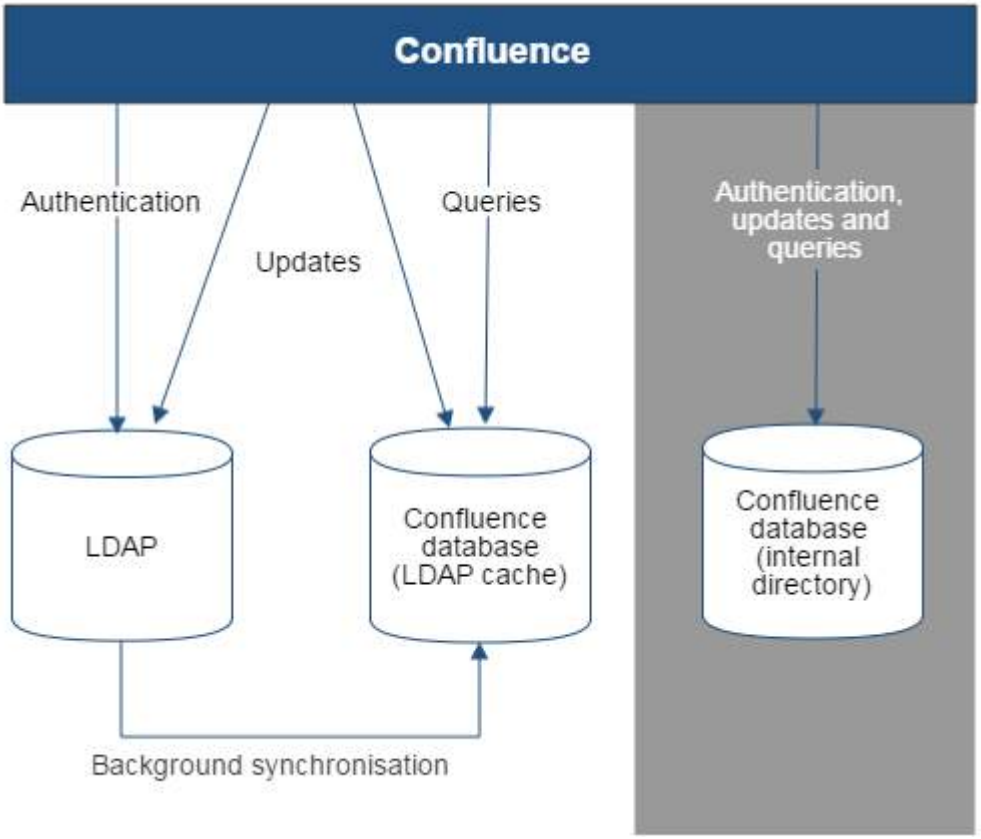

Diagram above: Confluence connecting to an LDAP directory.

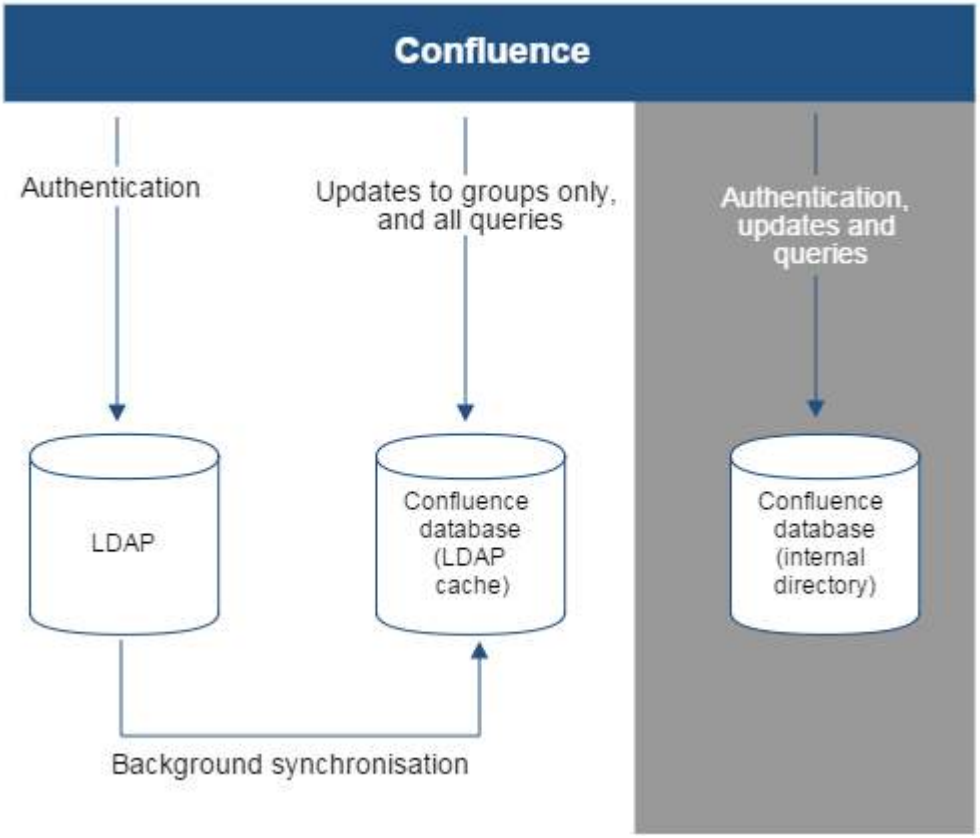

Diagram above: Confluence connecting to an LDAP directory with permissions set to read only and local groups.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Configuring the LDAP Connection Pool

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

Related pages:

- Connecting to an LDAP Directory
- Configuring User Directories

> **Use system properties instead of updating the settings in the application**
>
> Because of a bug, the Update Settings form used to specify the LDAP connection pool properties directly in the application wont work. Any values that you set will not have any effect. Use this page to configure the connection pool using system properties instead.

Configure the LDAP connection pool

To configure the JNDI connection pool:

1. Go to `<installation-directory>/bin`, and edit the `setenv.sh` (Linux) or `setenv.bat` (Windows) file.
2. Set the properties from the table below, for example:

```
-Dcom.sun.jndi.ldap.connect.pool.initsize=2
-Dcom.sun.jndi.ldap.connect.pool.prefsize=1
-Dcom.sun.jndi.ldap.connect.pool.maxsize=20
```

To view your LDAP connection pool:

1. Choose the**cog icon** , then choose**General Configuration**
2. Click '**User Directories**' in the left-hand panel.
3. Click '**LDAP Connection Pool Configuration**' in the 'Additional Configuration' section.

Pool properties

Setting	System property	Description	Default value
Initial pool size	com.sun.jndi.ldap.connect.pool.initsize	The number of LDAP connections created when initially connecting to the pool.	1
Preferred pool size	com.sun.jndi.ldap.connect.pool.prefsize	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	com.sun.jndi.ldap.connect.pool.maxsize	The max 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 server will be blocked. A value of 0 (zero) means that the number of connections is unlimited.	0
Pool timeout	com.sun.jndi.ldap.connect.pool.timeout	The length of time, in milliseconds, 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.	30000
Pool protocol	com.sun.jndi.ldap.connect.pool.protocol	Only these protocol types are allowed to connect to LDAP. 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	com.sun.jndi.ldap.connect.pool.authentication	Only these authentication types are allowed to connect to LDAP. 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.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Configuring an SSL Connection to Active Directory

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

On this page:

- Prerequisites
- Step 1. Install the Active Directory Certificate Services
- Step 2. Obtain the Server Certificate
- Step 3. Import the Server Certificate

Related pages:

- Connecting to an LDAP Directory
- Configuring User Directories

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

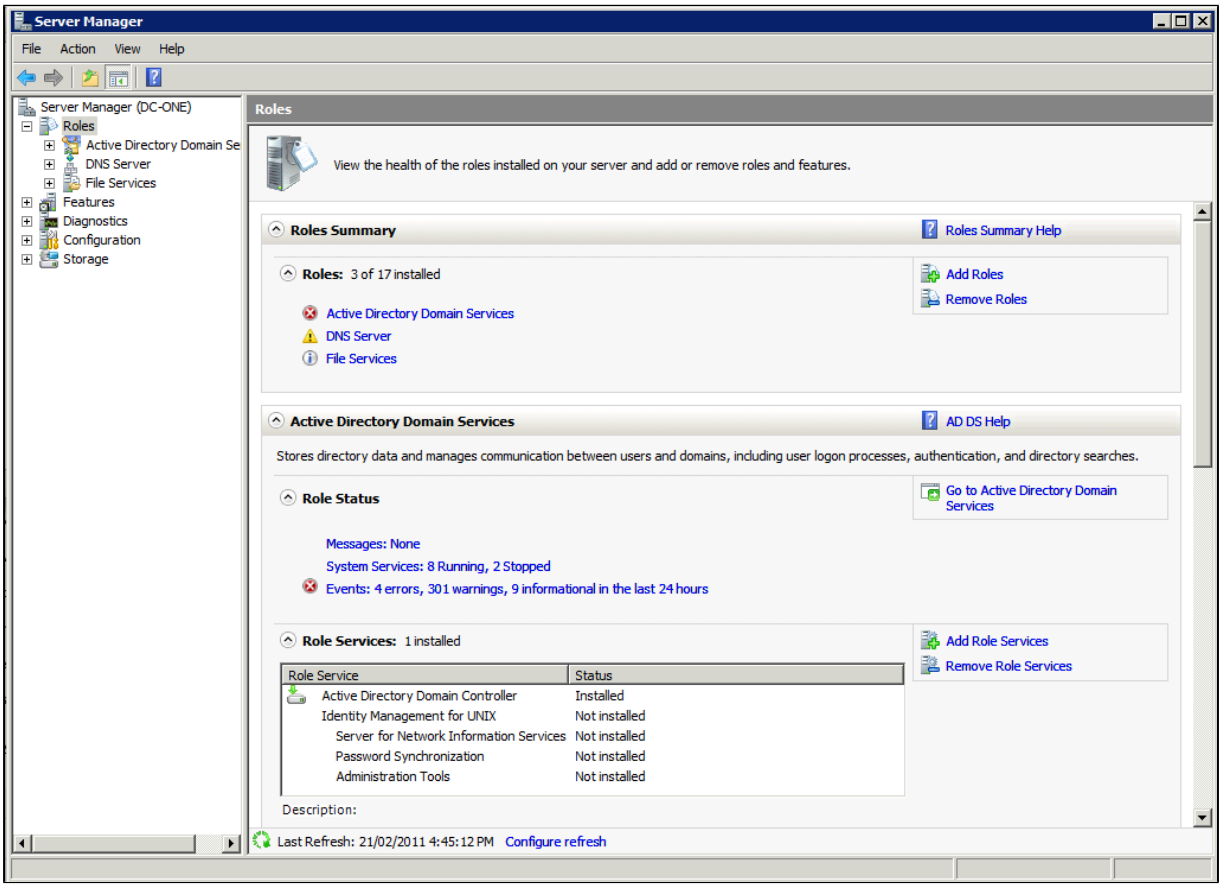

4. On the **Select Server Roles** page, select the **Active Directory Certificate Services** check box. Click **Next** twice.

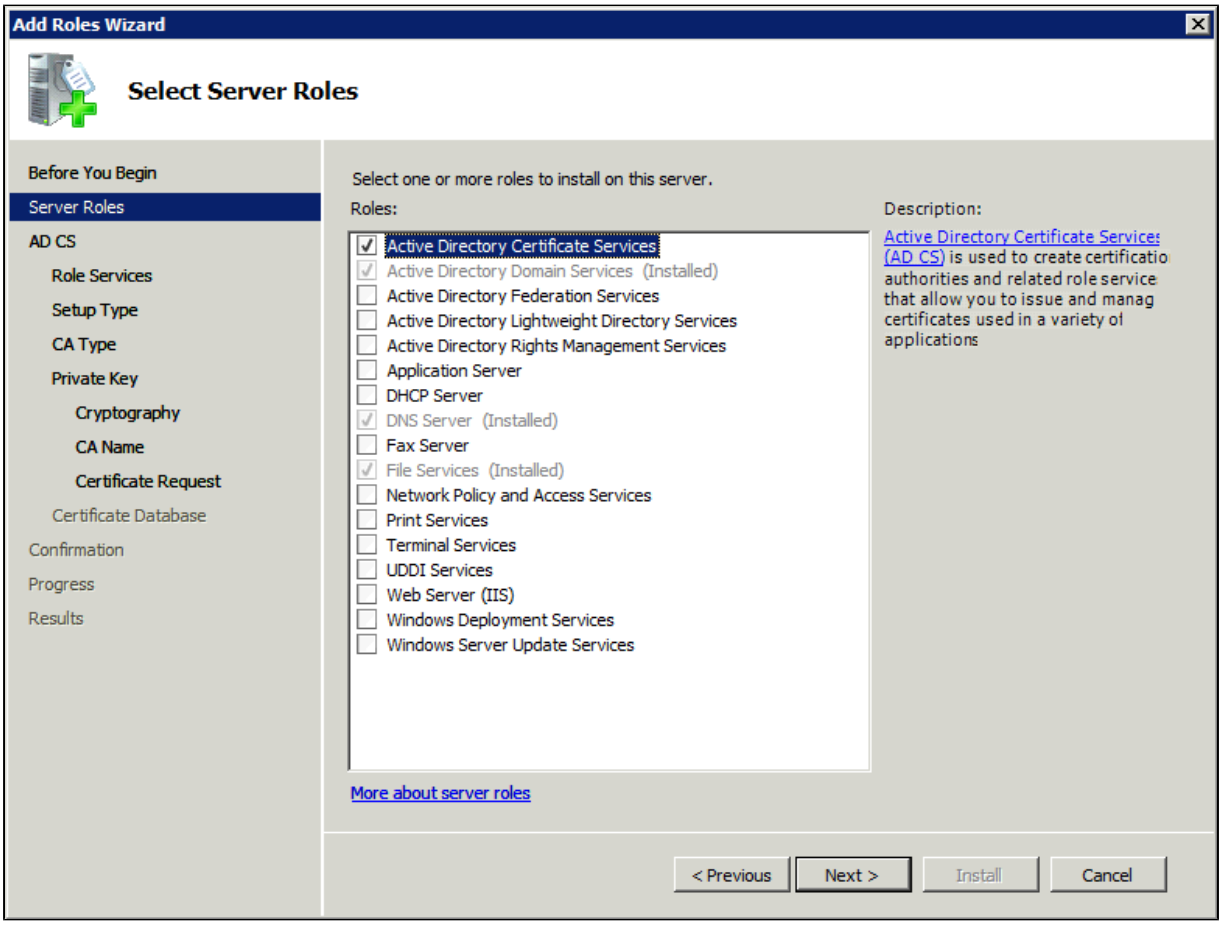

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

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

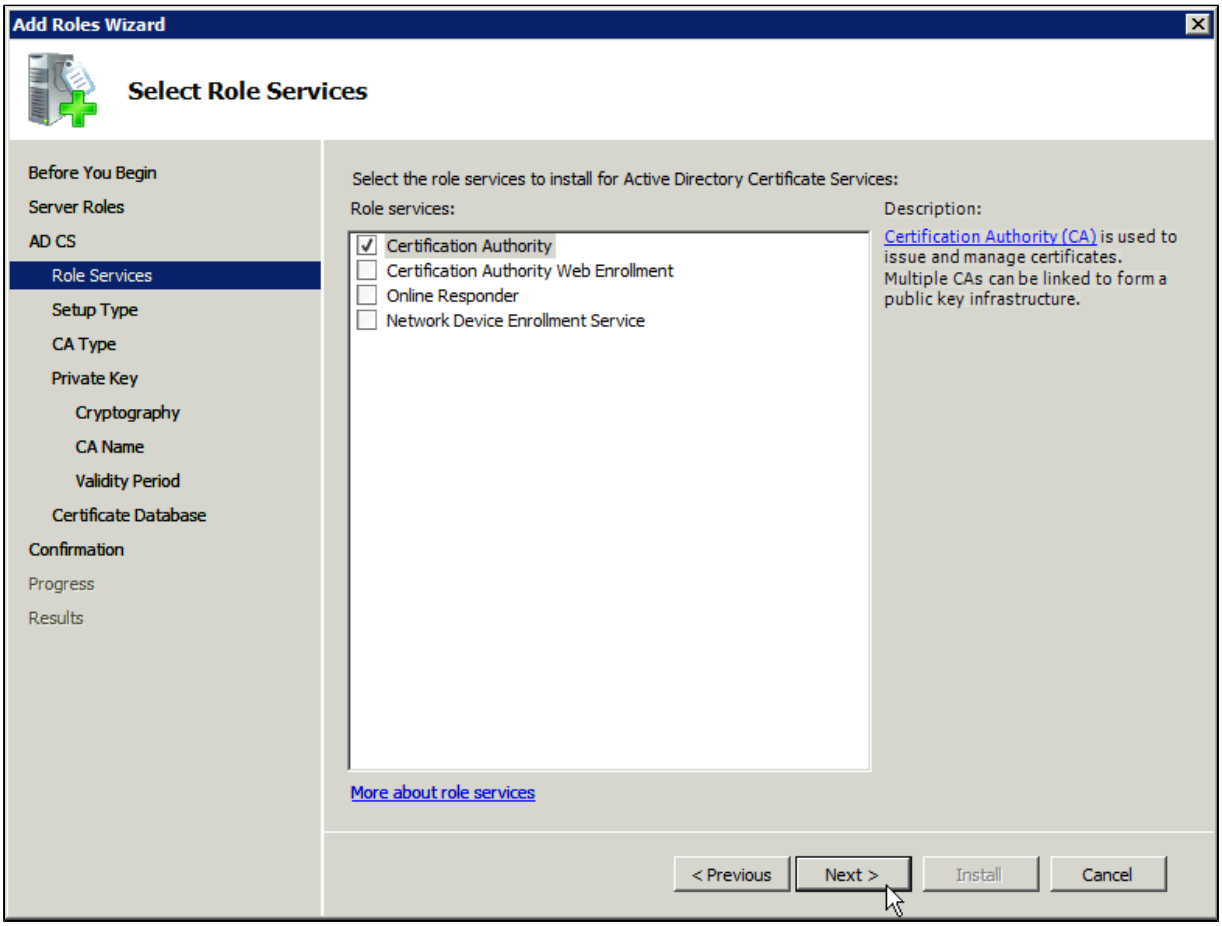

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

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

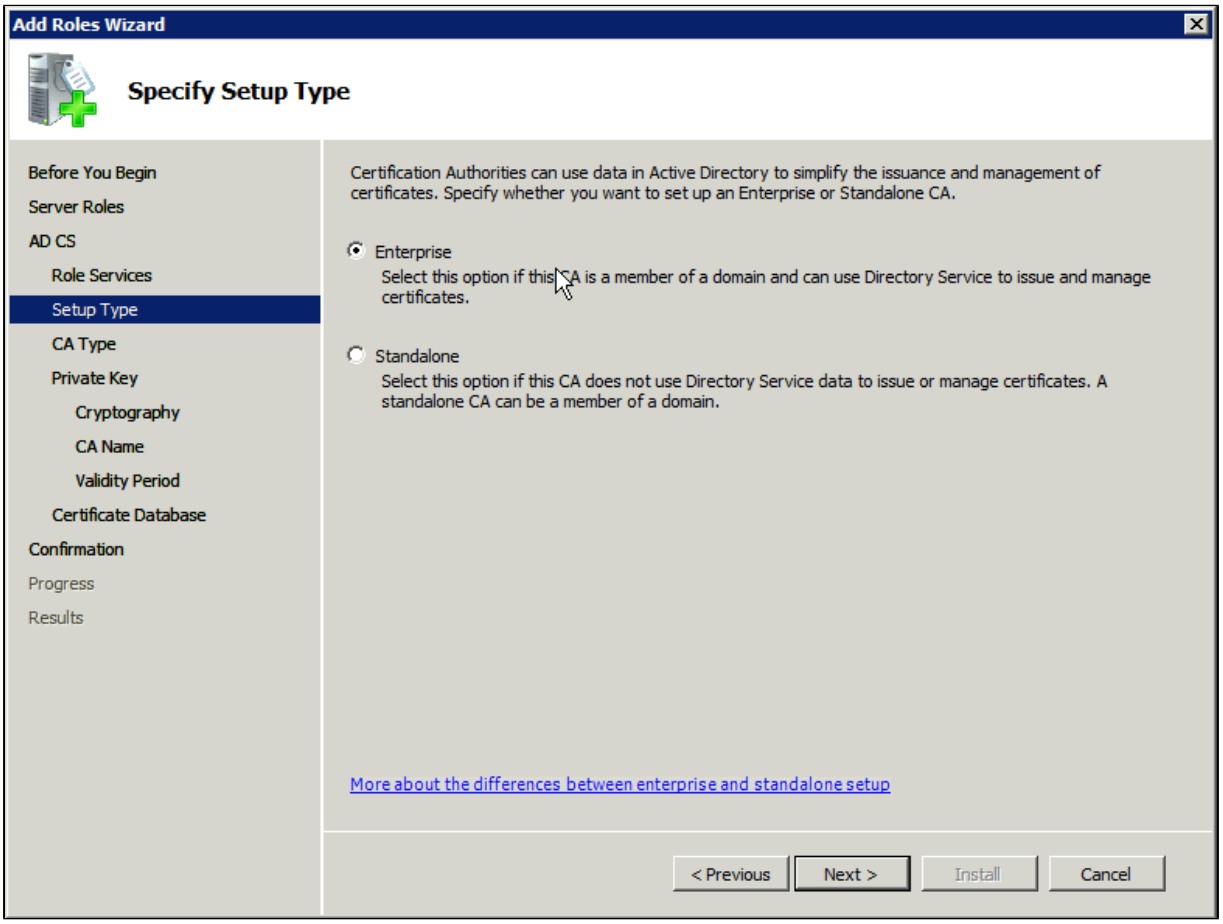

7. On the **Specify CA Type** page, click **Root CA**, and then click **Next**.

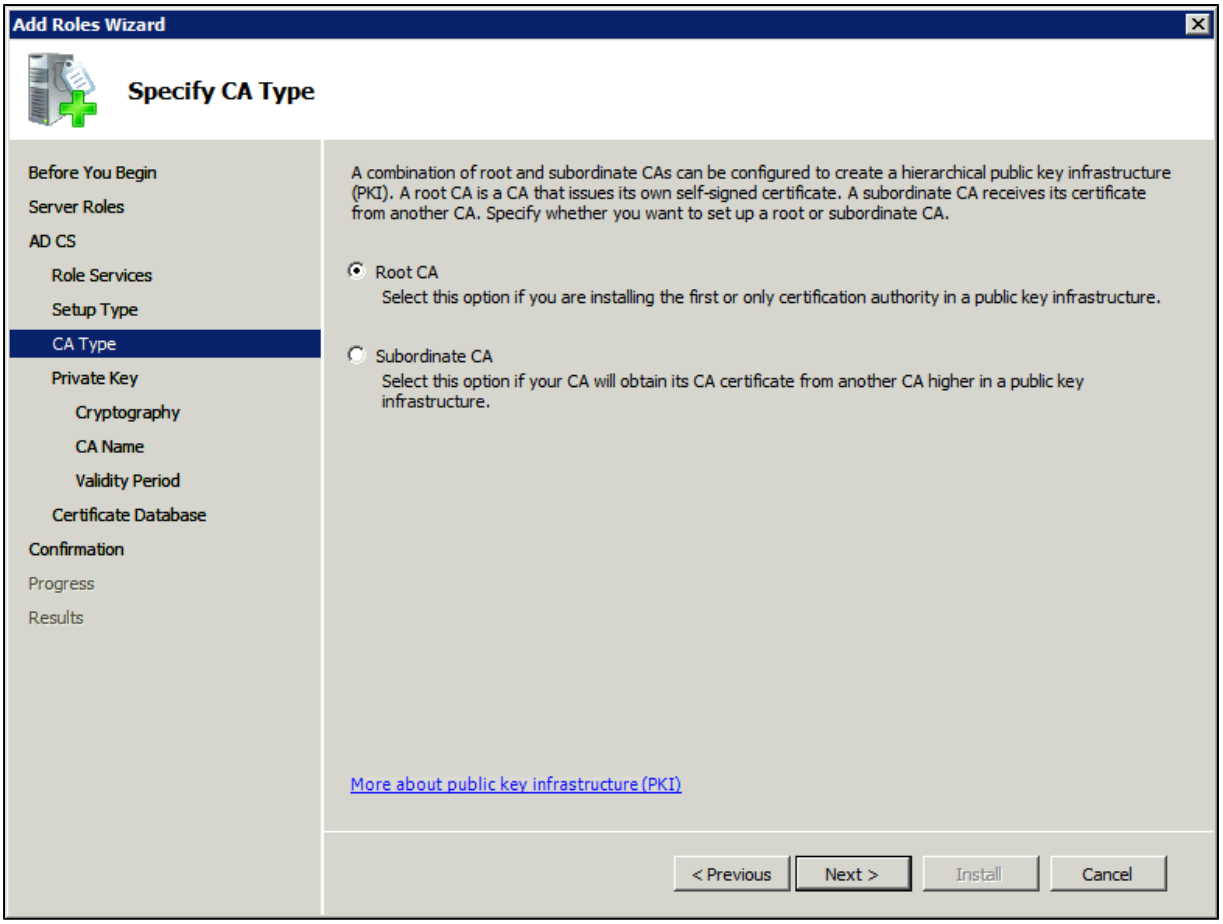

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

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.

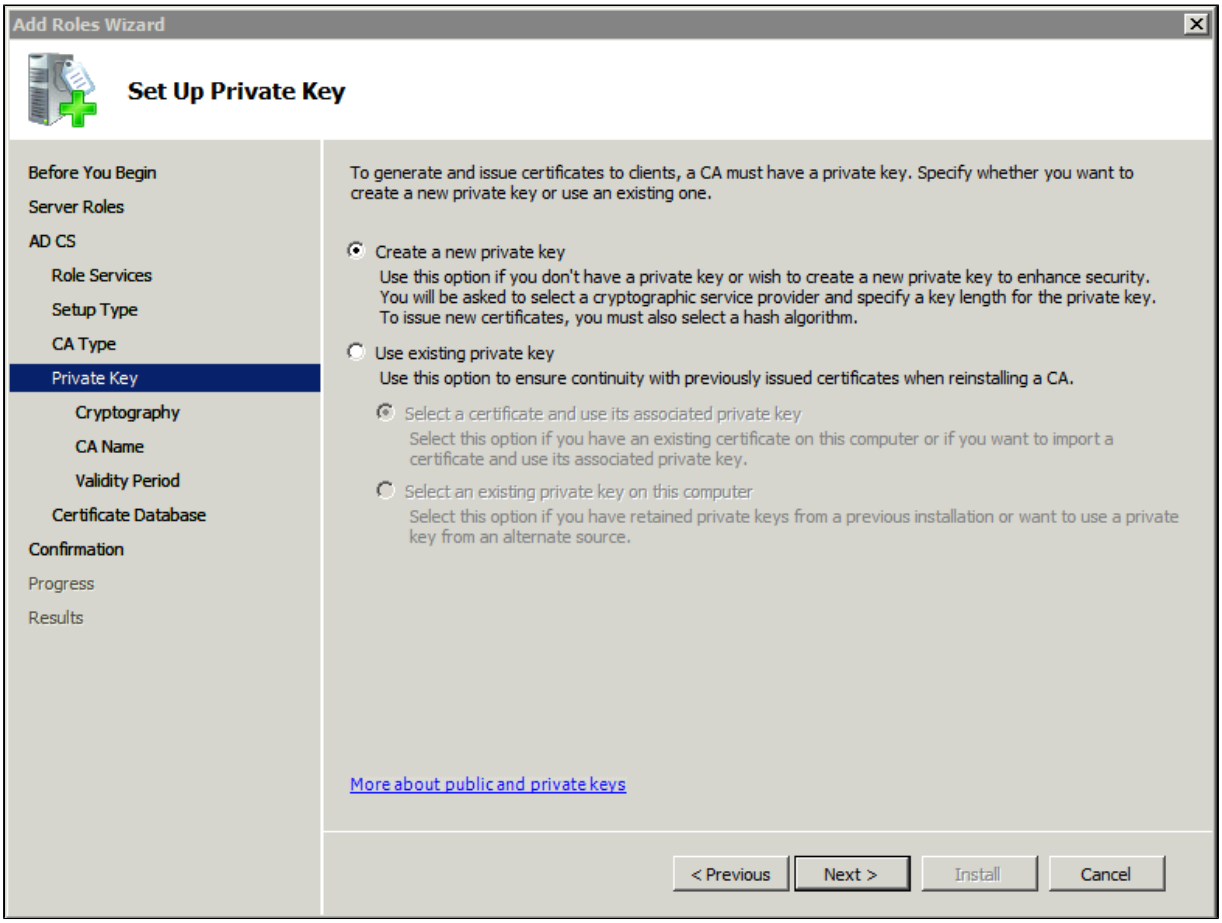

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

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

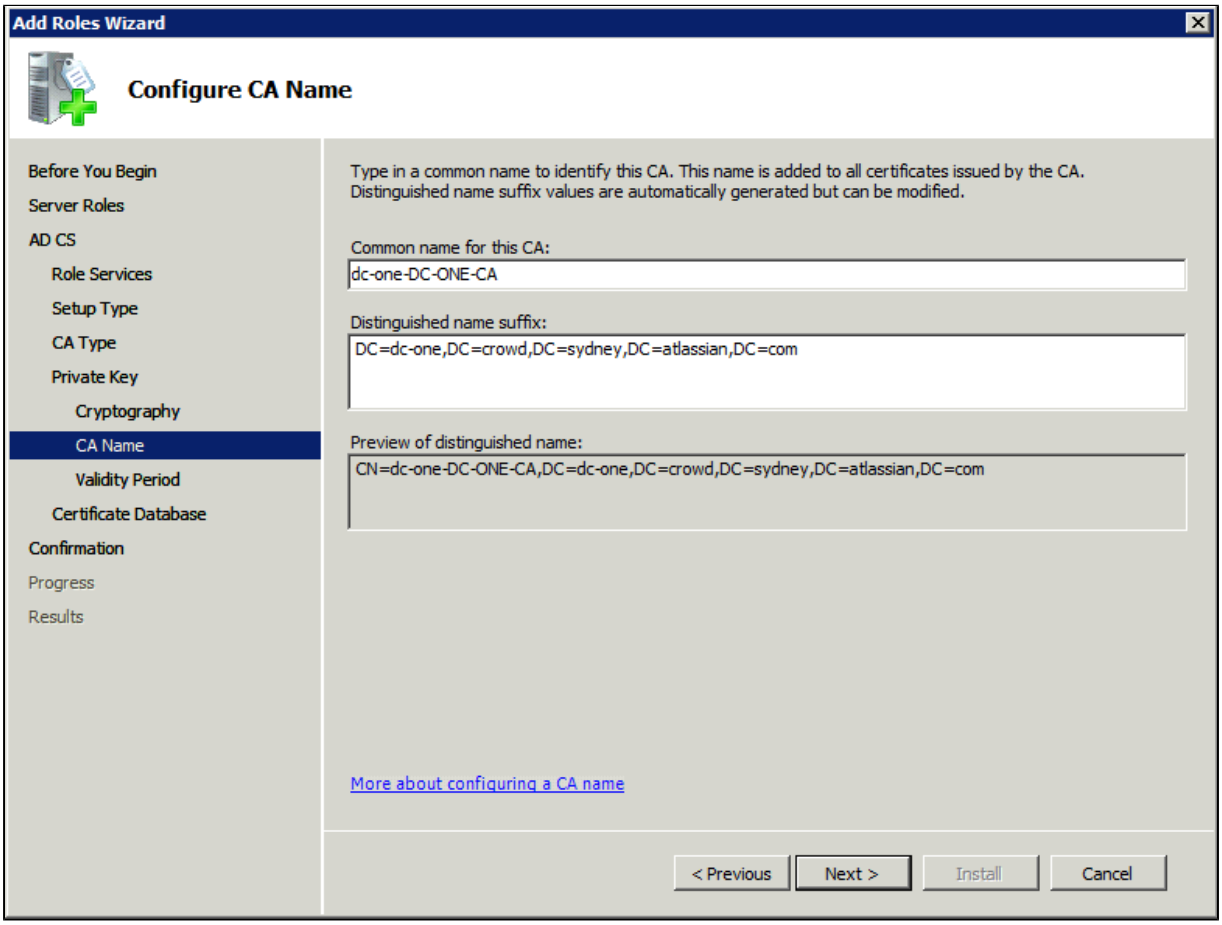

10. On the **Set Validity Period** page, accept the default values or specify other storage locations for the certificate database and the certificate database log, and then click **Next**.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

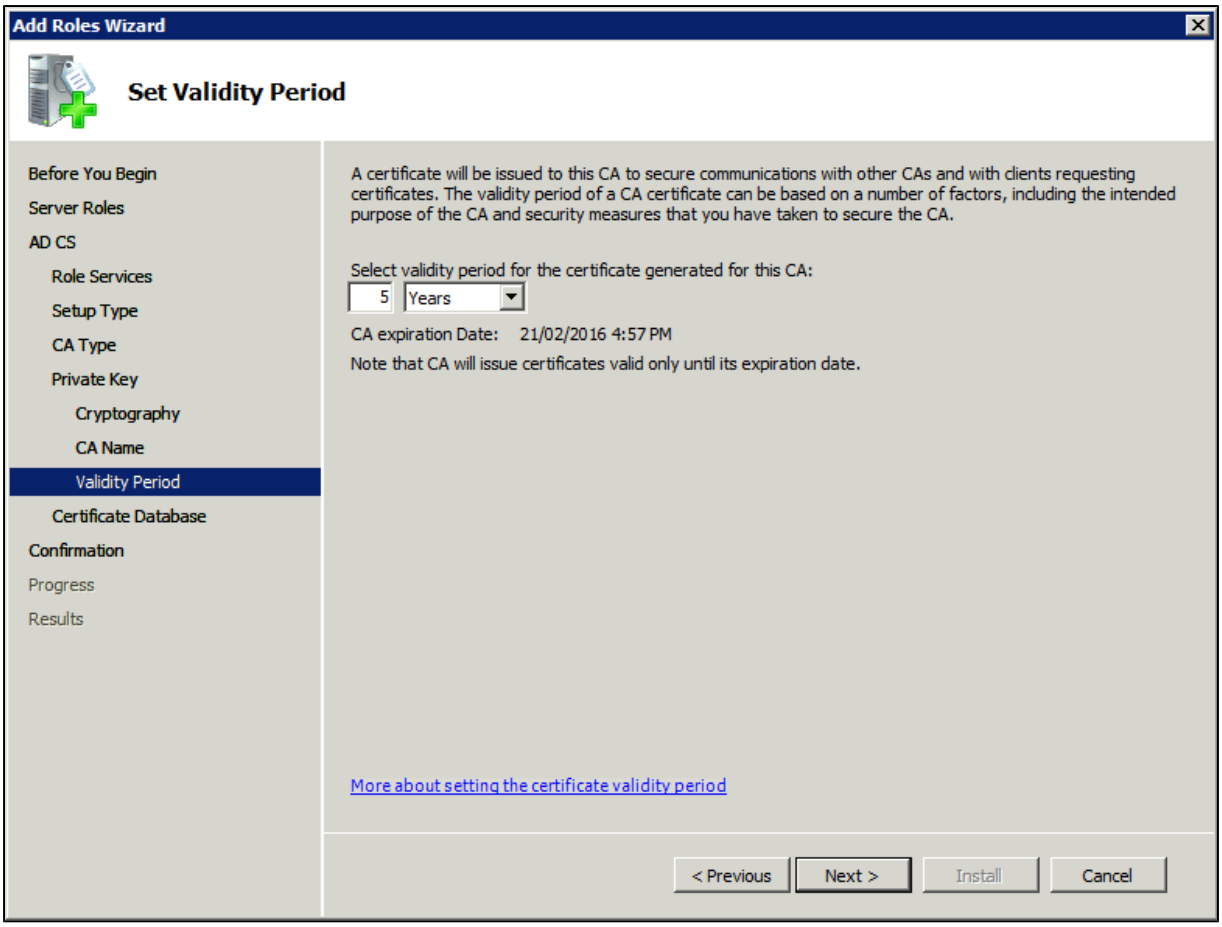

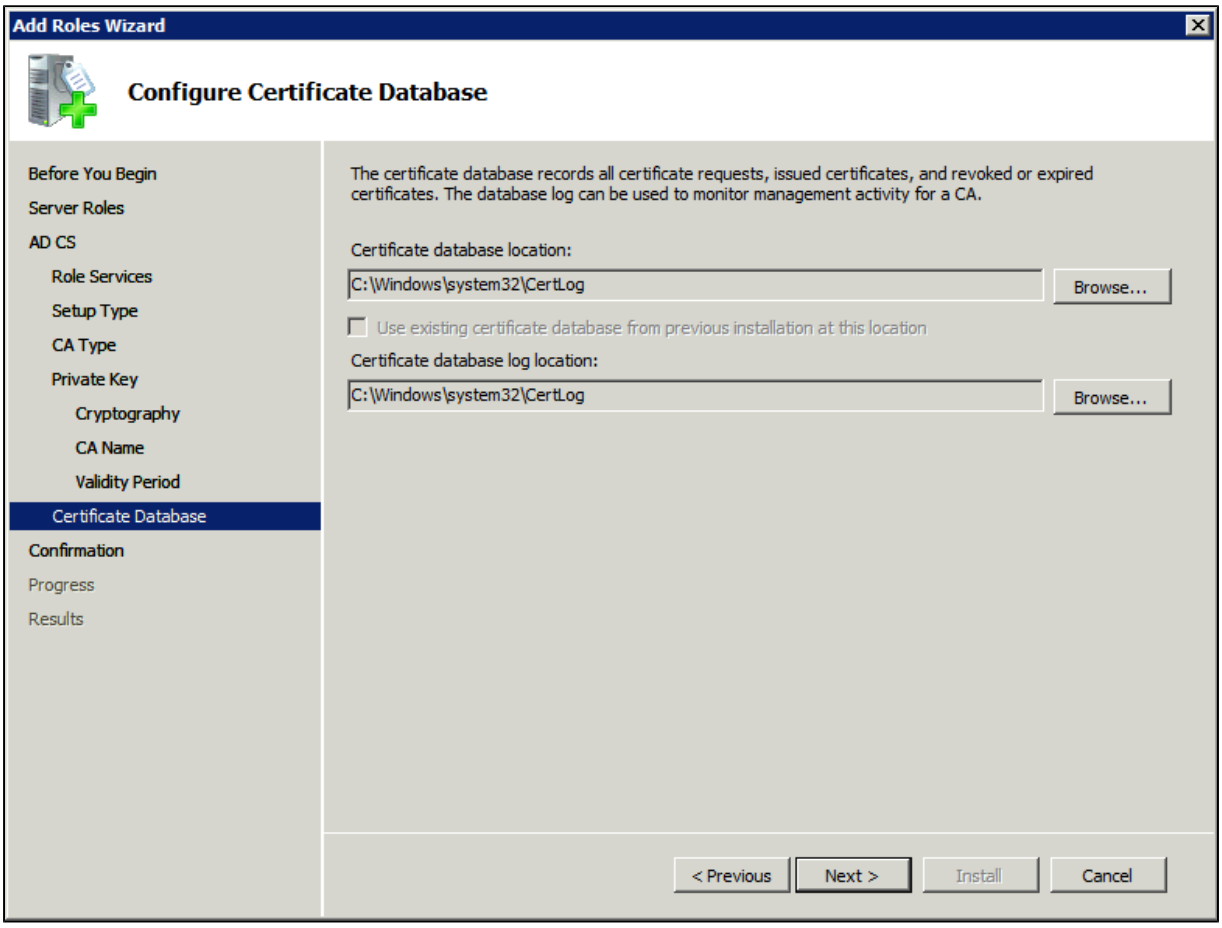

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

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

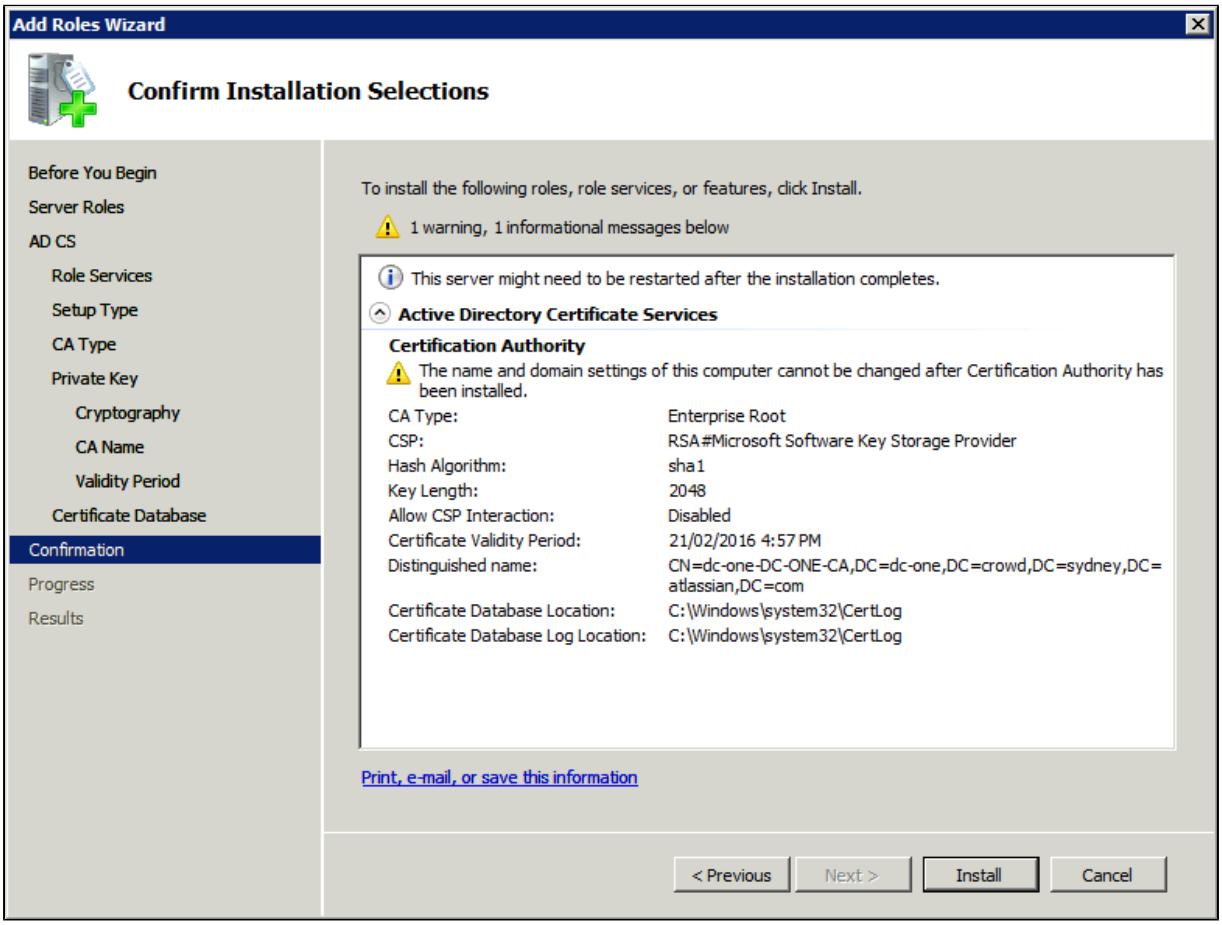

12. Review the information on the results screen to verify that the installation was successful.

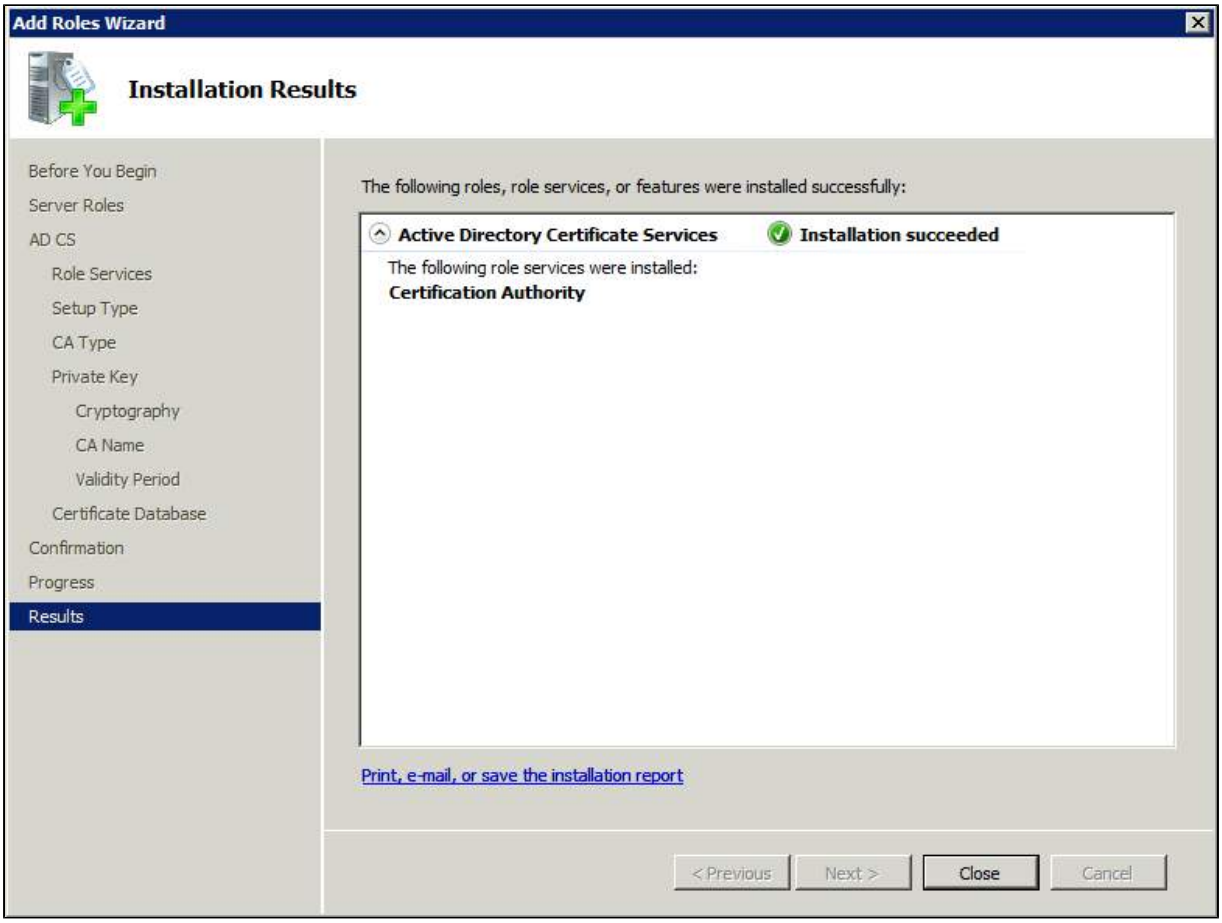

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Step 2. Obtain the Server Certificate

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

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

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

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

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

1. 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. Choose "DER encoded binary X.509 (.CER)" in step-11 ofExporting the LDAPS Certificate and Importing for use with AD DSsection. 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
   ```

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

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

6. You may nowchange '**URL**' to use LDAP over SSL (i.e.ldaps://<HOSTNAME>:636/) anduse the '**Secure SSL**' optionwhen 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
```

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

Mac OS X

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

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

2. Run the command below, where `server-certificate.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
```

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

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

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**
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 synchronized on each subsequent log in. If this box is not checked, the user's login will fail if the user wasn't already manually created in the directory. If you check this box the following additional fields will appear on the screen, which are described in more detail below: • Default Group Memberships • Synchronize Group Memberships • User Schema Settings (described in a separate section below)

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Update User attributes on Login	Whenever your users authenticate to the application, their attributes will be automatically updated from the LDAP server into the application. After you select this option, you won't be able to modify or delete your users directly in the application. • If you need to modify a user, do it on the LDAP server; it will be updated in the application after authenticating. • If you need to delete a user, do it on the LDAP server, but also in the application. If you delete the user only on the LDAP server, it will be rejected from logging in to the application, but it won't be set as inactive, which will affect your license. You'll need to disable the **Update User attributes on Login** option to delete the user, and then enable it again.
Default Group 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, authorization failures will result users will not be able to access the applications or functionality based on the intended group name. Examples: • `confluence-users` • `bamboo-users,jira-administrators,jira-core-users`
Synchronize 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 synchronized with the internal directory each time the user logs in. If you check this box the following additional fields will appear on the screen, both described in more detail below: • Group Schema Settings (described in a separate section below) • Membership Schema Settings (described in a separate section below)

Note: 'Copy Users on Login' must be enabled if you want to be able to change usernames.

Schema Settings

Setting	Description
Base DN	The root distinguished name (DN) to use when running queries against the directory server. Examples: • `o=example,c=com` • `cn=users,dc=ad,dc=example,dc=com` • For Microsoft Active Directory, specify the base DN in the following format: `dc=domain1, dc=local`. You will need to replace the `domain1` and `local` for your specific configuration. Microsoft Server provides a tool called `ldp.exe` which is useful for finding out and configuring the the LDAP structure of your server.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

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*. Nested groups simplify permissions by allowing sub-groups to inherit permissions from a parent group.
Use Paged Results	Enable or disable the use of the LDAP control extension for simple paging of search results. If paging is enabled, the search will retrieve sets of data rather than all of the search results at once. Enter the desired page size that is, the maximum number of search results to be returned per page when paged results are enabled. The default is 1000 results.
Follow Referrals	Choose whether to allow the directory server to redirect requests to other servers. This option uses the node referral (JNDI lookup `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`

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

User Last Name Attribute	The attribute field to use when loading the user's last name. Example: • `sn`
User Display Name Attribute	The attribute field to use when loading the user's full name. Example: • `displayName`
User Email Attribute	The attribute field to use when loading the user's email address. Example: • `mail`

Group Schema Settings

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

Setting	Description
Additional Group DN	This value is used in addition to the base DN when searching and loading groups. If no value is supplied, the subtree search will start from the base DN. Example: • `ou=Groups`
Group Object Class	This is the name of the class used for the LDAP group object. Examples: • `groupOfUniqueNames` • `group`
Group Object Filter	The filter to use when searching group objects. Example: • `(objectCategory=Group)`
Group Name Attribute	The attribute field to use when loading the group's name. Example: • `cn`
Group Description Attribute	The attribute field to use when loading the group's description. Example: • `description`

Membership Schema Settings

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

Setting	Description
Group Members Attribute	The attribute field to use when loading the group's members. Example: • `member`

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

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

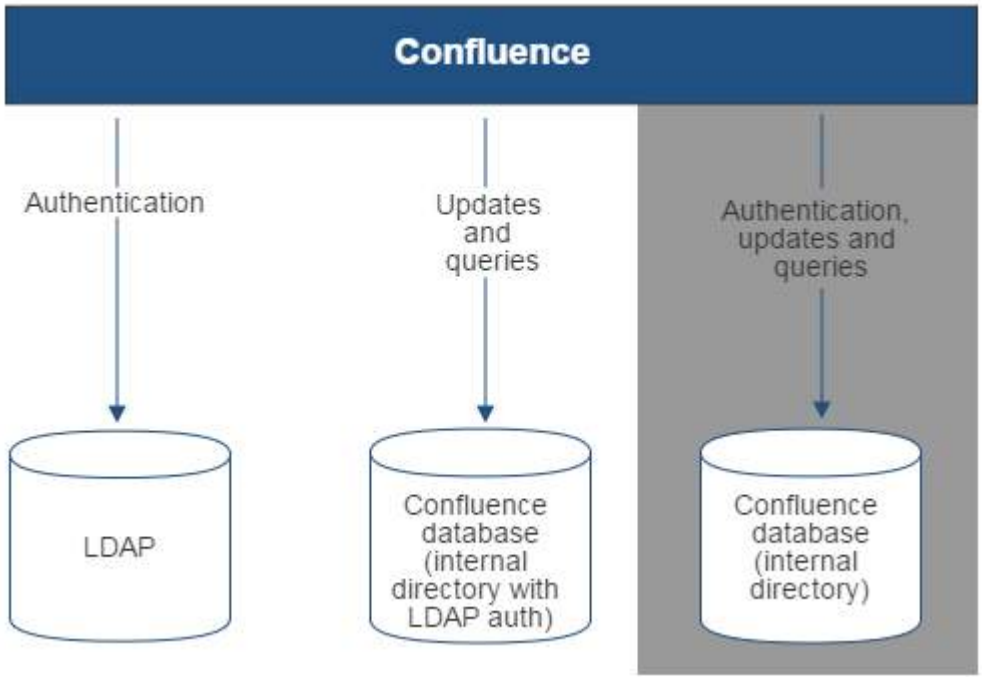

Diagram above: Confluence connecting to an LDAP directory for authentication only.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

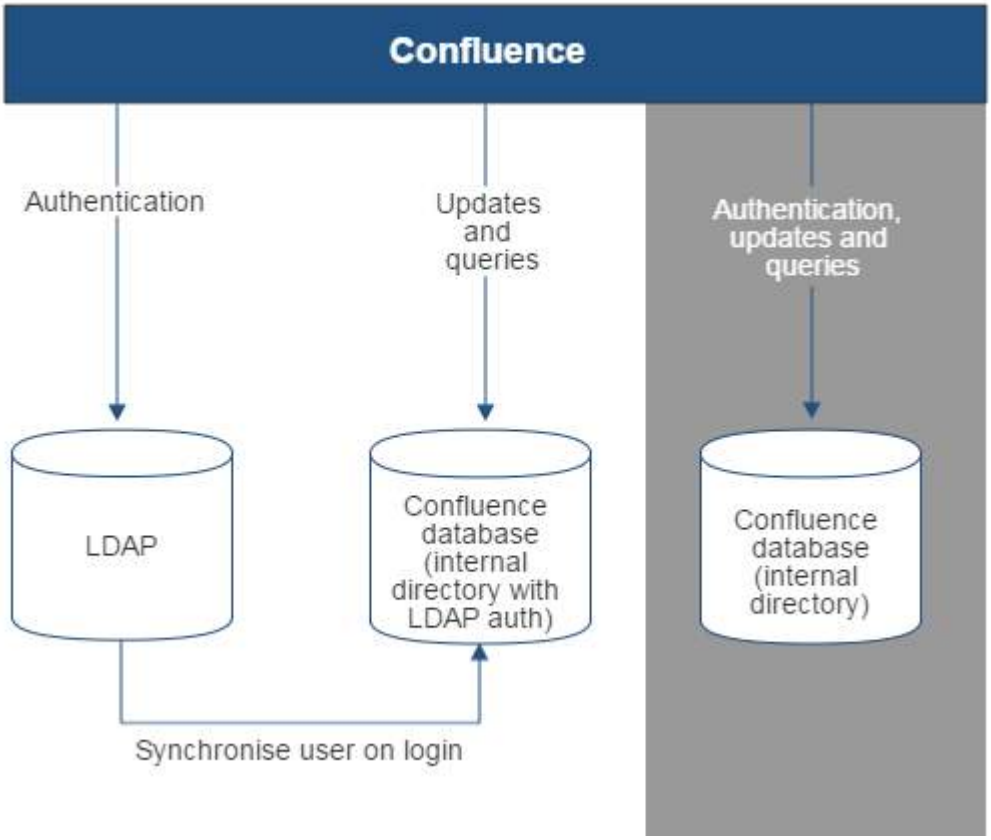

Diagram above: Confluence connecting to an LDAP directory for authentication only, with each user synchronized with the internal directory that is using LDAP authentication when they log in to Confluence.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Connecting to Crowd or Jira for User Management

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

You can't use Jira Cloud for user management.

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

Connecting Confluence to Crowd for User Management

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

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

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

To connect Confluence to Crowd:

1. Go to your **Crowd Administration Console** and define the Confluence application to Crowd. See the Crowd documentation: Adding an Application.
2. Go to > **General Configuration**> **User directories**.
3. **Add** a directory and select type '**Atlassian Crowd**'. Enter the settings as described below.
4. Save the directory settings.
5. Define the **directory order** by clicking the blue up- and down-arrows next to each directory on the '**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.
6. If required, configure Confluence to use Crowd for single sign-on (SSO) too. See the Crowd documentation: Integrating Crowd with Atlassian Confluence.

Crowd Settings in Confluence

Setting	Description

Name	A meaningful name that will help you to identify this Crowd server amongst your list of directory servers. Examples: • `Crowd Server` • `Example Company Crowd`
Server URL	The web address of your Crowd console server. Examples: • `http://www.example.com:8095/crowd/` • `http://crowd.example.com`
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
CONFSERVER-33979 - New JIRA/Crowd password not saved after test GATHERING IMPACT
when configuring Confluence to use Jira/Crowd as a external user directory.

Crowd Permissions

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

Advanced Crowd Settings

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

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

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

Connecting Confluence to Jira applications for User Management

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

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

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

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

1. In your Jira application go to**User Management** > **Jira User Server**.
   (For Jira 6.4 and earlier go to your Jira administration screen then **Users** > **Jira User Server**)
   - Click **Add Application**.
   - Enter the **application name** and **password** that Confluence will use when accessing Jira.
   - Enter the **IP address** or addresses of your Confluence server. Valid values are:
     - A full IP address, e.g. `192.168.10.12`.
     - A wildcard IP range, using CIDR notation, e.g. `192.168.10.1/16`. For more information, see the introduction to CIDR notation on Wikipedia and RFC 4632.
   - **Save** the new application.
2. Set up the Jira user directory in Confluence:
   - Go to > **General Configuration**>**User directories**.
   - **Add** a directory and select type '**Atlassian Jira**'.
   - Enter the settings as described below. When asked for the **application name** and **password**, enter the values that you defined for your Confluence application in the settings on Jira.
   - Save the directory settings.
   - ⚠Don't change the directory order until you have done the next step or you may accidentally lock yourself out of the Confluence admin console.
3. In order to use Confluence, users must be a member of the `confluence-users` group or have Confluence 'can use' permission. Follow these steps to configure your Confluence groups in your JIRA application:
   a. Add the `confluence-users` and `confluence-administrators` groups in your JIRA application.
   b. Add your own username as a member of both of the above groups.
   c. Choose one of the following methods to give your existing JIRA users access to Confluence:
      - Option 1: In your JIRA application, find the groups that the relevant users belong to. Add the groups as members of one or both of the above Confluence groups.
      - Option 2: Log in to Confluence using your JIRA account and go to the Confluence **Administration Console**. Click '**Global Permissions**' and assign the '**can use**' permission to the relevant JIRA groups.
4. In Confluence you can now define the**directory order**by clicking the blue up- and down-arrows next to each directory on the '**User Directories**' screen.Here is a summary of how the directory order affects the processing:
   - The order of the directories is the order in which they will be searched for users and groups.
   - Changes to users and groups will be made only in the first directory where the application has permission to make changes.

   For details seeManaging Multiple Directories.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

⚠ Ensure that you have added Confluence URL into Jira Whitelist in `Jira Administration >> System >> Security >> Whitelist`. For example: https://confluence.atlassian.com/ or refer to this guide: Configuring the whitelist.

Jira Settings in Confluence

Setting	Description
Name	A meaningful name that will help you to identify this Jira server in the list of directory servers. Examples: • `Jira Software` • `My Company Jira`
Server URL	The web address of your Jira server. Examples: • `http://www.example.com:8080` • `http://jira.example.com`
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 the 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 the inheritance of permissions from one group to its sub-groups.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Update group memberships when logging in	This setting enables updating group memberships during authentication and can be set to the following options:
	• **Every time the user logs in**: during the authentication, the users **direct** group memberships will be updated to match whats in the remote directory:
	  • Remove the user from all groups that the user no longer belongs to in the remote directory.
	  • Add the user to all the groups that the user belongs to in the remote directory. New groups with matching names and descriptions will be created locally if needed. The group will only contain the current user and other memberships will be populated when users who belong to the same group log in or when the synchronization happens.
	• **For newly added users only**: when a new user logs in for the first time, the users **direct** group memberships will be updated to match whats in the remote directory.
	  ⓘ Consider that the user's group memberships will be updated only if the user was created during the authentication.
	• **Never**: during the authentication, the user's group memberships wont change, even if the local state doesnt match whats in the remote.

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

Diagrams of Some Possible Configurations

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

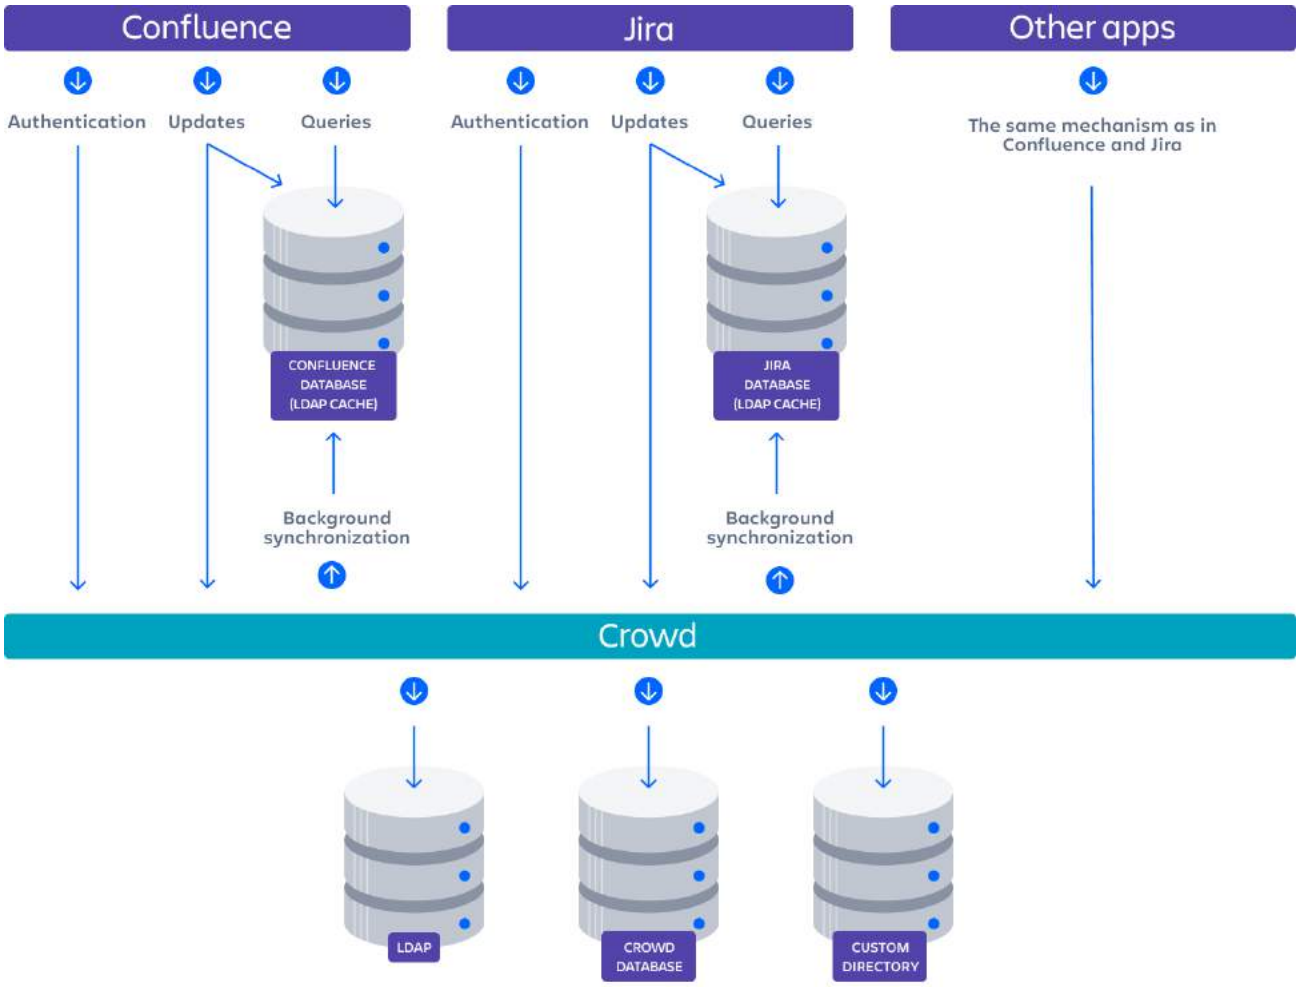

Diagram: Confluence, Jira and other applications connecting to Crowd for user management.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

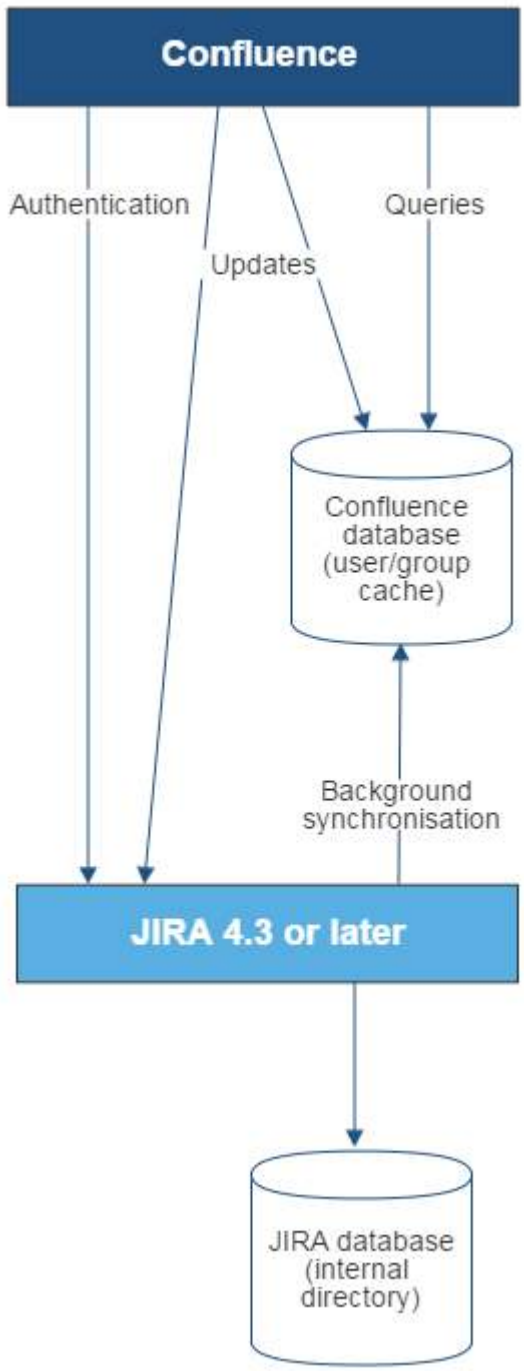

Diagram above: Confluence connecting to JIRA for user management.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

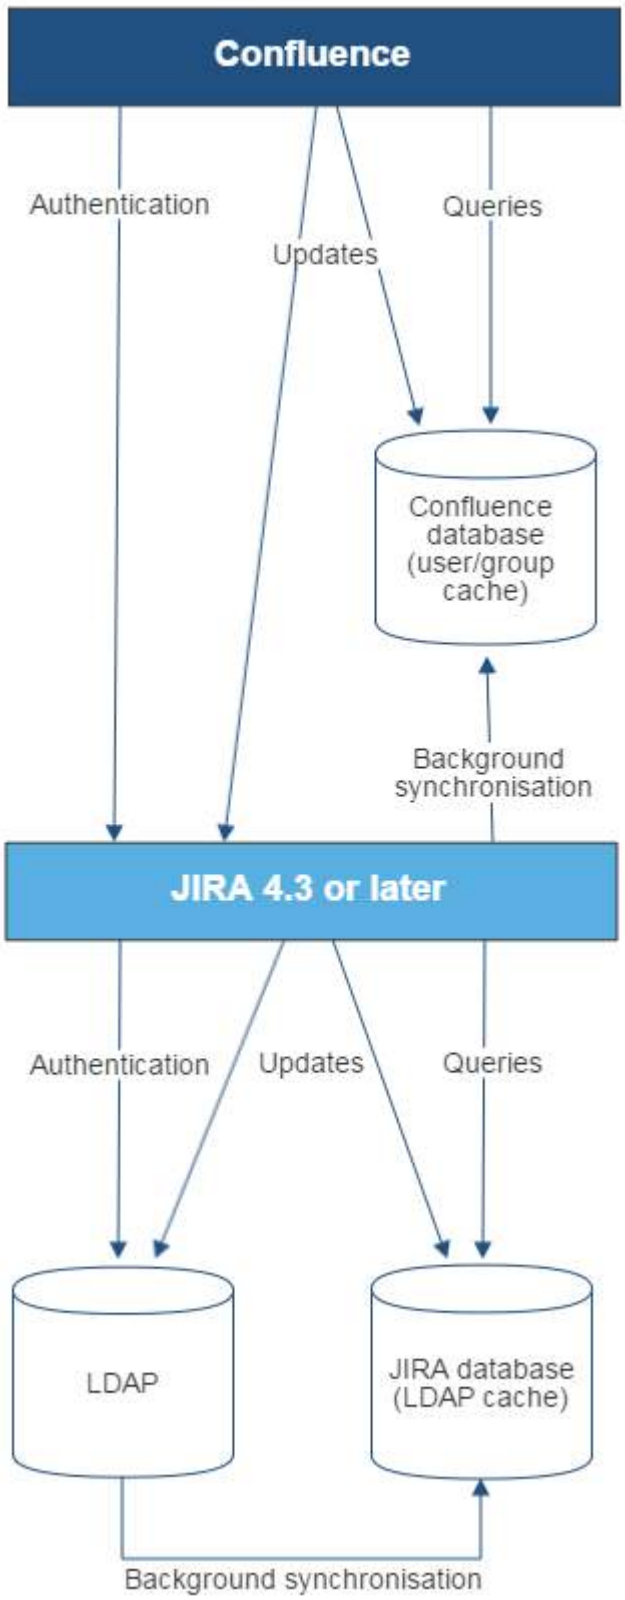

Diagram above: Confluence connecting to JIRA for user management, with JIRA in turn connecting to LDAP.

Troubleshooting

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

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

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.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Reverting from Crowd or Jira applications to Internal User Management

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

On this page:

- Option 1 Manually Recreate Users and Groups in Confluence
- Option 2 Transfer Crowd/Jira application Users and Groups to the Confluence Database

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

Option 1 Manually Recreate Users and Groups in Confluence

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

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

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

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

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

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

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

Step 1. Create Backups

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

1. From Confluence, create a full XML site backup including attachments.
2. Stop Confluence.
3. Make a backup copy of the Confluence home and installation directories.
4. Repeat the above steps for your External Application.
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>';
```

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

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

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

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

   Make a note of each of the usernames and groups returned. You'll need this in the next step.
2. In your **external directory**, remove the users from their respective groups. Their membership will still be retained in the Confluence internal directory.
3. Run the SQL query above again. Once it returns no results, you can move to the next step.

Move Groups to Confluence

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

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

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

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

```
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

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

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

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

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

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

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

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

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

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

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

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

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

Reset passwords

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

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Managing Multiple Directories

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

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

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

On this page:

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

Overview

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

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

Configuring the Directory Order

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

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

Effect of Directory Order

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

Login

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

Permissions

Aggregating membership (default)

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

Example:

- You have connected two directories: The Customers directory and the Partners directory.
- The Customers directory is first in the directory order.
- A 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*director ies, not the first one the user appears in. In most cases, this change will have no impact as users generally only exist in one directory, or their memberships are correctly synchronized between user directories. In some rare cases, where group memberships are out of synch, the change may lead to users gaining permissions to view spaces and pages (if they are a member a group in a user directory that was previously being ignored by Confluence).

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

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

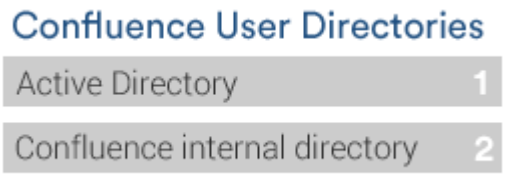

and Issac's group memberships in each directory looks like this:

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

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

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

To Confluence his group membership looks like this:

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

Non-aggregating membership

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

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

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

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

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

Example:

- You have connected two directories: The Customers directory and the Partners directory.
- The Customers directory is first in the directory order.
- A username`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.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

- 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.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Managing Nested Groups

Some directory servers allow you to define a group as a member of another group. Groups in such a structure are called *nested groups*. Nested groups simplify permissions by allowing sub-groups to inherit permissions from a parent group.

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

On this page:

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

Related pages:

- Configuring User Directories

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

Effect of Nested Groups

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

Login

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

Permissions

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

Viewing lists of group members

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

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

Adding and updating group membership

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

If you try to remove a user from a flattened list, the following will happen:

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

Examples

Example 1: User is member of sub-group

Imagine the following two groups exist in your directory server:

- staff
- marketing

Memberships:

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

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

Example 2: Sub-groups as members of the jira-developers group

In an LDAP directory server, we have the groups **engineering-group**and **techwriters-group**. We want to grant both groups developer-level access to the JIRA. We will have a group called**jira-developers**that has developer-level access.

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

Group memberships are now:

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

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

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

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

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

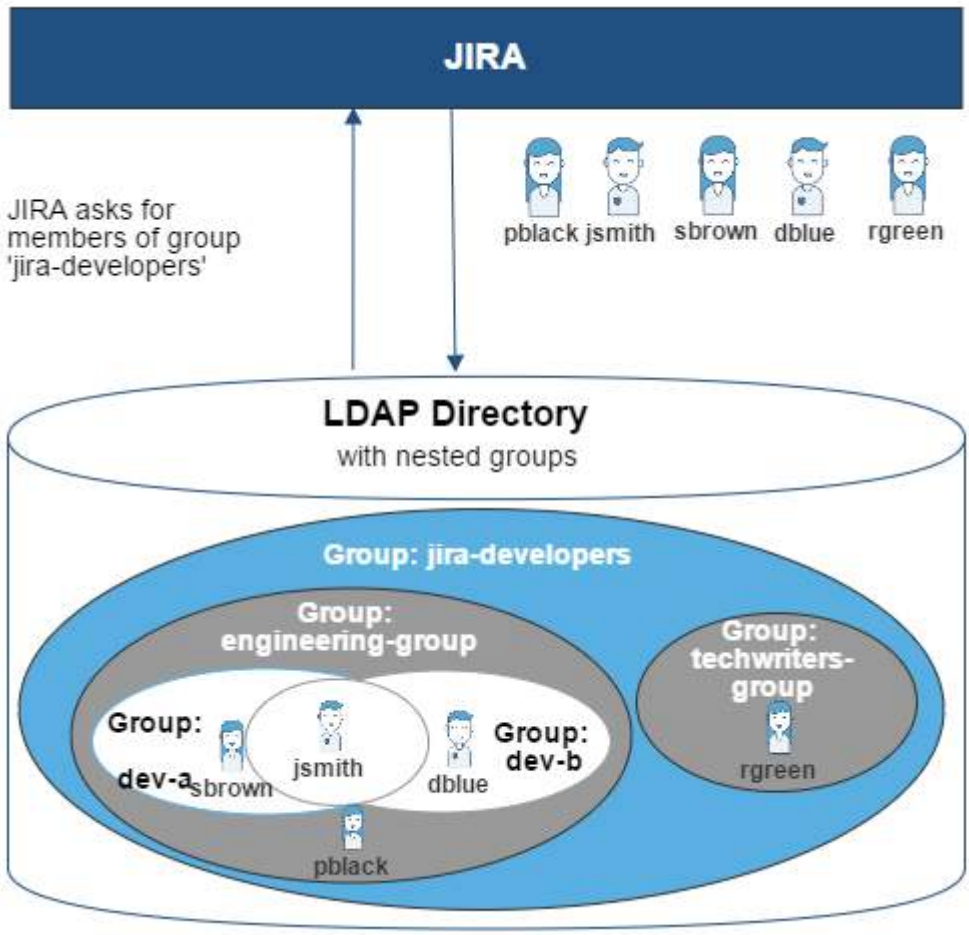

Notes

- **Possible impact on performance.** Enabling nested groups may result in slower user searches.
- **Definition of nested groups in LDAP.** In an LDAP directory, a nested group is a child group entry whose DN (Distinguished Name) is referenced by an attribute contained within a parent group entry. For example, a parent group **Group 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
```


Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Synchronizing Data from External Directories

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

On this page:

- Affected Directory Types
- How it Works
- Finding the Time Taken to Synchronize
- Manually Synchronizing the Cache
- Configuring the Synchronization Interval
- Unsynced users

Related pages:

- Configuring User Directories

Affected Directory Types

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

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

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

- **Internal Directory with LDAP Authentication**.
- **Internal Directory**.

How it Works

Here is a summary of the caching functionality:

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

Finding the Time Taken to Synchronize

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

Manually Synchronizing the Cache

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

Screen snippet: User directories, showing information about synchronization

Configuring the Synchronization Interval

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

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

The length you choose for your synchronization interval depends on:

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

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

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

Unsynced users

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

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

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

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

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

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

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

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

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

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

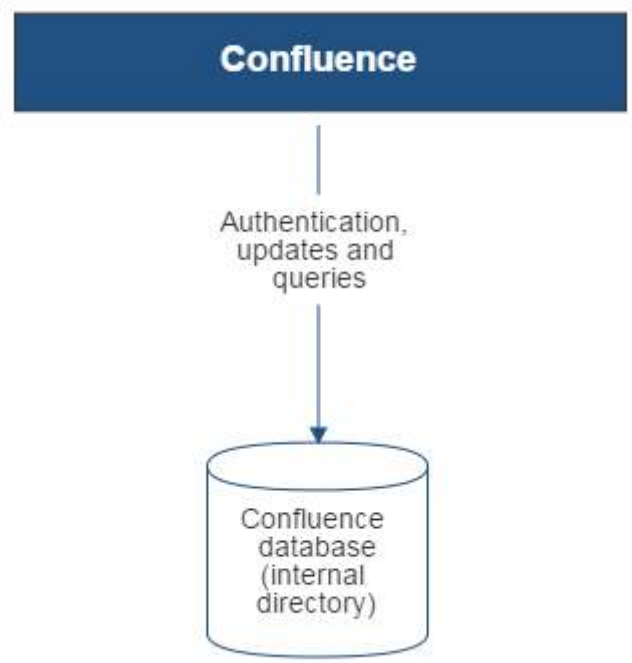

Diagram above: Confluence using its internal directory for user management.

Confluence with Read/Write Connection to LDAP

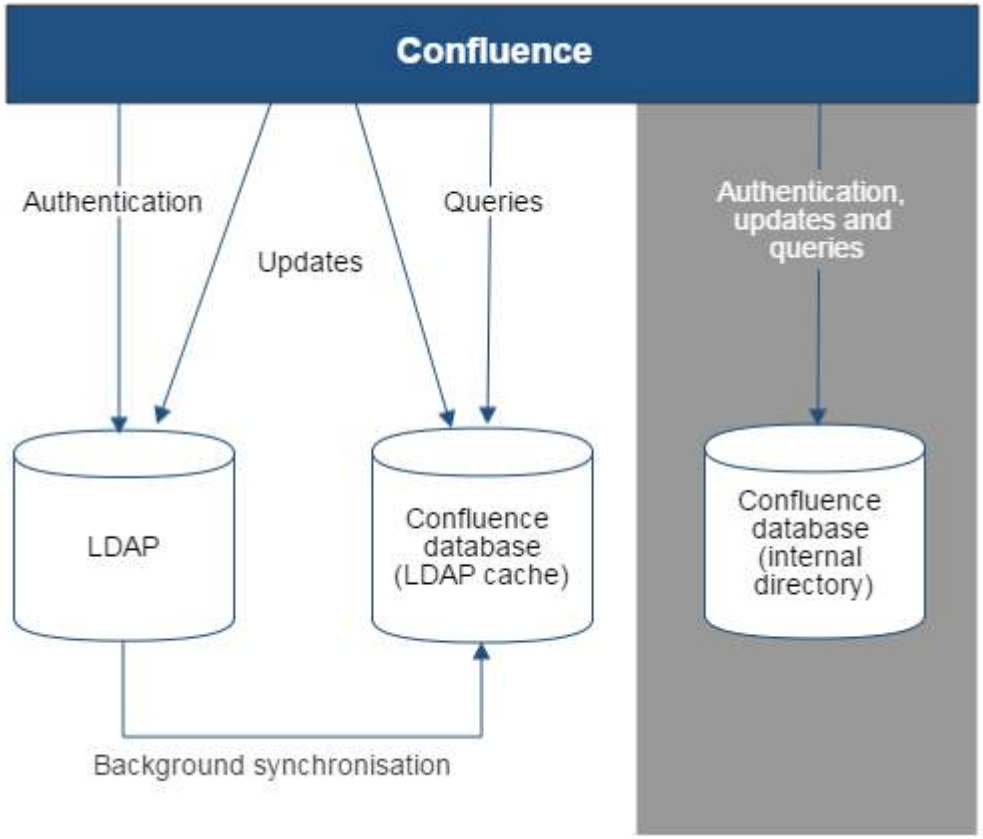

Diagram above: Confluence connecting to an LDAP directory.

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

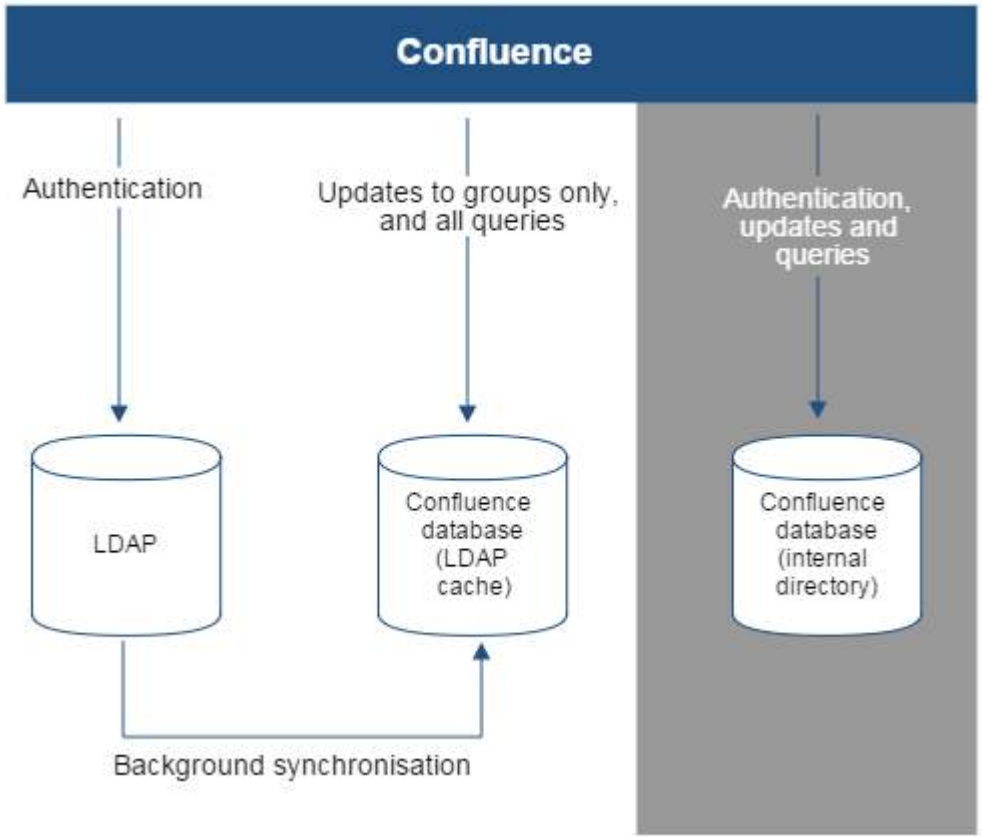

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

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

Confluence Internal Directory with LDAP Authentication

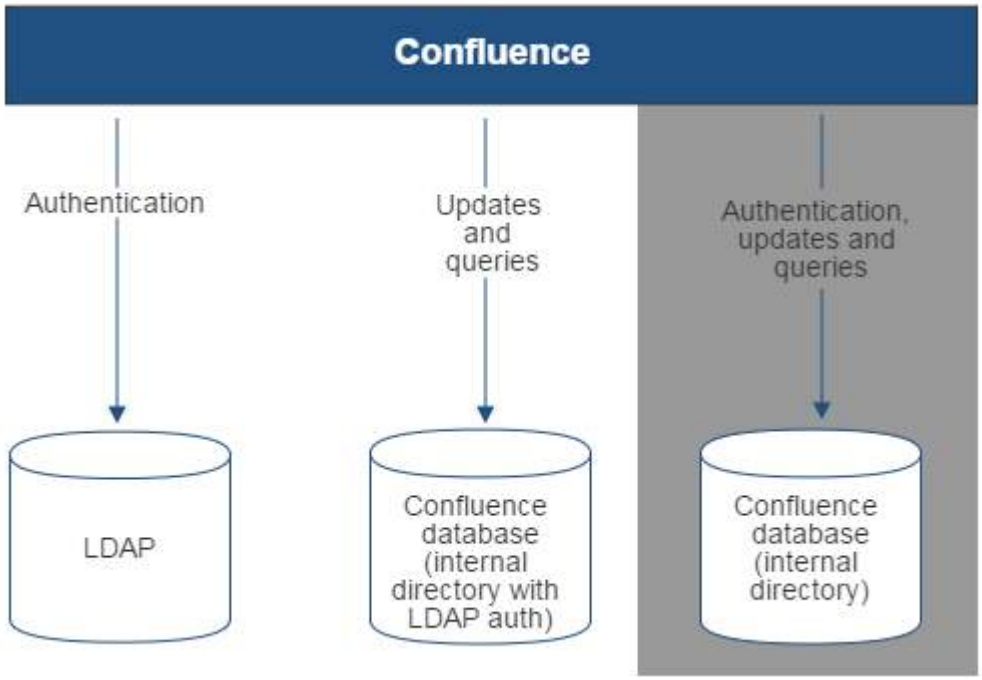

Diagram above: Confluence connecting to an LDAP directory for authentication only.

Confluence with LDAP Authentication, Copy Users on First Login

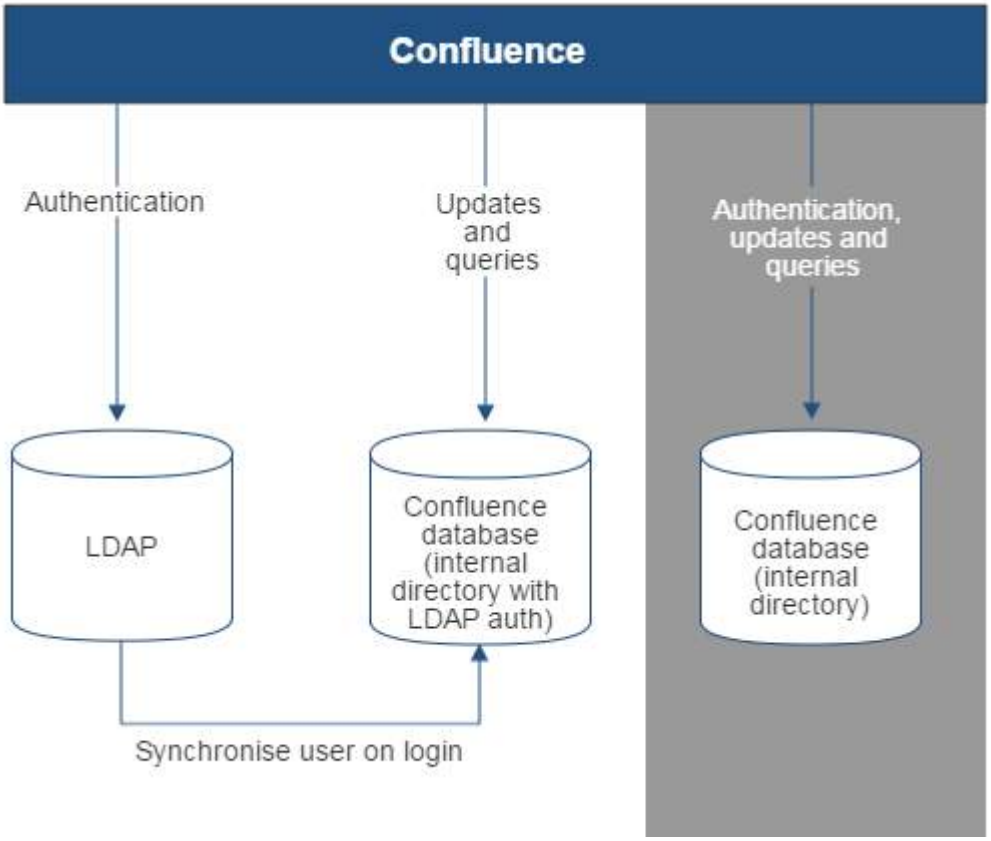

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

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

Confluence Connecting to Jira

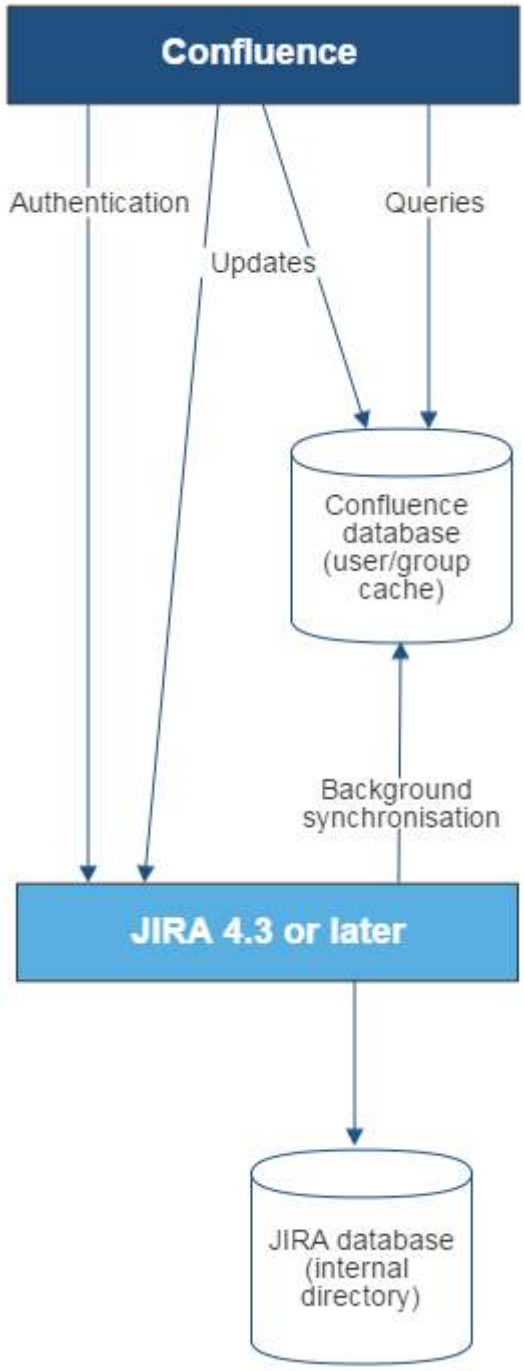

Diagram above: Confluence connecting to JIRA for user management.

Confluence Connecting to Jira and Jira Connecting to LDAP

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

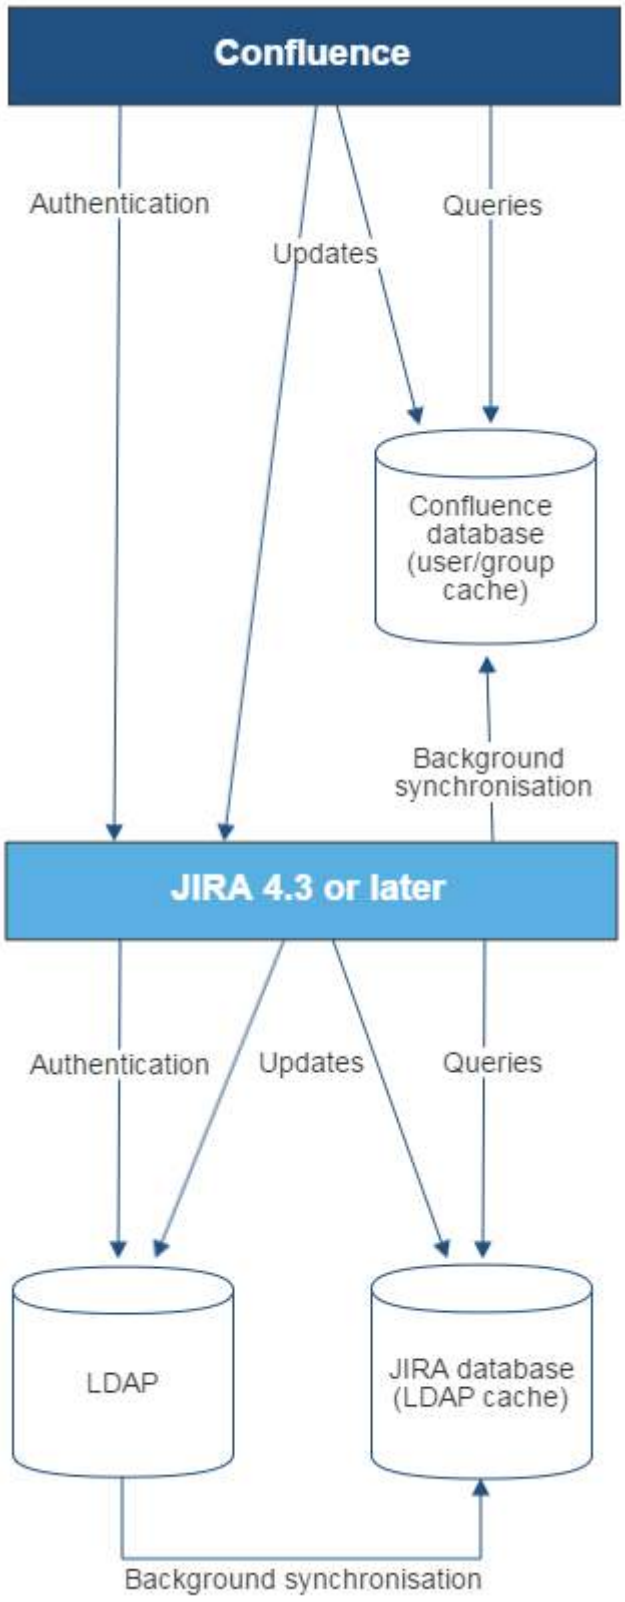

Diagram above: Confluence connecting to JIRA for user management, with JIRA in turn connecting to LDAP.

Confluence and Jira Connecting to Crowd

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

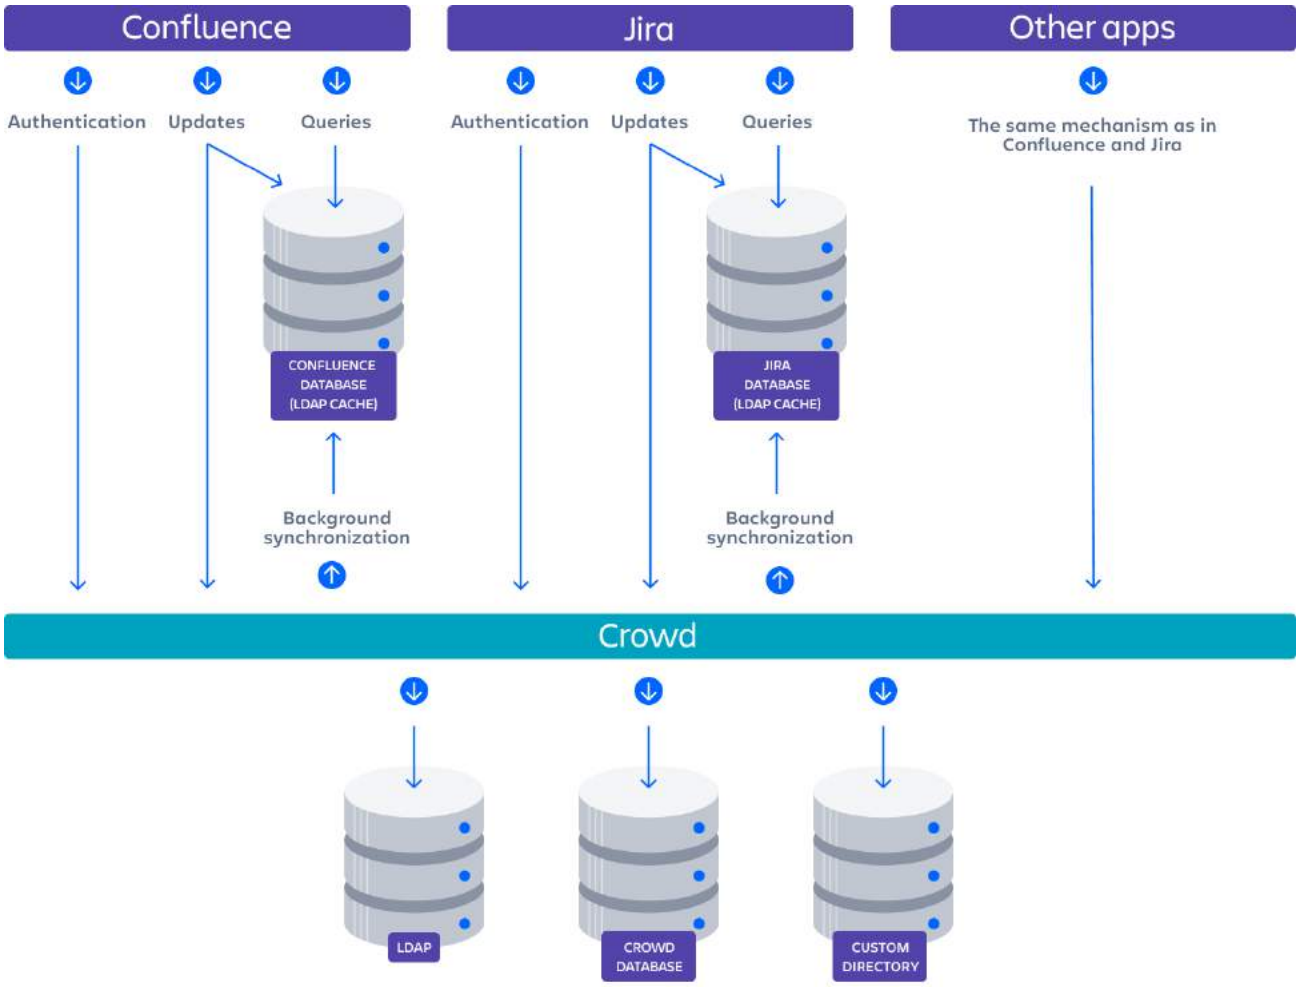

Diagram: Confluence, Jira and other applications connecting to Crowd for user management.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

User Management Limitations and Recommendations

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

On this page:

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

Related pages:

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

General Recommendations

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

Be careful when deleting users in remote directories.

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

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

Avoid hash, slash and question characters in usernames

There is a known issue where users with #, ? or / in their username cannot create spaces. See CONFSERVER-43494 GATHERING IMPACT and CONFSERVER-13479 GATHERING IMPACT for more information.

Recommendations for Connecting to LDAP

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

Optimal Number of Users and Groups in your LDAP Directory

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

This recommendation affects connections to LDAP directories:

- Microsoft Active Directory

- All other LDAP directory servers

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

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

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

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

Our Test Results

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

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

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

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

Redundant LDAP is Not Supported

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

Specific Notes for Connecting to Active Directory

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

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

1. **Externally moving objects out of scope or renaming objects causes problems in AD.** If you move objects out of scope in AD, this will result in an inconsistent cache. We recommend that you do not use the external LDAP directory interface to move objects out of the scope of the sub-tree, as defined on the application's directory configuration screen. If you do need to make structural changes to your LDAP directory, manually synchronize the directory cache after you have made the changes to ensure cache consistency.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

2. **Synchronizing between AD servers is not supported.** Microsoft Active Directory does not replicate the uSNChanged attribute across instances. For that reason, we do not support connecting to different AD servers for synchronization. (You can of course define multiple different directories, each pointing to its own respective AD server.)
3. **Synchronizing with AD servers behind a load balancer is not supported.** As with synchronizing between two different AD servers, Microsoft Active Directory does not replicate the uSNChanged attribute across instances. For that reason, we do not support connecting to different AD servers even when they are load balanced. You will need to select one server (preferably one that is local) to synchronize with instead of using the load balancer.
4. **You must restart the application after restoring AD from backup.** On restoring from backup of an AD server, the uSNChanged timestamps are reverted to the backup time. To avoid the resulting confusion, you will need to flush the directory cache after a Active Directory restore operation.
5. **Obtaining AD object deletions requires administrator access.** Active Directory stores deleted objects in a special container called cn=Deleted Objects. By default, to access this container you need to connect as an administrator and so, for the synchronization task to be aware of deletions, you must use administrator credentials. Alternatively, it is possible to change the permissions on the cn=Deleted Objects container. If you wish to do so, please see this Microsoft KB article.
6. **The User DN used to connect to AD must be able to see the uSNChanged attribute.**The synchronization task relies on the uSNChanged attribute to detect changes, and so must be in the appropriate AD security groups to see this attribute for all LDAP objects in the subtree.

Recommendations for Connecting to Jira for User Management

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

Single Sign-On Across Multiple Applications is Not Supported

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

Custom Application Connectors are Not Supported

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

Custom Directories are Not Supported

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

Load on your JIRA instance

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

JIRA Cloud applications not supported

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

Recommendations

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

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

• You do not need single sign-on (SSO) between your JIRA application and Confluence, or between two JIRA servers. • You do not have custom application connectors. Or, if you do have them, you are happy to convert them to use the new REST API. • You are happy to shut down all your servers when you need to upgrade your JIRA application.	
If **one or more** of the following are true: • If your JIRA application is already under high load. • You want to share user and group management across more than 5 applications. • You need single sign-on (SSO) across multiple applications. • You have custom applications integrated via the Crowd SOAP API, and you cannot convert them to use the new REST API. • You are not happy to shut down all your servers when you need to upgrade JIRA.	We recommend that you install Atlassian Crowd for user management and SSO.
If you are considering creating a custom directory connector to define your own storage for users and groups...	Please see if one of the following solutions will work for you: • If you have written a custom provider to support a specific LDAP schema, please check the supported LDAP schemas to see if you can use one of them instead. • If you have written a custom provider to support nested groups, please consider enabling nested groups in the supported directory connectors instead. • If you have written a custom provider to connect to your own database, please consider loading the data into the application's database instead. • If you need to keep the custom directory connection, please consider whether Atlassian Crowd meets your requirements. See the documentation on Creating a Custom Directory Connector.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Requesting Support for External User Management

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

The cause of such problems may be:

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

On this page:

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

Related pages:

- Troubleshooting Problems and Requesting Technical Support
- Configuring User Directories

Troubleshooting the Connection to your External User Directory

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

To test your directory connection:

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

Please refer to our knowedge base articles for troubleshooting user management and login issues.

If the above resources do not help, continue below.

Problems During Initial Setup

Raise a support request and include the following information.

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

Complex Authentication or Performance Problems

Raise a support request and include the following information.

Confluence Server

Log in to Confluence and access the Administration Console.

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

Confluence Configuration Files

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

User Management System

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

Diagnostics

- Enable profiling. See Performance Tuning.
- Enable detailed user management logging, by editing `confluence/WEB-INF/classes/log4j.properties`.
  Change this section:

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

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

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

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

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Disabling the Built-In User Management

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

You might disable Confluence's internal user management:

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

To disable management of users and groups within Confluence:

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

Single sign-on for Confluence Data Center

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

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

On this page:

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

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

Looking for a cross-domain SSO solution?

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

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

Managing System and Marketplace Apps

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

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

There are two main types of apps:

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

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

About the Universal Plugin Manager

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

The UPM allows you to:

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

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

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

SeetheUniversal Plugin Manager documentationfor more information on using the UPM.

Disable and uninstall apps

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

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

Go to > **General Configuration** > **Macro Usage**.

Writing User Macros

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

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

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

On this page:

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

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`$body`variable. Options for processing the macro body include: • **No macro body** Select this option if your macro does not have a body. • **Escaped** Confluence will add escape characters to the HTML markup in the macro body. Use this if you want to show actual HTML markup in the rendered page. For example, if the body is `<b>Hello World</b>` it will render as `<b>Hello World</b>`. • **Unrendered** HTML in the body will be processed within the template before being output. Ensure that HTML is ultimately output by the template. • **Rendered** Confluence will recognize HTML in the macro body, and render it appropriately. For example, if the body is `<b>Hello World</b>` it will render as **Hello World.**
Template	This is where you write the code that determines what the macro should do. • Use HTML and Confluence-specific XML elements in the macro template. • You can use theVelocitytemplating language. Here is more information onthe Velocity project. • If your macro has a body, your template can refer to the macro body text by specifying '$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:

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

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

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

Best practices

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

Add a descriptive header to your macro template

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

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

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

Expose your parameters in the macro browser

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

Supply default values for macro parameters

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

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

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

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

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

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

Consider security implications

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

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

Example user macros

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

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

Template	Enter the code below in the template field - this example will print the text straight onto the page.
	`## @noparams`
	`Hello World!`
	`$body`
	
	If you wanted the text to appear in a panel you could include the relevant AUI message class as shown here.
	`## @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.

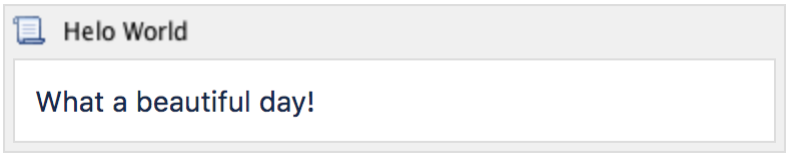

The result is:

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

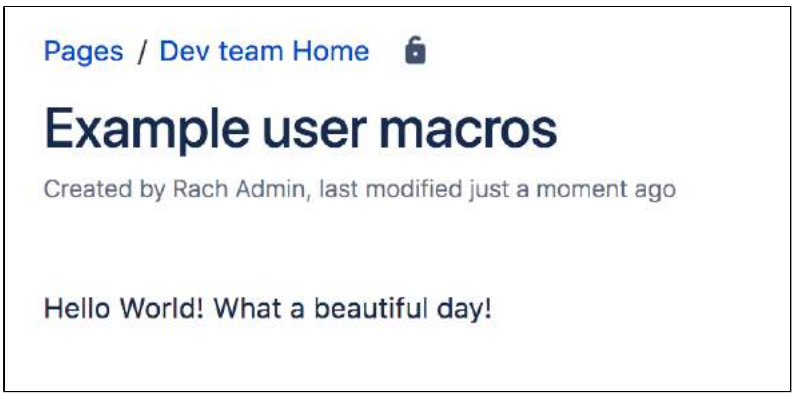

This example demonstrates how to create a user macro that can contain text that is visible when viewing a page, but does not print.

Field	Value
Macro name	`noprint`
Visibility	`Visible to all users in the Macro Browser`
Macro Title	`No Print`
Description	`Hides text from printed output.`
Categories	`Confluence Content`
Icon URL	You can leave this field blank
Documentation URL	You can leave this field blank
Macro body processing	`Rendered`

Template	Enter the code below in the template field.
	`## @noparams`
	`<div class="noprint">$body</div>`

Using the 'NoPrint' Macro on a page

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

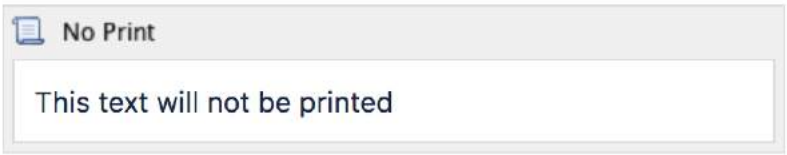

Making the PDF export recognize the NoPrint macro

SeeAdvanced PDF Export Customizations.

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

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

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>`

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

(Title)

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

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

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

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 macro. 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.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

`<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 color, border style, border color, border width and title color.
	
	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.

Do more with Confluence

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

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

Next Steps

Explore the power of user macros further by reading our Writing Advanced User Macrosguide.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

User Macro Template Syntax

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

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

On this page:

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

Related pages:

- Writing User Macros

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`hellow orld`.
Enter the following code in your template:

```
Hello World: $body
```

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

```
From Matthew
```

The wiki page will display the following:

> *Hello World: From Matthew*

Using parameters in your user macro

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

How your macro parameters are used on a Confluence page

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

Defining the parameters

A parameter definition in the template contains:

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

Format:

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

Additional notes:

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

The sections below describe each of the attributes in detail.

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

Parameter name

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

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

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

Parameter type

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

Parameter type	Description
boolean	Displays a checkbox to the user and passes the value 'true' or 'false' to the macro as a string.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

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

Using the parameters in your macro code

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

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

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

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.

Security consideration

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

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

Rendering HTML with variables

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

Example: `$outputHtml` instead of `$output`

Controlling parameter appearance in the editor placeholder

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

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

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

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

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

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

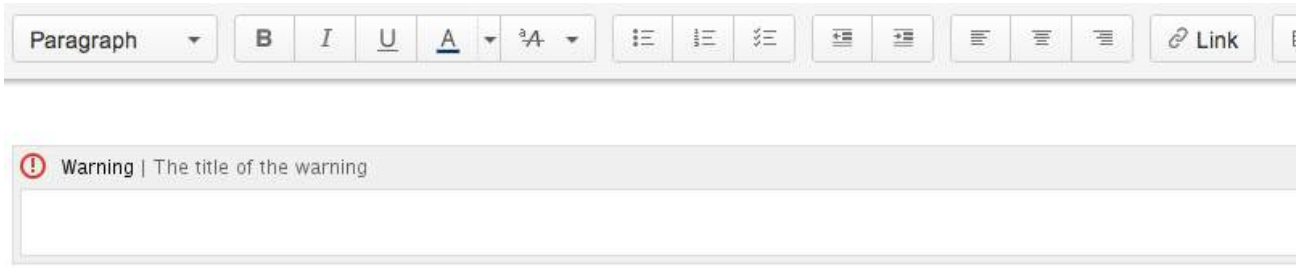

To configure the macro placeholder for a user macro, you will add attributes to the @param entry in the template.

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

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

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

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

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

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

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Customizing your Confluence Site

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

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

We've documented the customizations under two broad headings:

- You can change the **appearance** of Confluence by customizing the dashboard, adjusting the colors, adding a site logo, and more. See Changing the Look and Feel of Confluence.
- You can determine the **default behavior** by setting various options, or define the **default content** that appears in new spaces, on the dashboard, and in other Confluence locations. See Changing the Default Behavior and Content in Confluence.

Related pages:

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

Changing the Look and Feel of Confluence

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

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

Related pages:

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

Ways to customize the look and feel of your site:

- Add your own site logo. See Changing the Site Logo.
- Change the **color scheme** of the user interface. See Customizing Color Schemes.
- Use**themes**for advanced layout customization. SeeWorking with Themes.
- Change the **site or space layouts**, which determine how the controls are laid out in the site. This does not change the actual page layouts, but it does change the way the surrounding controls appear in the page. SeeCustomizing Site and Space Layouts.

Customizing the Confluence Dashboard

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

On this page:

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

Related pages:

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

Editing the site welcome message

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

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

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

Using a page as the site landing page

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

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

Advanced customizations

You can further customize the dashboard by editing the global layout file. SeeCustomizing Site and Space Layoutsfor more information on how to do this. You'll need some knowledge ofVelocityto 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 customizations are, the harder it may be to reapply the changes when upgrading. SeeUpgrading Customized Site and Space Layoutsto find out what will be involved before modifying the layouts.

Changing the Site Logo

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

You can change:

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

On this page:

- Changing the site logo
- Changing the site icon (favicon)
- Changing the default space logo
- Changing a specific space logo

Related pages:

- Changing the Look and Feel of Confluence

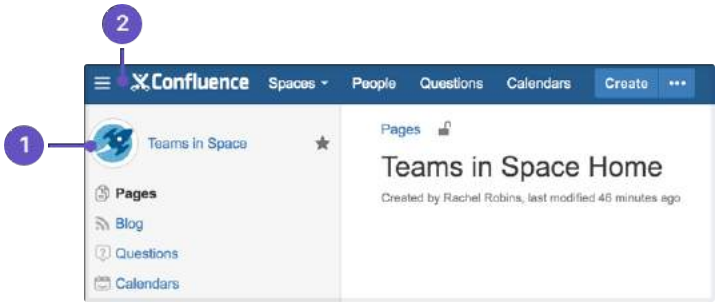

1. **Space logo:**appears in the sidebar and on the dashboard.
2. **Site logo:**always visible, click the logo to go to the dashboard (or site homepage).

Changing the site logo

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

To change the site logo:

1. Choose the**cog icon** , then choose**General Configuration**
2. Choose **Site Logo and Favicon**.
3. Choose **Browse** to upload a new logo.
4. Choose**Show Logo Only** or **Show Logo and Title** depending on whether you wish the Site Title to display in the header.
5. Choose **Save**.

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

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

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

Changing the site icon (favicon)

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

1. Go to > **General Configuration**>**Site Logo and Favicon**.
2. Locate your image file and choose**Upload**.

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

Changing the default space logo

The Space Logo appears in the sidebar and as an icon in the Sites Directory.The default space logo applies to all spaces that do not have a custom space logo applied - see Configure the Sidebar.

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

To change the default space logo:

1. Go to > **General Configuration**>**Default Space Logo**.
2. Choose **Logo:ON**
3. Choose**Browse**to upload a new logo
4. Choose **Upload Logo**
5. Choose**Save**.

Changing a specific space logo

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

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

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Customizing Color Schemes

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

To change the site's color scheme:

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

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

On this page

- Reset your color scheme after uploading a site logo

Related pages:

- Changing the Look and Feel of Confluence

Reset your color scheme after uploading a site logo

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

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

To reset the color scheme:

1. Choose the**cog icon** , then choose**General Configuration**
2. Choose **Color Scheme** in the left-hand panel
3. Click**Edit**
4. Hit**Reset**

Styling Confluence with CSS

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

On this page:

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

Related pages:

- Basic Styling Tutorial
- Styling Fonts in Confluence

Introduction

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

Note: By default, only system administratorscan edit the CSS for a space or for the site.To allow any user with Space Admin permissions to edit the CSS for a space, go to > **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 customizing web applications like Confluence. It is important to test each change that you make and ensure it works as expected in all areas of Confluence for example, on the Confluence dashboard as well as on regular pages.

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

Considerations for Using Custom CSS

CSS Knowledge is Required

If you are not familiar with CSS, see the links in the CSS Resources sectionbelow. 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 the various web browsers supported by Confluence.

CSS Customization is Not Supported

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

Getting Started

Editing the CSS

To edit a space's CSS style sheets:

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

To edit your global CSS stylesheet:

1. Choose > **General Configuration**> **Stylesheet**.
2. Choose **Edit**.
3. Paste your custom CSS into the text field.
4. Choose **Save**.

Note:

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

Follow the Tutorial

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

CSS Resources

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

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Basic Styling Tutorial

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

On this page:

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

Related pages:

- Styling Confluence with CSS

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.

Tutorial: Changing the Header Background

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

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

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

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

CSS Editing Tips

Begin With a Space Stylesheet

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

Use the Right Tools

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

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

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

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

Edit Simple Elements First

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

Notes

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

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Styling Fonts in Confluence

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

Related pages:

- Basic Styling Tutorial
- Styling Confluence with CSS

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

Changing the fonts

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

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

To edit a space's stylesheet:

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

```
.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-family: Times, "Times New Roman", serif;
 font-size: 14px;
}
```

Notes

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

Working with Themes

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

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

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

To see the themes installed in your site:

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

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

Note about the Documentation theme

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

Related pages:

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

Applying a Theme to a Site

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

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

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

Related pages:

- Apply a Theme to a Space

Creating a Theme

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

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

Related pages:

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

Customizing Site and Space Layouts

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

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

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

On this page:

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

Related pages:

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

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

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

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

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

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

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

Editing a site decorator file

To edit a site decorator:

1. Go to > **General Configuration**> **Layouts** (under **Look and Feel**)
2. Click **Create Custom**next to the decorator `.vmd`fileyou 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 customizations

Overriding Velocity templates

The `velocity` directory is at the front of Confluence's Velocity template search path. As such, you can override *any* of Confluence's Velocity templates by placing an identically named file in the right place. While we don't recommend you do this unless you know exactly what you're doing, it does give you complete control over the look of every aspect of Confluence. It also means that you can edit your templates in a 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.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Upgrading Customized Site and Space Layouts

As Confluence evolves, so do the default site and space layouts that drive the rendering of every page. As new functionality is added or current functionally is changed, the default layouts are modified to support these changes.

Related pages:

- Customizing Site and Space Layouts

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

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

To reapply your custom layouts, you need to:

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

Step 1. Obtain your Custom Layouts

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

If not, you should be able to find your customizations using the following method. This method extracts all site- and space-level layouts from your Confluence site as a single output. From this output, you should be able to identify your customizations.

✓ This method is handy to use if you have:

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

Custom layouts are stored in the `DECORATOR` table within your Confluence database. You can `SELECT` for the source of the layout using SQL like this:

```
mysql> select SPACEKEY,DECORATORNAME,BODY from DECORATOR;
+----------+--------------------+------+
| SPACEKEY | DECORATORNAME      | BODY |
+----------+--------------------+------+
| NULL     | decorators/main.vmd | ...  |
+----------+--------------------+------+
1 row in set (0.03 sec)
```

*This example was tested on **MySQL**, but should be applicable to all SQL databases.*

Step 2. Reapply your Customizations

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

What are 'major' and 'minor' releases? Major release upgrades are ones where the 1st digit of Confluence's version number or the 1st digit after the 1st decimal place differ after the upgrade, for example, when upgrading from Confluence 3.0 to 3.1, or 2.8 to 3.0. Minor release upgrades are ones where the 1st digit of Confluence's version number and the 1st digit after the 1st decimal place remain the same after the upgrade, for example, when upgrading Confluence 3.0 to 3.0.1.

If you have made Confluence site-wide layout customizations:

1. Choose the**cog icon** , then choose**General Configuration**
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 customizations available (preferably in a form which can be copied and pasted).
4. Click **Reset Default** next to the layout whose customizations need to be reapplied.
5. Click **Create Custom** next to the same layout and reapply your customizations (by copying and pasting them) into the appropriate locations within the new default layout.
6. Click the **Save** button.
7. Repeat this procedure from step 4 for each layout whose customizations need to be reapplied.

If you have made space-specific layout customizations:

1. Go to the space and choose**Space tools**>**Look and Feel**from the bottom of the sidebar
2. Choose**Layout**. The decorators are grouped under **Site**, **Content** and **Export** layouts.
3. Ensure you have all your customizations available (preferably in a form which can be copied and pasted).
4. Click **Reset Default** next to the layout whose customizations need to be reapplied.
5. Click **Create Custom** next to the same layout and reapply your customizations (by copying and pasting them) into the appropriate locations within the new default layout.
6. Click the **Save** button.
7. Repeat this procedure from step 5 for each layout whose customizations need to be reapplied.

Step 3. Test your Modifications Carefully

Changes may interact unpredictably with future versions of Confluence. When upgrading, you should always test your custom modifications thoroughly before deploying them on a live site. It's beyond the scope of Atlassian Support to test and deploy these changes.

Turning Off Caching

Velocity is configured to cache templates in memory. When you edit a page from within Confluence, it knows to reload that page from disk. If you are editing the pages on disk, you will either have to turn off Velocity's caching temporarily in `WEB-INF/classes/velocity.properties`, or restart the server to make your changes visible.

The `velocity.properties` file is available in the `confluence-x.x.x.jar` file, 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.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Working With Decorator Macros

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

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

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

Macro	Usage
`#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 colored 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 customized 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.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Custom Decorator Templates

About Decorators

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

What this means for Confluence is that you can customize the look and feel of parts of your Confluence site through editing decorators, for example:

- The "Main" decorator defines the generic header and footer
- The "Page" decorator defines how a page is displayed
- The "Printable" decorator defines the look and feel of the printable versions of pages.

You can view and edit these decorators from within Confluence. Changes to the decorators will affect all spaces in that Confluence installation.

The decorator that is used to draw Confluence's administrative pages cannot be edited from within Confluence. This means that if you make a mistake that renders the rest of the site unuseable, the administrative pages should still be available for you to fix the template.

Browsing the Default Decorators

At any time, you can browse the default decorators that come packaged with Confluence by following the "View Default" links on the "Site Layouts" page. The template browser also allows you to view the "#parsed" templates that are included within the template when it is compiled. While you can't edit these included templates, you will probably have to copy some or all of them into your custom template as you do your customization.

Editing Custom Decorators

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

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Customizing a Specific Page

If you'd like to change the appearance of a specific page, you can modify the corresponding Velocity template. Here's how to find out which one:

1. Access the page. Note the name of the action. For example, the "Contact Administrators" page is `<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.

Customizing the Login Page

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

Notes:

- Customizations to the Confluence login page will need to be reapplied when you upgrade Confluence. Consider this before making drastic changes to the layout, and be sure to keep a list of what you have changed for your upgrade process later.
- Please test your changes on a **test** Confluence site first.

Only administrators with access to the server where Confluence is running can modify the Confluence login page.

Related pages:

- Changing the Site Logo
- Velocity Template Overview
- Customizing Site and Space Layouts
- Changing the Look and Feel of Confluence
- Modify Confluence Interface Text

To change the login page:

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

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

Modify Confluence Interface Text

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

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

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

Replace "x.x.x" with your Confluence version, for example for 4.3.2, it will be named "confluence-4.3.2.
jar".
Within this File, the relevant file to edit is :\com\atlassian\confluence\core\ConfluenceActionSupport.
properties.
```

Refer to Editing jar files for reference.

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

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

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

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

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

Steps For Modification

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

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

4. Restart Confluence

Modify Keyboard Shortcuts

Confluence provides a set of keyboard shortcuts. You could customize the shortcuts by making modifications inside the `ConfluenceActionSupport.properties` file.

- To disable a particular shortcut, you can simply just comment out a respective line of code. One may like to disable the shortcut to one of the navigation links: *View, Edit, Attachments, Info* . For instance, to disable shortcut to *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.

Customizing Email Templates

> Customizing the Confluence email templates is **not supported**. If you do decide to edit the templates we strongly recommend you use a test instance of Confluence.
>
> Any customizations you make to the Confluence email notification templates will need to be reapplied after upgrading Confluence.

Email notification templates are contained within the `confluence-email-notifications`plugin, which is a system app (plugin) that is installed automatically when you install Confluence.

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

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

To change the email notification templates:

1. In the Confluence web application folder, find the file`/confluence/WEB-INF/atlassian-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 theguide on editing files within .jar archives for other methods).
5. Drop the new jar file into the`/confluence/WEB-INF/atlassian-bundled-plugins`folder (replacing the original file - you might want to make a copy of the original file for easy roll back) and then restart your instance.
6. Test your changes carefully before installing the updated plugin in production.

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

RELATED TOPICS

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

Changing the Default Behavior and Content in Confluence

Confluence comes with some handy default settings that determine what people see when they first enter the Confluence site, and the default content that is put into new spaces and other areas of Confluence.

Confluence administrators can change the settings to customize the behavior and the default content of their Confluence site:

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

Related pages:

- Changing the Look and Feel of Confluence

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 TemplatesandBlueprints.

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 Customizing Default Space Content).

Administrators can also disable templates and blueprints, to stop them appearing in the Create and Create Space dialogs anywhere in their Confluence site.

To **disable** a template or blueprint across the entire Confluence site:

- Go to > **General Configuration**>**Global Templates and Blueprints**.
- Select**Disable** next to the template, page blueprint or space blueprint you wish to disable.

Administrators can re-enable these templates and blueprints at any time.

Changing the Site Title

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

To change the title of your Confluence site:

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

Related pages:

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

Choosing a Default Language

Administrators can define a default language to be applied to all spaces in your Confluence site. Note that individual users can select a language preference for their session.

Related pages:

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

Setting the default language

To change the default language for the Confluence site:

1. Choose the**cog icon** , then choose**General Configuration**
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.

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

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

The following languages are still bundled, but we no longer translate new features for these languages.

- Eesti (Eesti | Estonia)
- slenska (sland | Iceland)
- Slovenina (Slovensk republika | Slovak Republic)
- Romn (Romnia | Romania)

Other settings that affect the language

Individual users can choose the language that Confluence will use to display screen text and messages. Note that the list of supported languages depends on the language packs installed on your Confluence site.

The language used for your session will depend on the settings below, in the following order of priority from highest to lowest:

- The language preference defined in your user profile. Note that you need to be logged in for this setting to take effect.
- The language that you choose by clicking an option at the bottom of the Confluence login screen. Confluence stores this value in a cookie. When the cookie expires, the setting will expire too.

- The language set in your browser. The browser sends a header with a prioritized list of languages. Confluence will use the first supported language in that list.Confluence administrators can disable this option by setting the `confluence.browser.language.enabled` system property to `false`.
- The default language for your site, as defined by your Confluence site administrator.

Showing User Interface Key Names for Translation

This feature is useful if you are troubleshooting translations of the Confluence user interface. After opening the Confluence dashboard, you can add the following action to the end of your Confluence URL:

```
?i18ntranslate=on
```

For example`http://myconfluencesite.com?i18ntranslate=on`

This will cause each element of the user interface to display its special **key name**. This makes it easier to find the context for each key within the user interface.

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

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

```
?i18ntranslate=off
```


Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Configuring the Administrator Contact Page

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

See the explanation of Confluence Groups for Administrators.

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

On this page:

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

Related pages:

- Configuring Captcha for Spam Prevention

Customizing the Administrator Contact Message

You can customize the message that is presented to the user on the '**Contact Site Administrators**' page.
To edit the administrator contact message:

1. Choose the**cog icon** , then choose**General Configuration**
2. 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 orConfluence 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

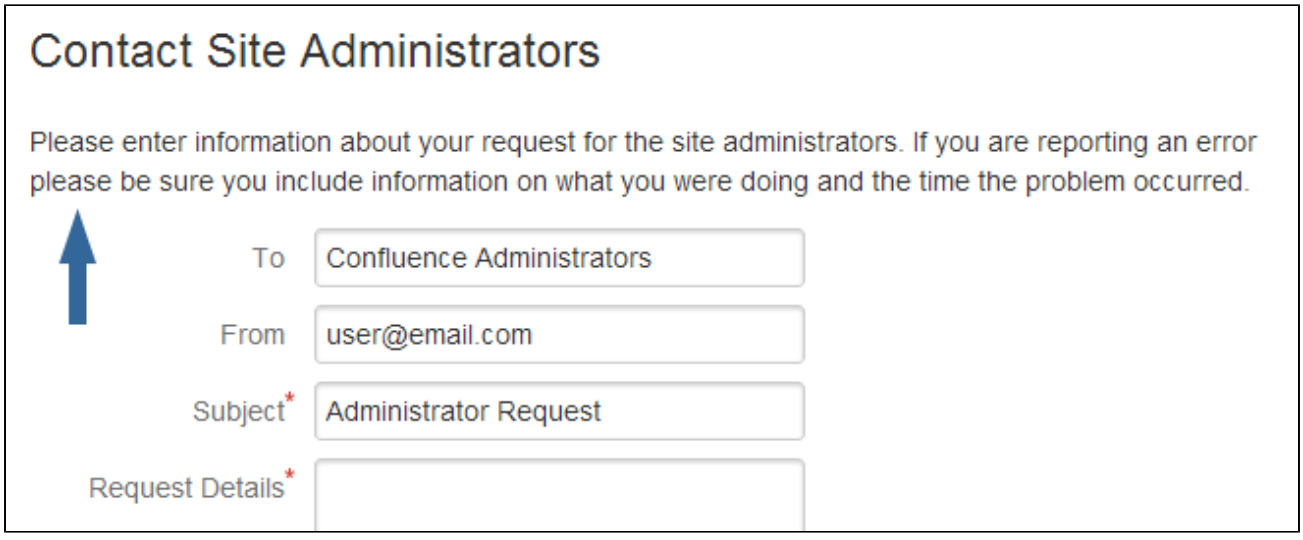

To restore the message to its default simply remove the custom message you entered when following the instructions above, so that the 'Custom Contact Administrators Message' field is empty.

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**
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.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Configuring the Site Home Page

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

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

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

Related pages:

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

To use a page as your site home page:

1. Go to > **General Configuration**>**Further Configuration**.
2. 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.

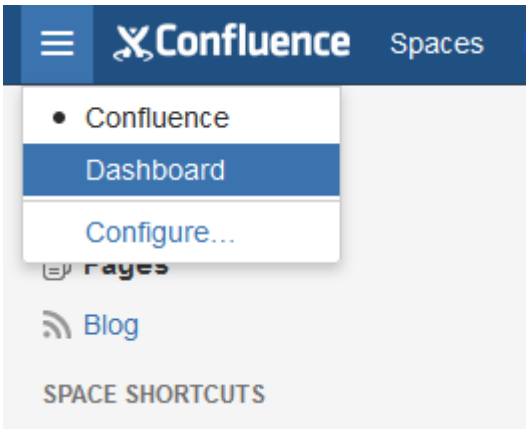

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Customizing Default Space Content

Confluence Administrators can edit the template that is used to create the home page for new sites. This default content appears on the home page when a new space is created. There is a different template for site spaces, personal spaces and space blueprints.

The default content in the template only appears for new spaces (those that are created after you have defined the content). Changes to the template do not affect existing home pages.

On this page:

- Edit the default home page for a blank space
- Reset the original default content

Related pages:

- Spaces
- Page Templates

Edit the default home page for a blank space

To edit the default (blank) space content template:

1. Choose the**cog icon** , then choose**General Configuration**
2. Choose **Global Templates and Blueprints** in the left-hand panel.
3. Choose **Edit** next to 'Default Space Content' or 'Default Personal Space Content' depending on whether you want to customize the content for new site space or personal space home pages.
4. Enter the content that you want to appear on the home page for new blank spaces. you can add variables, macros and other content in the saw way as edited a page template.
5. Choose **Save**.

The following variables are available to be added to the default space content templates.

- **$spaceKey** - inserts the space key into the site space homepage
- **$spaceName** - inserts the space name into the site space homepage
- **$userFullName** - inserts the user (owner of the personal space) into the personal space homepage
- **$userEmail** - inserts the email address of the user (owner of the personal space) into the personal space homepage.

Default space templates differ from ordinary page templates in that they do not present the user with a form to complete, so variables should be limited to those listed in the **Variables** menu.

Some macros, such as the Table of Contents macro, may not display correctly when you preview the template as they are designed to work on a page. The macros will display correctly on the home page when you create a new space. For more information on editing a template, including adding macros see - Adding Content to a Template.

Reset the original default content

To reset the original default content:

1. Choose the**cog icon** , then choose**General Configuration**
2. Choose **Global Templates and Blueprints** in the left-hand panel.
3. Choose **Reset to default**next to the template you wish to reset.

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

Editing the Site Welcome Message

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

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

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

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

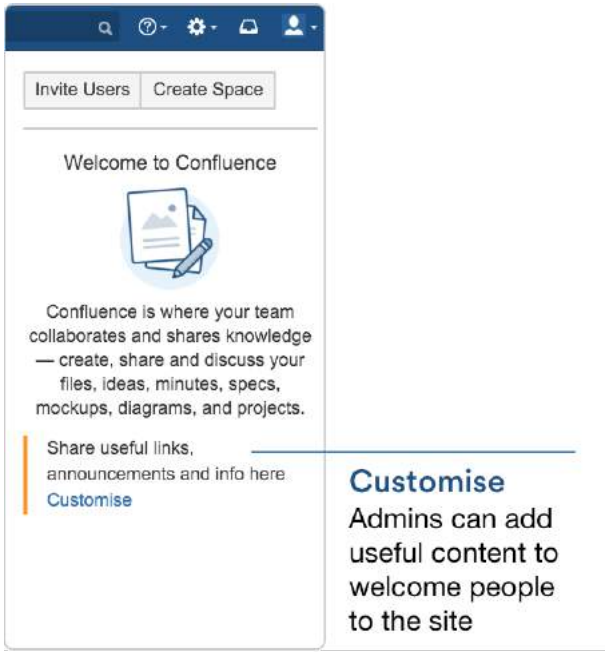

Hints for using the template editor

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 Pagemacro to include content from elsewhere in your site, rather than adding content directly to the template.

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

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

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

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Integrating Confluence with Other Applications

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

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

Related Topics

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

Linking to Another Application

Application links is a bundled plugin that allows you to link Confluence to other Atlassian products or external applications. Thanks to this, they can exchange information or give access to certain resources or functionalities. In the case of Confluence, the most common integration is the one with Jira it lets you easily display information about Jira issues on Confluence pages, link pages to issues in Jira, or use other features created specifically for app links.

You can also link Confluence to external applications using either OAuth 1.0 or OAuth 2.0. These integrations are typically used for internal integrations and require that your application is compatible with application links.

View application links

To view application links:

1. Go to > **General Configuration**> **Applications**.
2. Youll see the following page:

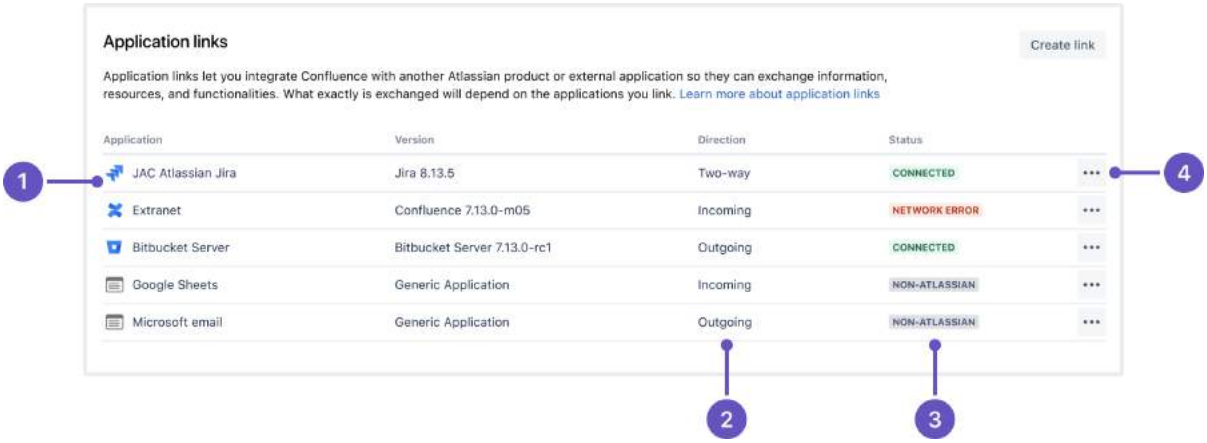

1. **Application** -Name of the linked application and its version. For external applications, it always shows Generic application.
2. **Direction**- Communication direction, either Incoming, Outgoing, or Two-way. For Atlassian products, you should configure two-way communication, but some external applications wont need it.
3. **Status** -Connection status. For external applications, it always shows Non-Atlassian.
4. **Actions** -Actions you can do on your links, such as edit or delete. For OAuth 2.0 connections, you can additionally view your OAuth credentials.

Link to Atlassian products or external applications using OAuth 1.0

When you link to other Atlassian products, the communication is using OAuth 1.0. You can also use this option to link to external applications, although we recommend that you update your integrations to use OAuth 2.0.

To link to other Atlassian products or external applications using OAuth 1.0:

1. In application links, select **Create link**.
2. Select **Atlassian product** as the link type.
3. Enter the **Application URL** of your Atlassian product or external application.
4. Follow the steps in the wizard. Youll be redirected between Confluence and the product youre linking to to authorize the two-way connection.

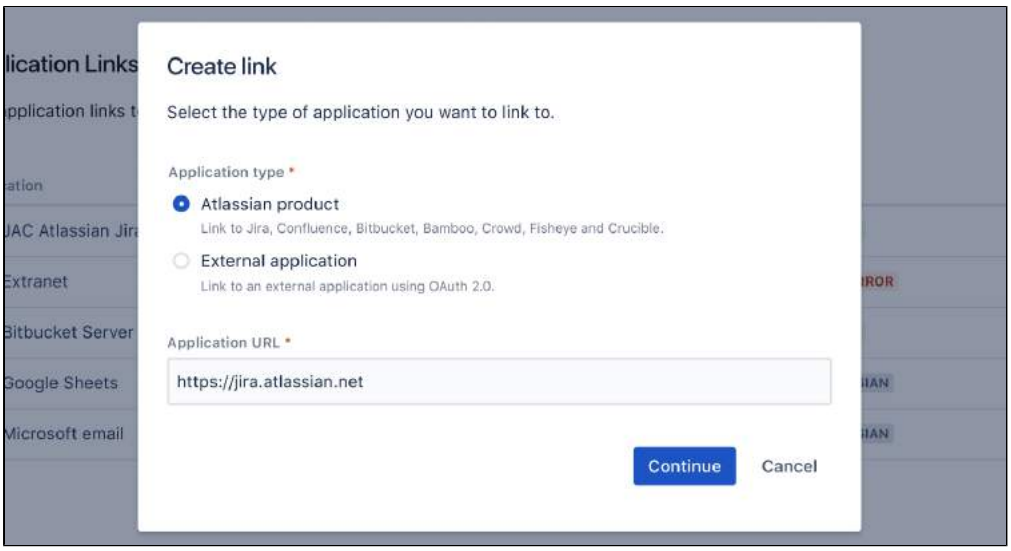

Link to external applications using OAuth 2.0

You can link Confluence to external applications using OAuth 2.0 in both directions, either making Confluence act as a client (outgoing link) or provider (incoming link).

Configure Confluence as an OAuth 2.0 client (outgoing link)

In this scenario, Confluence acts as an OAuth client, requesting data from the external application.

For more information, see Configure an outgoing link.

Configure Confluence as an OAuth 2.0 provider (incoming link)

In this scenario, Confluence acts as an OAuth provider, allowing the external application to access its data.

For more information, see Configure an incoming link.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Configure an outgoing link

When you configure an outgoing link to an external application, Confluence requests data from this application, which means that it acts as the OAuth client. To learn more about the type of links and additional details, see Linking to Another Application.

Before you begin

You need to ensure the following:

- Your server needs to run over HTTPS If it doesnt, you will not be able to configure OAuth 2.0.
- Your base URL needs to be configured correctly. This is important as the redirect URL youll need to provide is based on the Confluences base URL.

Create an outgoing link using application links

To create an outgoing link:

1. Go to > **General Configuration**>**Application links**.
2. Select **Create link**.
3. Select **External application**, and then choose **Outgoing** as the direction.
4. Fill in the details as described in the sections below.

Configure your outgoing link

Follow these steps to configure your link.

1. Choose a service provider

Choose one of the following providers that you want to configure. For Google or Microsoft, some of the fields will be pre-filled.

- Google
- Microsoft
- Custom (for internal tools or other providers)

2. Copy the Redirect URL and register it in your external application

Copy the Redirect URL and register it in your external application to obtain the client ID and client secret required to complete the configuration.

If youre using Google or Microsoft as service providers, youll be able to copy the Redirect URL right away. For custom providers, you need to first provide the Authorization endpoint and Token endpoint. For more info on registering the URL with Google or Microsoft, see:

- OAuth 2.0 in Google
- OAuth 2.0 in Microsoft

Different providers might have different requirements related to the redirect URL. For example, Google does not allow it to be a private IP address. Make sure you provide an external URL (for example of a load balancer for Data Center).

3. Provide remaining application details

Provide the remaining details. Here you can find descriptions for all the fields:

Name	Description
Client ID	The client ID generated by the external application after registering Confluences Redirect URL. This is the public identifier of the application.

Client secret	The client secret generated by the external application after registering Confluences Redirect URL. This is the shared secret between Confluence and the application, which ensures the authorization is secure.
Scopes	The required OAuth 2.0 scopes (permissions) that control what Confluence can do in the external application.
Authorization endpoint	The HTTPS URL where authorization to use OAuth 2.0 is started.
Token endpoint	The HTTPS URL where refresh token requests are sent. As OAuth 2.0 tokens have an expiry, Confluence will periodically update the token.
Redirect URL	The Redirect URL that must be registered in the external application to obtain its client ID and client secret. This redirects the authentication flow back to Confluence.

4. Save your outgoing link

After you save the link, it will appear on the list together with other application links.

Troubleshooting

I fail to get an OAuth 2.0 refresh token

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Configure an incoming link

When you configure an incoming link with an external application, you allow this application to access Confluence data, which means that Confluence acts as the OAuth provider. To learn more about the type of links and additional details, see Linking to Another Application.

Before you begin

- If youre creating an OAuth 2.0 integration and want to use Confluence as the provider, you can find the details of our OAuth 2.0 implementation in Confluence OAuth 2.0 provider API.
- You can configure additional details using system properties.

Create an incoming link using application links

To create an incoming link:

1. Go to > **General Configuration**>**Application links**.
2. Select **Create link**.
3. Select **External application**, and then choose Incoming as the direction.
4. Fill in the details as described in the sections below.

Provide application details

In this type of link, you only need to provide the Redirect URL (also known as Callback URL) from your external application. After authorizing the application, the user will be redirected to this URL with the authorization code.

Provide application permissions

Select permissions the application can have on your instance. You can choose the following permission scopes:

- Read
- Write
- Admin
- System admin

Note that even if you grant higher permissions, the application wont be able to do more than the user authorizing it. For more info on what each of these scopes do, see OAuth 2.0 scopes for incoming links.

Copy OAuth credentials to the application

After providing the Redirect URL and selecting the scopes, Confluence will generate the OAuth credentials that include these details. You need to copy the credentials to your external application to complete the link.

At this point, the application link has already been created in Confluence. You can view its details in Application links, including the OAuth credentials in case you needed to access them later.

View OAuth credentials for an existing link

If you lose your OAuth credentials, you can access them any time in the details of the application link you created.

To view OAuth credentials:

1. Go to > **General Configuration**>**Application links**.
2. Find the application link youre interested in, and select **More actions** > **View credentials**.

OAuth 2.0 scopes for incoming links

When creating incoming links from external application, you need to select scopes, which are permissions the application can have on your instance.

What the application can do with scopes

As an admin, you can select which scopes the application can request from the authorizing user, but the actual permissions will always be capped at what this user can do. For example, even if you select the ADMIN permissions, the application wont be able to use them if the authorizing user only has WRITE permissions.

Scopes

Here are the scopes you can select when configuring the link. The same scopes will be displayed to users when they authorize the integration. They can later be accessed in their user profile in Authorized applications, where they can also revoke the granted access.

Scope	Description
READ	**View content** View content your account can view, including spaces, pages, blog posts, custom content, attachments, comments, and templates. Also view your user profile.
WRITE	**Create, update, and delete content** Create, update, and delete content your account can change, including spaces, pages, blog posts, custom content, attachments, comments, and templates. Also change your user profile.
ADMIN	**Administer Confluence** Perform most administrative functions on the entire Confluence instance, excluding functions such as backups, imports, and infrastructure settings which are limited to system administrators.
SYSTEM_ADMIN	**Administer Confluence system** Perform all administrative functions on the entire Confluence instance, including functions such as backups, imports, and infrastructure settings.

Configuring Workbox Notifications

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

Possible configurations:

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

On this page:

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

Notes:

- **Workbox includes notifications and tasks:** When you enable in-app notifications, personal tasks are also enabled in the workbox. When you disable in-app notifications, the workbox no longer appears and personal tasks are therefore not available on this server.

Which notifications are included?

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

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

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

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

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

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

Configuring the polling intervals

The polling intervals are used by the Confluence server that displays in-app notifications and tasks in its workbox.

Option	Description

Active polling interval	This is the number of seconds that Confluence will wait before checking (polling) for new notifications relevant to the page that the user is currently viewing. This setting applies to the page open in the browser tab that currently has focus. It does not matter whether the user has the workbox open or not.
Inactive polling interval	This is the number of seconds that Confluence will wait before checking(polling) for new notifications relevant to all pages that are not currently in focus. These pages may be on the Confluence server that displays the workbox, or on other Confluence or Jira servers that send their notifications to this server. This setting defines an upper limit. For inactive pages, Confluence starts with a polling interval equal to the active polling interval, then gradually increases the interval between polls until it reaches the limit defined here.

Including notifications from Jira

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

To include notifications from a Jira application:

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

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

Notes:

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

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

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

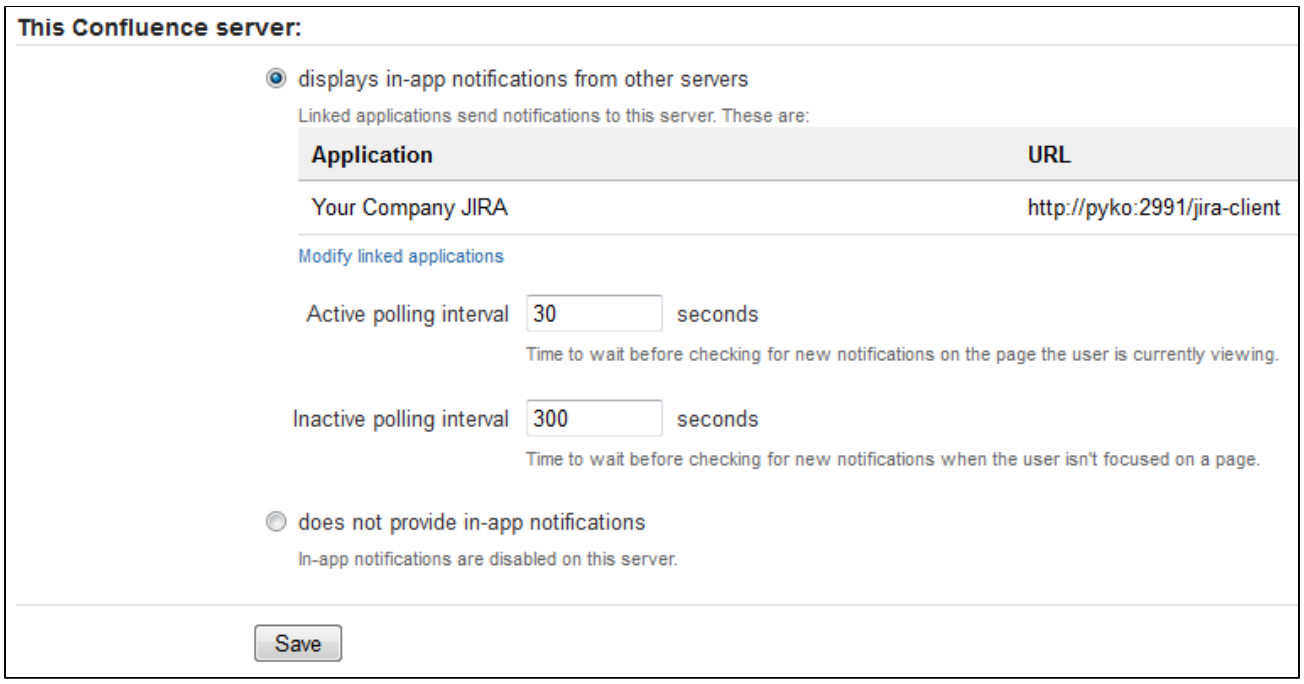

Stopping Jira applications from sending notifications to Confluence

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**
   - Choose**Application Links**in the left-hand panel.
   - Set up the link as described inLinking to Another Application.
2. Configure the notification settings in*ConfluenceChatty*:
   - Choose**In-app Notifications**in the left-hand panel of the Confluence administration console.
   - Choose**displays in-app notifications from other servers**.
3. Configure the notification settings in*ConfluenceQuiet*:
   - Choose**In-app Notifications**in the left-hand panel of the Confluence administration console.
   - Choose**sends in-app notifications to another server**.
   - Select the Confluence server that will display the workbox in our example, this is*ConfluenceChatty*. (The entry for*ConfluenceChatty*will appear here only if you have already configured*ConfluenceChatty*to display in-app notifications.)

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

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

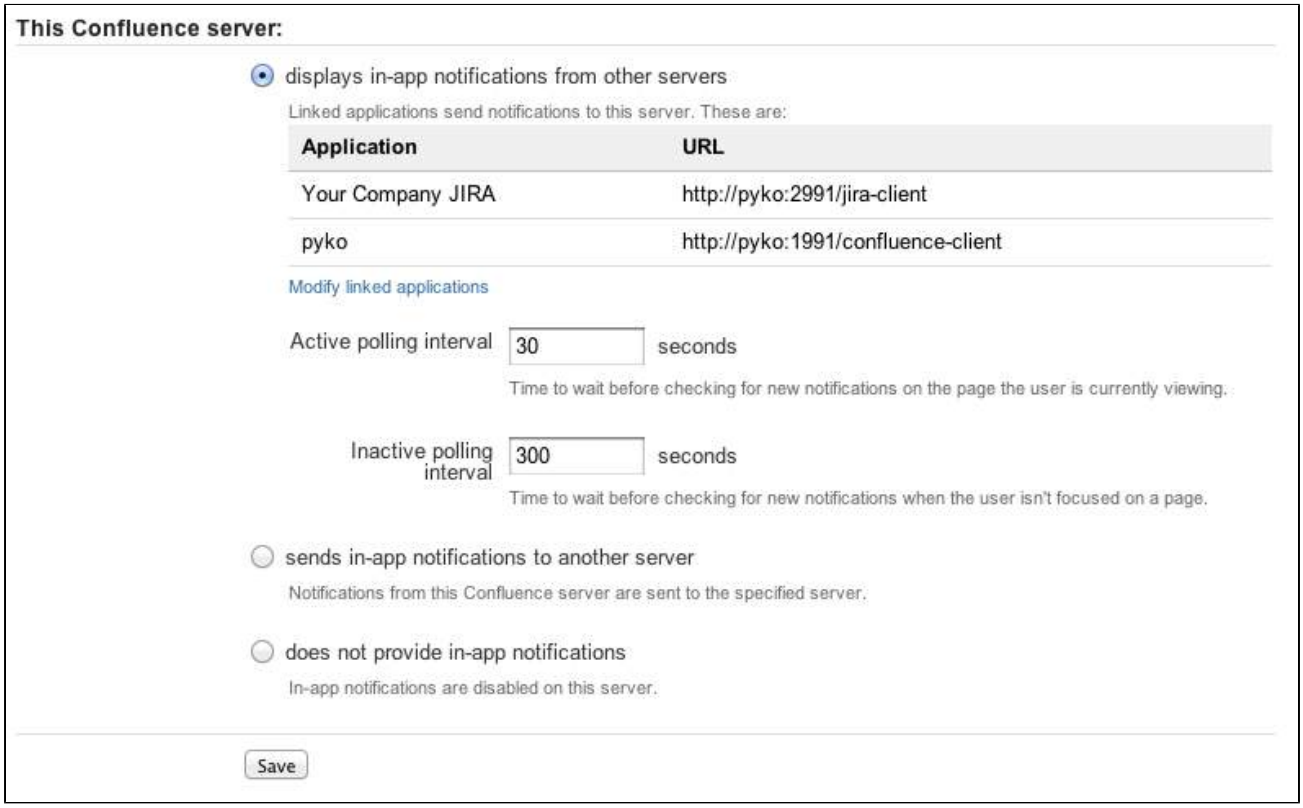

Sending Confluence notifications to another Confluence server

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*Confluence Quiet*above.

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

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

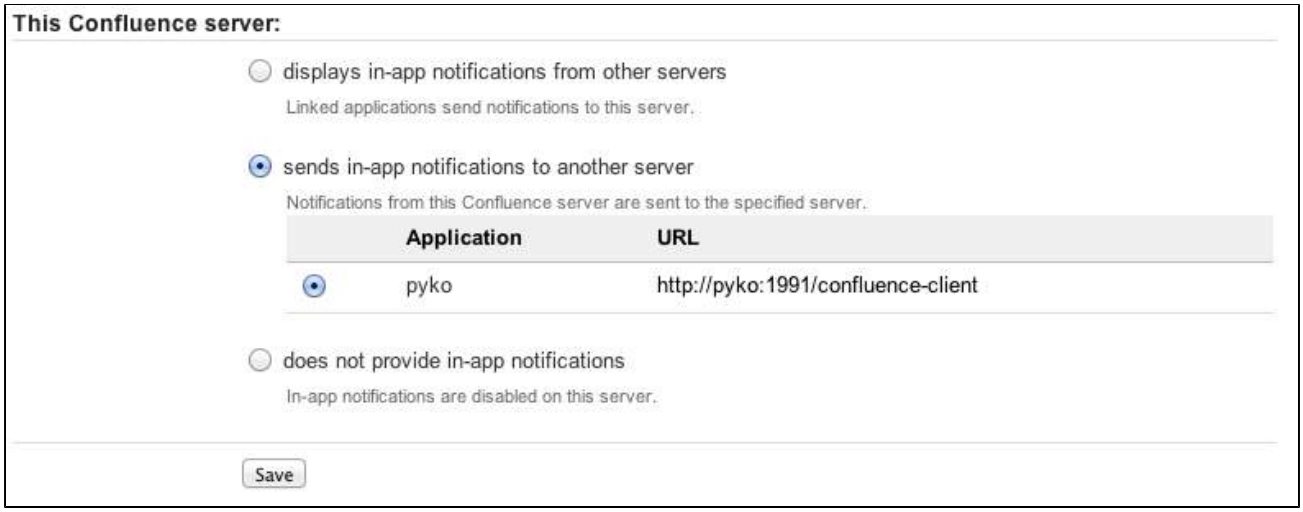

Disabling workbox and in-app notifications in Confluence

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.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Integrating Jira and Confluence

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

Learn more about what you can do with Jira and Confluence

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

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. SeeCan Multiple Atlassian Products Be Deployed in a Single Tomcat Container?for more information.

Use Jira and Confluence together

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

Delegate user management to Jira

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

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

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

Connect Jira and Confluence with an application link

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

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

Having trouble integrating your Atlassian products with application links?

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

Registering External Gadgets

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

There's two ways to register external gadgets:

- **Subscribe to all of the external application's gadgets:**You can add all the gadgets from your Jiraapplication, Bamboo, 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 Allowlist
- The big list of Atlassian gadgets
- Linking to Another Application

Setting up a trust relationship with the other application

In addition to registering the external gadgets, we recommend that you set up an OAuth or Trusted Application relationship between the application that serves the gadget (the service provider) and Confluence (the consumer). The trust relationship is required for gadgets that access restricted data from the external web application.

See how to configure OAuth or Trusted Applications Authentication, using Application Links.

If the external web application provides anonymous access to all the data you need in the gadgets, then you do not need a trust relationship.

For example, if your gadgets will retrieve data from Jira and your Jira server includes projects and issues that are restricted to logged-in users, then you will need a trust relationship between Confluence and Jira. If you do not set up the trust relationship, then the gadgets will show only the information that Jira makes visible to anonymous users.

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

Subscribing to all of the application's gadgets

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

To subscribe to another site's gadgets:

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

Screenshot: Subscribing to a gadget feed

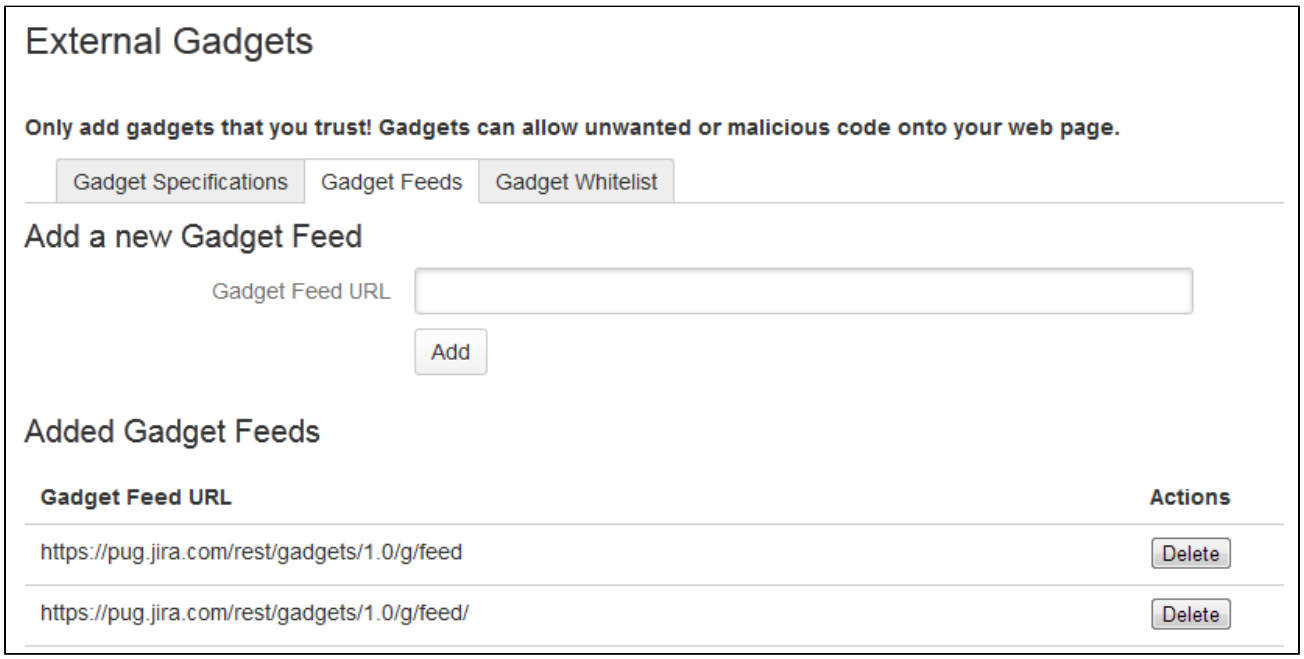

Registering individual gadgets

If you cannot subscribe to an application's gadgets, you will need to register the gadgets one by one. This is necessary for applications and websites that do not support gadget subscription, and for applications where you cannot establish a trusted relationship via Application Links.

First you will need to get the gadget URL and copy it to your clipboard.

Getting a gadget's URL from an Atlassian application

If your application is another Atlassian application:

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

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

If the gadget is supplied by a plugin, the URL will have this format:
`http://my-app.my-server.com:port/rest/gadgets/1.0/g/my-plugin.key:my-gadget/my-path/my-gadget.xml`
For example:
`http://mycompany.com/jira/rest/gadgets/1.0/g/com.atlassian.streams.streams-jira-plugin:activitystream-gadget/gadgets/activitystream-gadget.xml`

To find a gadget's URL in JIRA:

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

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

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

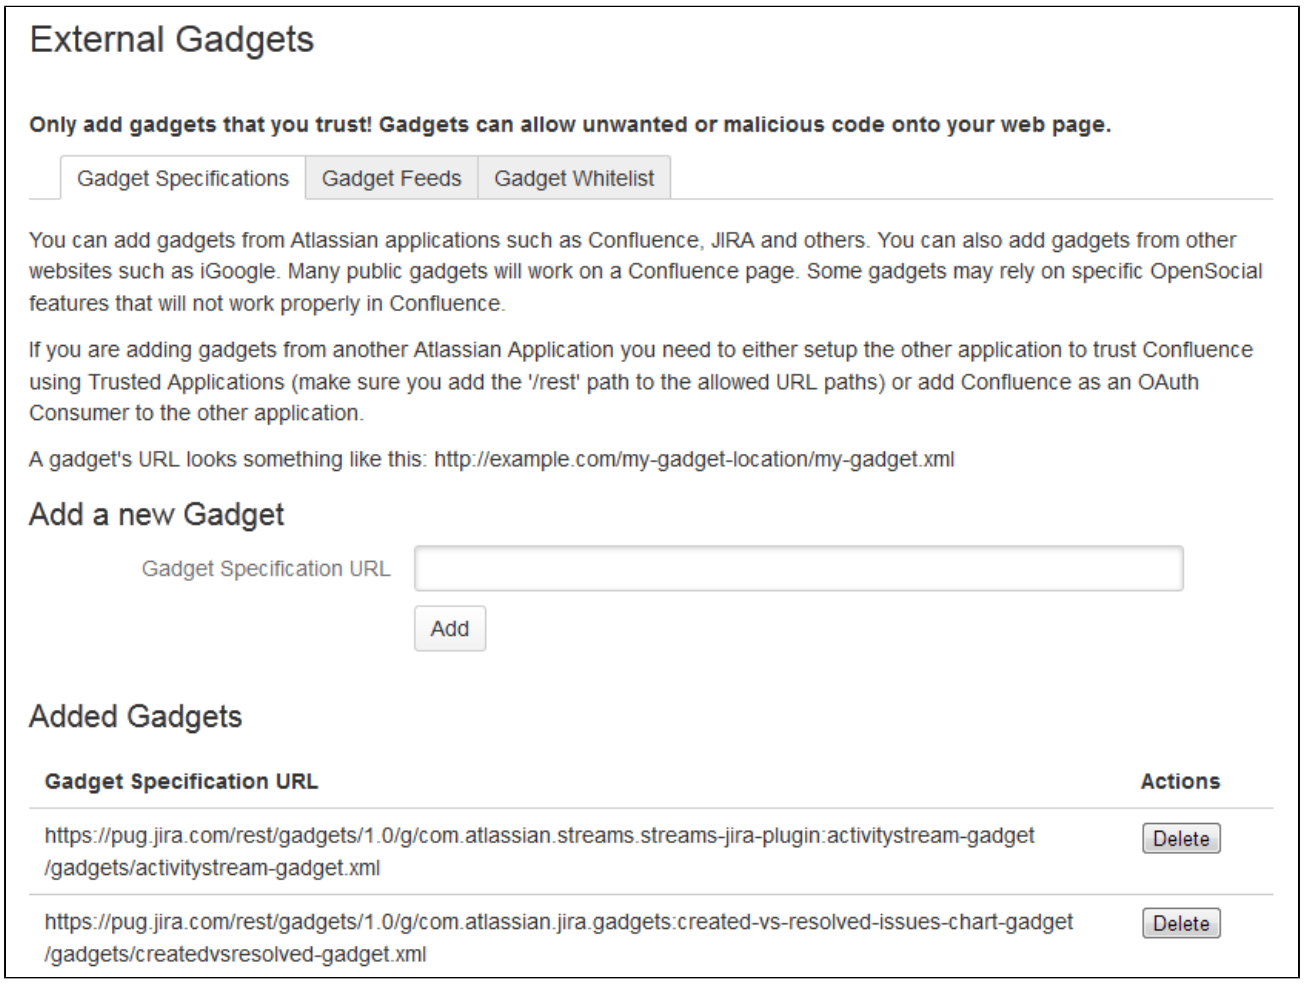

Removing access to external gadgets

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

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.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Configuring the Office Connector

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

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

On this page:

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

Related pages:

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

Enabling and disabling the Office Connector

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

To enable or disable the Office Connector modules:

1. Go to > **Manage apps**
2. Choose**System**from the filter drop down and then search for**Office Connector**
3. Expand the Office Connector listing. From here you can:
    - Choose**Configure**to specify preferences for the Office Connector (this opens the configuration screen described below)
    - Click**Disable**to disable all modules of the app
    - Expand the**modules**list to enable or disable selected Office Connector modules.

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

Configuring the Office Connector options

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

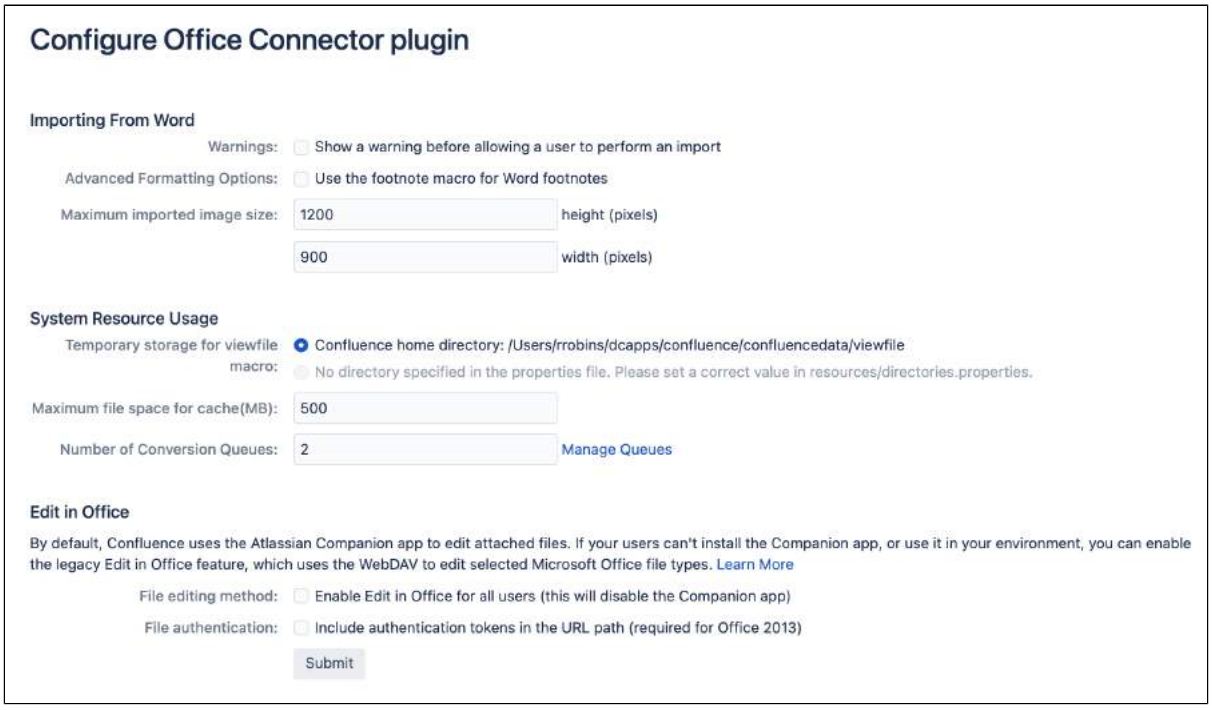

To set the configuration options for the Office Connector:

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

Option	Default Value	Description
Warnings: Show a warning before allowing a user to perform an import	Disabled	If this option is enabled, the user will receive a warning when importing a Word document. The warning will tell the user when they are about to overwrite existing content.
Advanced Formatting Options: Use the footnote macro for Word footnotes	Disabled	⚠ **Note:** This feature requires a third party app. If this option is enabled, a Confluence page created from an imported Word document will use the {footnote} macro from Adaptavist to render any footnotes contained in the document. Note that you will need to install the Content Formatting for Confluence app from Adaptavist to get this macro.
Maximum imported image size	1200x900	Users will be prevented from importing a Word document if it contains images that exceed the maximum imported image size. The size is based on the size the image is displayed in Word, not the resolution of the image. Defaults to 1200 wide x 900 high.
Temporary storage for viewfile macro	The Confluence Home directory.	The {viewfile} macro will cache data temporarily. This option allows you to set the location of the cache. Available settings are: **Confluence home directory** The temporary file will be stored in your Confluence Home directory. **A directory specified in the `directories.properties` file** You can specify a location by editing the Office Connector's `directories.properties` file: 1. Locate the `OfficeConnector-x.xx.jar` file (where x.xx is the version number) in your Confluence Home directory and copy it to a temporary location 2. Unzip the JAR file and find the `resources/directories.properties` file. The content of the file looks like this: `#Complete the following line to set a custom cache directory. #If resetting to blank, don't delete anything before or including the '=' com.benryan.confluence. word.edit.cacheDir=` 3. Edit the last line, adding the path to your required temporary location directly after the '=' character. For example: On Windows:com.benryan.confluence.word.edit. cacheDir=c: \my\path\ On Linux:com.benryan.confluence.word.edit. cacheDir=/home /myusername/my/path 4. Save the file, recreate the JAR and put it back in your Confluence Home directory, overwriting the original JAR.
Maximum file space for cache (MB)	500	This is the maximum size of the cache used by the {viewfile} macro. (See above.)

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Number of Conversion Queues	2	This is the maximum number of threads used to convert PowerPoint, Excel files or PDF slide shows. You can use this setting to manage Confluence performance, by limiting the number of threads so that the Office Connector does not consume too many resources. Click **Manage Queues** to view attachments that are still pending conversion.
File editing method:Enable Edit in Office for all users (this will disable the Companion app)	Disabled	This allows administrators to disable the Companion app method for editing files, and instead use the Office Connector to edit compatible files.
File authentication:Include authentication tokens in the URL path (required for Office 2013)	Disabled	If this option is enabled, the Office Connector will use authentication tokens in the URL. This needs to be enabled to edit Office 2013 documents.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Managing Webhooks

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

On this page:

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

A webhook consists of:

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

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

Securing the webhook

Confluence uses webhook secrets to authenticate the payload.Combined with HTTPS, it helps ensure the message transmitted is the one that Confluence intended to send, and that the contents were not tampered with.

When you define a secret for a webhook, each request is signed via a Hash-based Message Authentication Code (HMAC). The default for this algorithm is HMACSha256. The header X-Hub-Signature is defined and contains the HMAC.

To authenticate the validity of the message payload, the receiver can perform the HMAC algorithm on the received body with the secret as the key to the HMAC algorithm.If the results don't match, it may indicate there was a problem with transmission that has caused the message payload to change.

Create a new webhook

You need Confluence Administrator or System Administrator global permissions to create a webhook.

To create a new webhook:

1. Go to > **General Configuration** >**Webhooks**.
2. Enter a **title** for your webhook.
3. Enter the **URL** of the application or server.
4. Enter a **secret**. This is a string of up to 255 characters that you define.
5. Select **Test connection** to check you can reach the application or server.
6. Choose the events that should trigger the webhook.
7. Select **Active** to make your webhook available immediately.
8. Select **Create**.

Screenshot: Creating a webhook to notify a chat application when a space is created or removed.

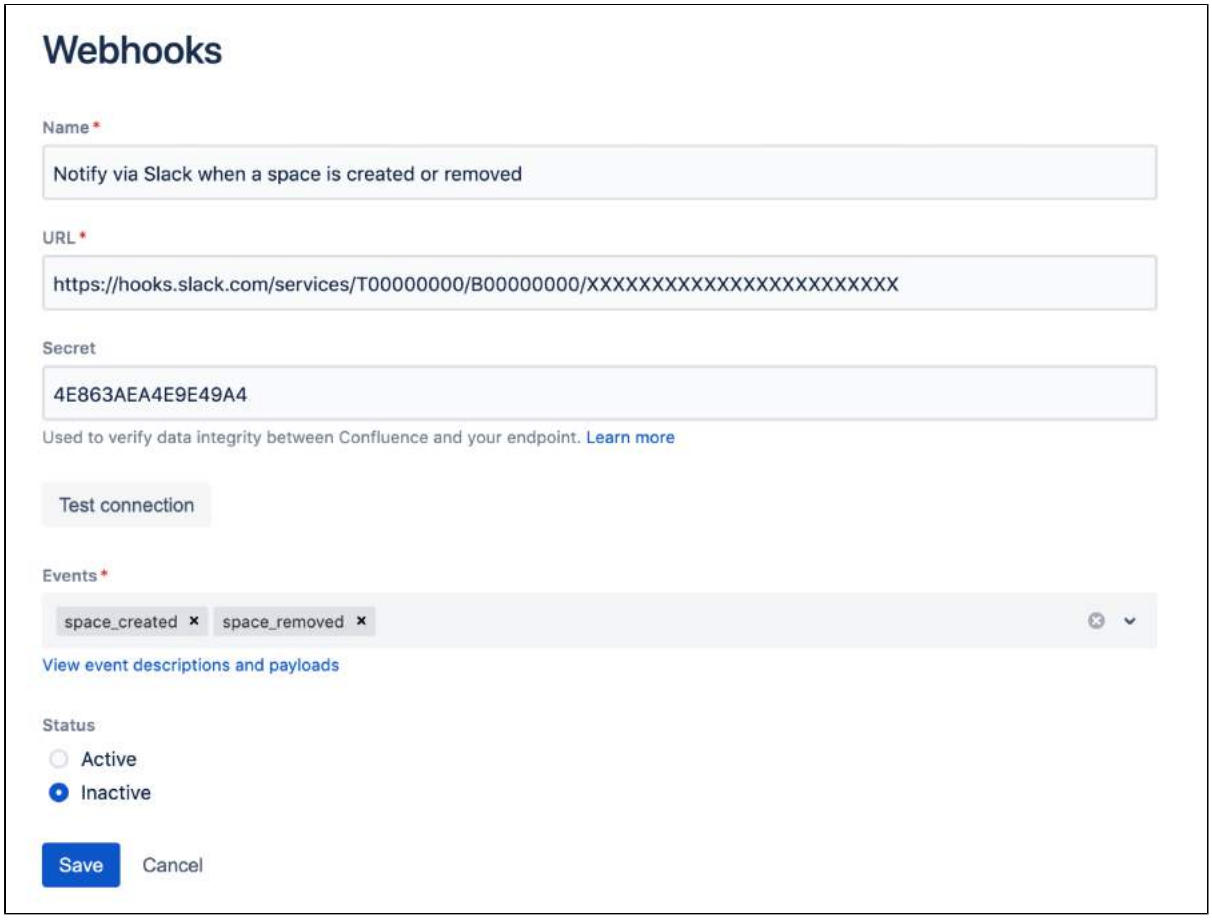

You can also create a webhook using the API. SeeWebhooks in the Confluence developer documentation.

Triggering webhooks

You can configure your webhook to be triggered by the following events.

Event	Triggered when...
attachment_created	a file is attached to a page or blog post
attachment_removed	a file is deleted (sent to the trash) from the attachments page (not triggered when a version is deleted from the file history)
attachment_restored	a file is restored from the trash
attachment_trashed	a file is purged from the trash
attachment_updated	a new file version of is uploaded directly or by editing the file
blog_created	a blog post is published
blog_removed	a blog post is deleted (sent to the trash)
blog_restored	a blog post is restored from the trash
blog_trashed	a blog post is purged from the trash
blog_updated	a blog post is edited
blueprint_page_created	a page is created from a blueprint (such as meeting notes, decision, or how-to)

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

comment_created	a page comment, inline comment or file comment is made
comment_removed	a page comment, inline comment, or file comment is deleted
comment_updated	a page comment, inline comment, or file comment is edited
content_created	a page, blog post, attachment (file), comment (page, inline, or file), or other file (such as a space logo) is created or uploaded.
content_restored	a page, blog post, or attachment (file) is restored from the trash
content_trashed	a page, blog post, or attachment (file) is purged from the trash
content_updated	a page, blog post, attachment (file), or comment (page, inline, and file) is edited.
content_permissions_updated	a view or edit restriction is applied or removed from a page or blog post
group_created	a new group is created
group_removed	a group is deleted
label_added	an existing label is applied to a page, blog post, or space
label_created	a label is added for the first time (did not already exist)
label_deleted	a label is removed from the last page, blog post, or space, and so ceases to exist
label_removed	a label is removed from a page, blog post, or space
page_children_reordered	the default ordering of pages is changed to alphabetical in the Space Tools > Reorder pages tab (is not triggered when you drag a page, or move a page, to change the page order)
page_created	a page is published for the first time, including pages created from a template or blueprint
page_moved	a page is moved to a different position in the page tree, to a different parent page, or to another space
page_removed	a page is deleted (sent to the trash)
page_restored	a page is restored from the trash
page_trashed	a page is purged from the trash
page_updated	a page is edited (triggered at the point the unpublished changes are published)
space_created	a new space is created
space_logo_updated	a new file is uploaded for use as the logo of a space
space_permissions_updated	space permissions are changed in the Space Tools > Permissions tab (is not triggered when you edit space permissions using Inspect Permissions)
space_removed	a space is deleted
space_updated	the space details (title, description, status) is updated in the Space Tools > Overview tab
theme_enabled	a specific theme or default theme is applied to to a space or the whole site

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

user_created	a new user account is created
user_deactivated	a user account is disabled
user_followed	someone follows a user
user_reactivated	a disabled user account is enabled
user_removed	a user account is deleted

Event payloads

Here's an example of the event payload for the `page_trashed` event. This is the raw data that's sent, in JSON format, to your endpoint.

```
{
  "timestamp":1596182511300,
  "event":"page_trashed",
  "userKey":"ff80818154ec9913015501e194f601d8",
  "page":{
     "id":309264476
  }
}
```

You'll note that the content is comprised mostly of IDs. This is to ensure that identifiable information is not stored by third party services, or leaked to users who do not have permission to see it.

Once received, you can use the REST API to interpret these IDs. SeeConfluence Server Rest API.

Circuit breaking

To help protect your Confluence site, any webhooks that fail consistently, are skipped for a period of time. By default, if a webhook fails five times, it is considered unhealthy and is skipped, initially for 10 seconds. If it continues to fail, it will be gradually shipped for longer periods, up to 10 hours.

A webhook may also be skipped if there are too many webhooks in flight. If there are 500 webhooks being invoked, further requests will be skipped until the number in flight drops below 500.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Managing your Confluence License

Your license entitles you to run Confluence and be eligible for support and upgrades for a specified period. It also defines the number of users who are entitled to use Confluence.

To quickly check the status of your license you can go to > **General Configuration** > **Troubleshooting and support tools**.

You'll need need Confluence Administrator or System Administrator permissions to view and edit your license.

On this page:

- Viewing your license details
- Updating your license
- Understanding the user count for your license
- Exceeding your licensed user count
- Reducing your user count
- Downgrading your license
- Finding your Support Entitlement Number (SEN)
- What happens when your maintenance or subscription expires

Related pages:

- Upgrading Beyond Current Licensed Period
- Confluence installation and upgrade guide
- Confluence administrator's guide

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).
- The number of users on your license that count towards your total licensed user count.
- Your license expiry date, for support and upgrade eligibility.
- Your server ID which is generated when you install Confluence for the first time and remains the same for the life of the installation (including after upgrades or changes to your license).
- Your support entitlement number (SEN).

Updating your license

If you change your license (for example to a license with more users), or migrate from Confluence Cloudand you will need to update your license.

To update your Confluence license:

1. Go to > **General Configuration**>**License Details**
2. Enter your new license in the**License key**field.
3. Choose**Save**.

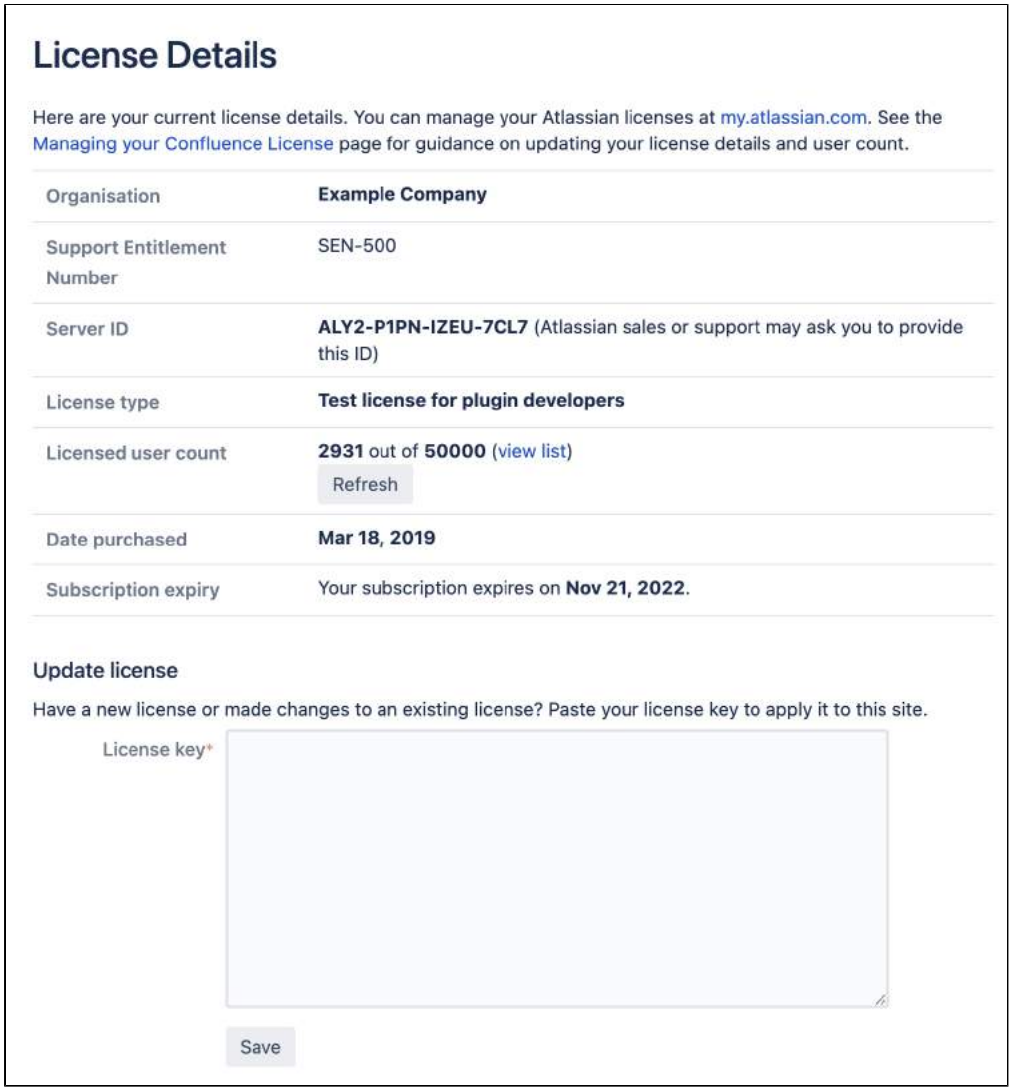

⚠️If you run Data Center in a cluster:

- the license will automatically propagate to all online nodes. However, any node that is offline won't be updated, and you may need to apply the license on this node when you bring it back online. See CONFSERVER-56375 LONG TERM BACKLOG
- in AWS, the license is not automatically written to the `confluence.cfg.xml` file in the shared home directory, which means new nodes aren't provisioned with the new license. See ~~CONFSERVER-60045~~ CLOSED

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 that count towards your licensed user count.

- It includes only users who have the 'can use' global permissionsfor the Confluence site.
- It does not include anonymous users, who may access your Confluence site if you have allowed anonymous access.
- It does not include deactivated users.

To view a list of your licensed users:

1. Go to > **General Configuration**>**License Details**

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

2. Select **view list**(1).
3. On the **Users** page, you will see a filtered list of Licensed users only (2), including details like Last login (3) to help you manage your license usage.

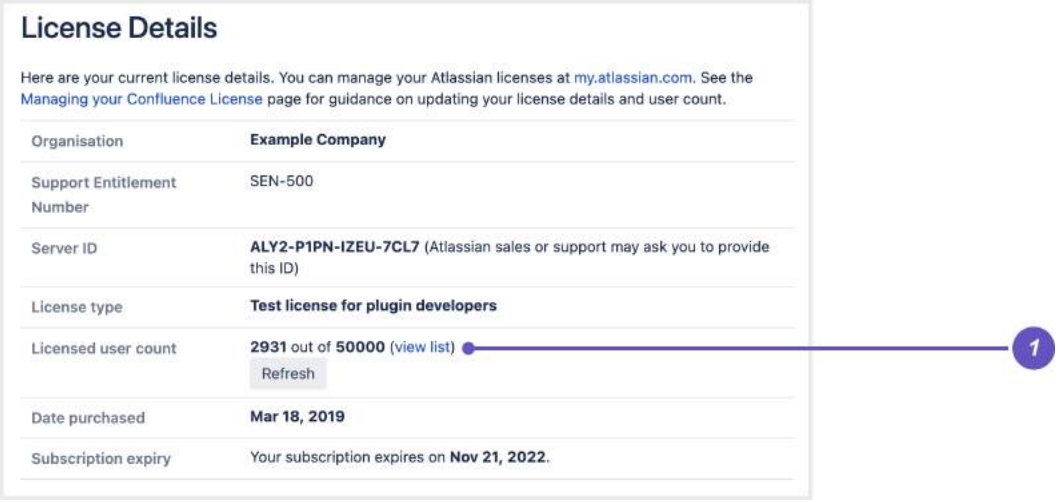

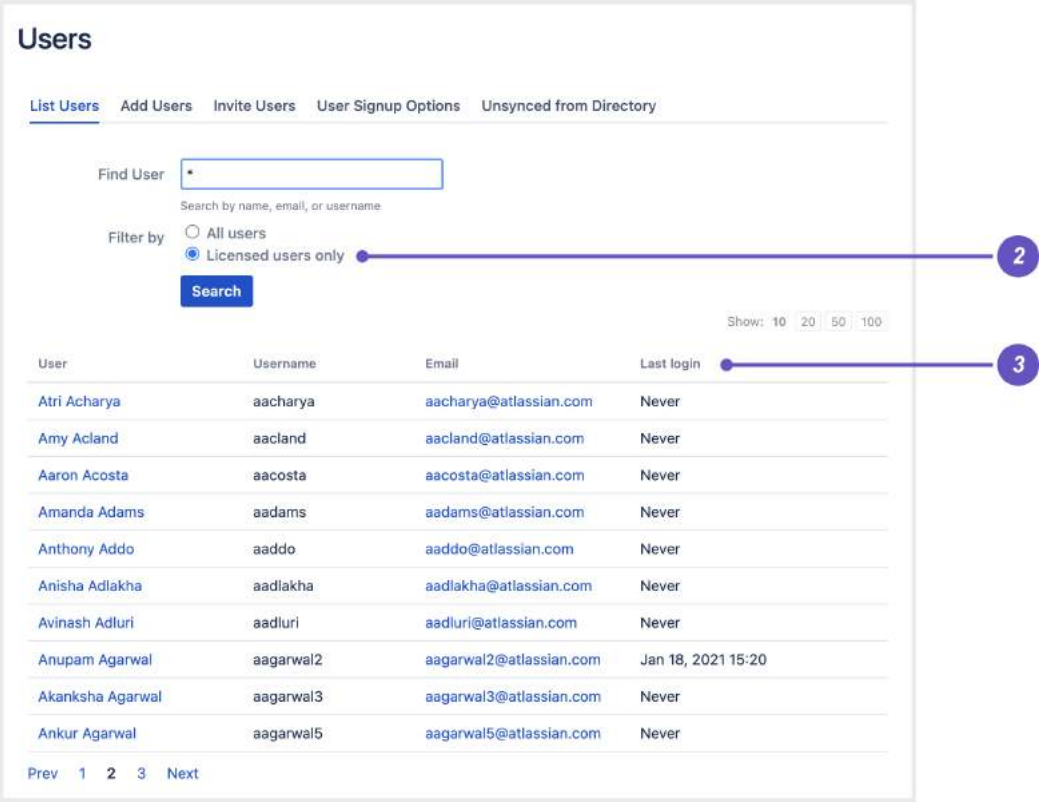

Exceeding your licensed user count

If you exceed the number of users included in your license, your Confluence instance will become read-only, that means no users will be able to create or edit content until you reduce the number of users.

Reducing your user count

You can reduce your user count by removing or deactivating users who do not require access to Confluence. SeeDelete or Disable Users.

If you haveconnected Confluence to an LDAP directory, you may want configure Confluence to only synchronize a subset of users from LDAP rather than all users. SeeHow to change the number of users synchronized from LDAP to Confluencein the Knowledge Base. This can be a complicated process and we recommend that you only use this method if necessary.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Downgrading your license

If you decide to downgrade your Confluence license to pay for fewer users you need to ensure that your licensed user count does not exceed the total allowed, before applying for the new license.

If you have more users than your new license allows, you will need to reduce your user count before applying for the new license.

Finding your Support Entitlement Number (SEN)

You can find your Support Entitlement Number (SEN) in three places:

- In Confluence - go to > **General Configuration** > **License Details**)
- At **my.atlassian.com**
- On your Atlassian invoice.

See How to find your Support Entitlement Number (SEN)for more general information about how Atlassian Support uses this number.

What happens when your maintenance or subscription expires

Important changes to our server and Data Center products

Weve ended sales for new Server licenses, and will end support for server on February 15, 2024. Were continuing our investment in Data Center with several key improvements. Learn what this means for you

Server licenses

If you have a Confluence Server license, your maintenance entitles you to access Atlassian support, and upgrade Confluence.

When your maintenance expires, you can still use Confluence, but you'll no longer be able to contact Atlassian support, or upgrade to a version of Confluence released after your maintenance expiry date.

Data Center licenses

Confluence Data Center is offered as a subscription (also known as a fixed term license), which includes access to support and version upgrades.

If your subscription expires, Confluence will become read-only, which means you'll be able to view pages, but not create or edit them.

Our licensing policy can change from time to time, so it's best to check our Purchasing and Licensing FAQ for the latest information.

Need more information about your Server or Data Center license? Get in touch

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

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
- Backup and Restore
- Attachment Storage Configuration
- Confluence Data Model
- Finding Unused Spaces or Pages
- Data Import and Export
- Import a Text File
- Auditing in Confluence
- Set retention rules to delete unwanted data
- Data pipeline

Related pages:

- Managing System and Marketplace Apps
- Integrating Confluence with Other Applications
- Getting Started as Confluence Administrator
- Confluence administrator's guide

Check the amount of data in your site

To check the amount of data in your site:

1. Go to > **General Configuration** > **System information**.
2. Scroll down to the **Confluence usage** section.

This will tell you the number of:

- Total Spaces - total number of spaces, including site spaces and personal spaces
- Site Spaces
- Personal Spaces
- Content (All Versions) - total number of content versions (includes the latest version and all historical versions). Content includes pages, comments, blogs, attachments, personal information, space description, personal space description, spaces, and drafts. Content in your trash will also contribute to this count
- Content (Current Versions) - total number of content (only counts the latest version). See earlier point for what is considered content. Content in your trash will also contribute to this count
- Local Users
- Local Groups

Clean up unwanted data

As your team grows, so does the data being created and stored in Confluence. Find outthe ways you can declutter and reduce the overall size of your Confluence site in our Cleanup guide.

Database Configuration

This document provides information on connecting Confluence to an external database.

On this page:

- Choosing an external database
- About the embedded H2 database
- Database setup
- Database drivers
- Database connection methods
- Database troubleshooting

Related pages:

- Database JDBC Drivers
- Supported Platforms
- Embedded H2 Database
- Managing Confluence Data

Choosing an external database

Note: Take time to choose your database wisely. The XML backup built into Confluence is not suited for migration or backup of large data sets. If you need to migrate later, you may need to use a third party database migration tool.

Below is more information on selecting and migrating to an external database:

- Migrating to a Different Database
- Supported Databases
- Database Troubleshooting

About the embedded H2 database

The embedded H2 database is**only**supported for testing and app development purposes onnon-clustered (single node)Confluence Data Center installations.

To find out if you are still using the embedded database, go to > **General Configuration** > **Troubleshooting and support tools**.

Database setup

To find out how to set up your database, see:

- Database Setup for Oracle
- Database Setup For MySQL
- Database Setup for PostgreSQL
- Database Setup for SQL Server
- Configuring Confluence Data Center to work with Amazon Aurora

Database drivers

You must use a supported database driver. SeeDatabase JDBC Drivers for the drivers we support.

If you attempt to use an unsupported or custom JDBC driver (or `adriverClassName`from an unsupported or custom driver in your JINDI datasource connection) collaborative editing will fail.

Database connection methods

Confluence connects to your database using a JDBC URL. The Confluence Setup Wizard will establish this connection by default (this won't be shown) as this is the recommended connection method.

If you want to use a JNDI datasource,seeConfiguring a datasource connectionfor the steps you'll need to take before you set up Confluence, asthe setup wizard will only provide the option to use a datasource if it detects a datasource in your Tomcat configuration.

Database troubleshooting

For database-related problems seeDatabase Troubleshooting.

If you need more help, check outTroubleshooting Problems and Requesting Technical Support.

Database JDBC Drivers

This page provides the download links for the JDBC drivers for all supported databases.

Due to licensing constraints, we are not able to bundle MySQL or Oracle database drivers with Confluence, so you will need to manually download and install the driver listed below before you can set up Confluence.

If you use PostgreSQL or Microsoft SQL Server, the drivers are bundled with Confluence, so you're ready to.

On this page:

- Adding your database driver (MySQL and Oracle)
- Supported drivers

Related pages:

- Database Configuration
- Supported Platforms

Adding your database driver (MySQL and Oracle)

The Confluence setup wizard will stop you at the Database configuration step if it can't find an appropriate driver for the database you select.

To make your database driver available to Confluence:

1. Stop Confluence.
2. Download and extract the appropriate driver from the list below.
3. Drop the .jar file in your`<installation-directory>/confluence/WEB-INF/lib` directory.
4. Restart Confluence then go to `http://localhost:<port>`in your browser to continue the setup process.

The setup wizard will return to the database configuration step, and you're back on your way.

Supported drivers

Database	Driver bundled?	JDBC drivers	Notes	More information
PostgreSQL	✅	**Postgres JDBC driver download (latest)**	We recommend that you use the bundled JDBC **42.5.3**driver. If you want to use a later driver, you can download it from the PostgreSQL website.	Database Setup for PostgreSQL
Microsoft SQL Server	✅	**Microsoft JDBC Driver for SQL Server download**	We recommend that you use the bundled **Type 4** JDBC driver. If you decide to use a later version, we may not be able to provide support for any problems you encounter.	Database setup for Microsoft SQL Server
	❌	**jTDS 1.3.1 driver download**	⚠️ This driver is deprecated. New Confluence installations use the Microsoft JDBC Driver for SQL Server (above). If you're upgrading an existing Confluence site to Confluence 6.4 you should continue to use the bundled jTDS driver. We'll help you migrate to the Microsoft driver in a later release.	

MySQL 8.0	✗	**Connector\J 8.0 driver download**	Due to licensing constraints, MySQL drivers are not bundled with Confluence. Confluence is currently tested with the **8.0.33**driver.	Database setup for MySQL
Oracle	✗	**JDBC driver downloads**	Due to licensing constraints, Oracle drivers are not bundled with Confluence. For Oracle 19c you can use either ojdbc8.jar or ojdbc10.jar. We recommend using the thin drivers only. See the Oracle JDBC driver FAQ.	Database setup for Oracle

If you attempt to use an unsupported or custom JDBC driver (or a `driverClassName`from an unsupported or custom driver in your JINDI datasource connection) collaborative editing will fail. You must use a supported driver.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Database Setup for Oracle

This page provides instructions for configuring Confluence to use an Oracle database.

On this page:

- Before you start
- 1. Install Oracle
- 2. Create database user
- 3. Install Confluence
- 4. Download and install the Oracle thin driver
- 5. Enter your database details
- Troubleshooting

Related pages:

- Database Configuration
- Known Issues for Oracle
- Confluence installation and upgrade guide

Before you start

- SeeSupported Platformstocheck your version of Oracle is supported. You may need to upgrade your database before installing Confluence.
- If you're switching from another database, including the embedded evaluation database, readMigrating to Another Databasebefore you begin.

> You'll need an experienced Oracle database administrator (DBA) to set up and maintain your database.
>
> Our support team can assist with Confluence problems, but are unable to help you administer your Oracle database.
>
> If you don't have access to an experienced Oracle DBA, consider using a different supported database.

1. Install Oracle

If you don't already have an operational Oracle server,downloadand install it now. See theOracle documentationfor instructions.

When setting up your Oracle server:

- **Character encoding** must be set toAL32UTF8 (this the Oracle equivalent of Unicode UTF-8).
- **Collation** should be set to `BINARY`.

2. Create database user

To create the user and assign its privileges:

1. Use the`sqlplus` command to access Oracle via the command line

   ```
   sqlplus user/password <as sysdba|as sysoper>
   ```

   If you're logging in with the user 'sys' you'll need to include the "as sysdba" or "as sysoper" to determine which sys role you want to use.
2. Create a Confluence user (for example `confluenceuser`). It's important that this user is only granted the required privileges:

   ```
   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>;
   ```

- Specify the `tablespace` for the table objects as shown above.
- The `connect` role is required to set up a connection.
- The `resource` role is required to allow the user to create objects in its own schema. The `resource` role includes `create table`, `create sequence`, and `create trigger` by default. If you've altered the `resource` role to remove these, you'll need to grant these privileges to the user directly, or through some other role.
- Don't grant the `select any table` permission as this can cause problems with other schemas.

3. Install Confluence

Check out the Confluence Installation Guide for step-by-step instructions on how to install Confluence on your operating system.

4. Download and install the Oracle thin driver

Due to licensing restrictions, we're not able to bundle an Oracle driver with Confluence. To make your database driver available to Confluence:

1. Stop Confluence.
2. Head to Database JDBC Drivers and download the appropriate driver. The driver file will be called something like `ojdbc8.jar`
3. Drop the .jar file in your `<installation-directory>/confluence/WEB-INF/lib` directory.
4. Restart Confluence then go to `http://localhost:<port>` in your browser to continue the setup process.

5. Enter your database details

The Confluence setup wizard will guide you through the process of connecting Confluence to your database.

Use a JDBC connection (default)

JDBC is the recommended method for connecting to your database.

The Confluence setup wizard will provide you with two setup options:

- **Simple** - this is the most straightforward way to connect to your database.
- **By connection string** - use this option if you want to specify additional parameters and are comfortable constructing a database URL.

Depending on the setup type, you'll be prompted for the following information.

Setup type	Field	Description
Simple	Hostname	This is the hostname or IP address of your database server.
Simple	Port	This is the Oracle port. If you didn't change the port when you installed Oracle, it will default to `1521`.
Simple	Service name	This is the service name (of your confluence database.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

By connection string	Database URL	The database URL is entered in this format: `jdbc:oracle:thin:@//<HOST>:<PORT>/<SERVICE>` <SERVICE> can be either the SID or Service Name. For example: `jdbc:oracle:thin:@//localhost:1521/confluence` By default, we use the new style URL provided by the thin driver. You can also use the `tnsnames` style.
Both	Username	This is the username of your dedicated database user. In the example above, this is `confluenceuser`.
Both	Password	This is the password for your dedicated database user.

To determine the host, port, service name, and/or SID, execute the following command as the user running Oracle (usually 'Oracle'):

```
lsnrctl status
```

Here's an example of the output:

```
SNRCTL for Linux: Version 11.2.0.2.0 - Beta on 29-JUN-2012 15:20:59
Copyright (c) 1991, 2010, Oracle.  All rights reserved.
Connecting to (DESCRIPTION=(ADDRESS=(PROTOCOL=IPC)(KEY=EXTPROC_FOR_XE)))
STATUS of the LISTENER
------------------------
Alias                     LISTENER
Version                   TNSLSNR for Linux: Version 11.2.0.2.0 - Beta
Start Date                06-JUN-2012 08:36:34
Uptime                    23 days 6 hr. 44 min. 25 sec
Trace Level               off
Security                  ON: Local OS Authentication
SNMP                      OFF
Default Service           XE
Listener Parameter File   /u01/app/oracle/product/11.2.0/xe/network/admin/listener.ora
Listener Log File         /u01/app/oracle/diag/tnslsnr/<HOSTNAME>/listener/alert/log.xml
Listening Endpoints Summary...
  (DESCRIPTION=(ADDRESS=(PROTOCOL=ipc)(KEY=EXTPROC_FOR_XE)))
  (DESCRIPTION=(ADDRESS=(PROTOCOL=tcp)(HOST=<HOSTNAME>)(PORT=1521)))
  (DESCRIPTION=(ADDRESS=(PROTOCOL=tcp)(HOST=<HOSTNAME>)(PORT=8080))(Presentation=HTTP)(Session=RAW))
Services Summary...
Service "PLSExtProc" has 1 instance(s).
  Instance "PLSExtProc", status UNKNOWN, has 1 handler(s) for this service...
Service "XE" has 1 instance(s).
  Instance "XE", status READY, has 1 handler(s) for this service...
Service "XEXDB" has 1 instance(s).
  Instance "XE", status READY, has 1 handler(s) for this service...
The command completed successfully
```

- The host and port are determined by the line containing `PROTOCOL=tcp` (the line **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, if 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 will be a direct JDBC connection.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Seethe Oracle JDBC FAQ for more information on Oracle JDBC URLs.

6. Test your database connection

In the database setup screen, hit the**Test connection**button to check:

- that Confluence can connect to your database server
- that the database character encoding is correct
- that your database user has appropriate permissions for the database
- that your database user has NOT been granted the SELECT ANY TABLE privilege

Once the test is successful, hit**Next**to continue with the Confluence setup process.

Troubleshooting

- If Confluence complains that it is missing a class file, you may have placed the JDBC driver in the wrong folder.
- The following page contains common issues encountered when setting up your Oracle database to work with Confluence:Known Issues for Oracle.
- There's a known issue when running Oracle with Native Network Encryption that can cause Confluence to become unresponsive. See ~~CONFSERVER-60152~~ CLOSED for more detailsandConfluence Unresponsive Due to High Database Connection Latency for some suggested mitigation strategies.
- There's a known issue when username or schema names contain dots. See CONFSERVER-60274 GATHERING IMPACT for more information.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Database Setup for PostgreSQL

This page provides instructions for configuring Confluence to use a PostgreSQL database.

On this page:

- Before you start
- 1. Install PostgreSQL
- 2. Create a database user and database
- 3. Install Confluence
- 4. Enter your database details
- Troubleshooting

Related pages:

- Database Configuration
- Known issues for PostgreSQL

Before you start

- SeeSupported Platformstocheck your version of PostgreSQL is supported. You may need to upgrade your database before installing Confluence.
- If you're switching from another database, including the embedded evaluation database, readMigrating to Another Databasebefore you begin.

1. Install PostgreSQL

If you don't already have PostgreSQL installed,downloadand install it now.

A few tips when installing PostgreSQL:

- The **password** you provide during the installation process is for the 'postgres' account, which is the database root-level account (the super user). Remember this username and password asyou'll need it each time you log in to the database.
- The **default port** for PostgreSQL is 5432. If you decide to change the default port, make sure it does not conflict with any other services running on that port.
- Choose the **locale** that best matches your geographic location.
- Don't launch **Stack Builder** at the completion of the installer.

2. Create a database user and database

Once you've installed PostgreSQL:

1. Create a database user, for example`confluenceuser`.
    - Your new user must be able to**create database objects**and must have **can login** permission.
2. Next, create a database (for example`confluence`):
    - **Owner**is your new database user (for example`confluenceuser`)
    - **Character encoding**must be set to`utf8`encoding.
    - **Collation** must also be set to `utf8`. Other collations, such as "C", are known to cause issues with Confluence.
      If you are running PostgreSQL on Windows use the equivalent character type and collation for your locale, for example`English_United States.1252`
    - In Linux systems, if the locale is not utf8, include LC_CTYPE as utf8 during database creation.

You can usepgAdminas an alternative to the command line to complete this step.

3. Install Confluence

Check out theConfluence Installation Guidefor step-by-step instructions on how to install Confluence on your operating system.

4. Enter your database details

The Confluence setup wizard will guide you through the process of connecting Confluence to your database. Be sure to select "**My own database**".

Use a JDBC connection (default)

JDBC is the recommended method for connecting to your database.

The Confluence setup wizard will provide you with two setup options:

- **Simple**- this is the most straightforward way to connect to your database.
- **By connection string**- use this option if you want to specify additional parameters and are comfortable constructing a database URL.

Depending on the setup type, you'll be prompted for the following information.

Setup type	Field	Description
Simple	Hostname	This is the hostname or IP address of your database server.
Simple	Port	This is the PostgreSQL port. If you didn't change the port when you installed Postgres, it will default to `5432`.
Simple	Database name	This is the name of your confluence database. In the example above, this is `confluence`

By connection string	Database URL	The database URL is entered in this format:
		`jdbc:postgresql://<server>:<port>/<database>`
		
		For example:
		`jdbc:postgresql://localhost:5432/confluence`
		
		If you need to connect to an SSL database, add the `sslmode=require` parameter in the database URL. For example:
		`jdbc:postgresql://localhost:5432/confluence?sslmode=require`

Both	Username	This is the username of your dedicated database user. In the example above, this is `confluenceuser`.
Both	Password	This is the password for your dedicated database user.

5. Test your database connection

In the database setup screen, hit the **Test connection** button to check:

- that Confluence can connect to your database server
- that the database character encoding is correct
- that your database user has appropriate permissions for the database

Once the test is successful, hit **Next** to continue with the Confluence setup process.

If Confluence and PostgreSQL are hosted on different servers, see the PostgreSQL documentation on how to set up pg_hba.conf to make sure Confluence and PostgreSQL can communicate remotely.

Troubleshooting

- If Confluence complains that it is missing a class file, you may have placed the JDBC driver in the wrong folder.
- If you're unable to connect to the database from Confluence and they are on different machines, most likely you have a firewall in between the two machines or your `pg_hba.conf` file is misconfigured.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

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.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Database Setup for SQL Server

This page provides instructions for configuring Confluence to use a Microsoft SQL Server database.

On this page:

- Before you start
- 1. Install SQL Server
- 2. Create a database and database user
- 3. Install Confluence
- 4. Enter your database details
- Database driver changes
- Troubleshooting

Related pages:

- Database Configuration
- Known issues for SQL Server
- Confluence installation and upgrade guide

Before you start

Check the following before you start:

- SeeSupported Platformstocheck your version of SQL Server is supported. You may need to upgrade your database before installing Confluence.
- If you're switching from another database, including the embedded evaluation database, readMigrating to Another Databasebefore you begin.

1. Install SQL Server

If you don't already have Microsoft SQL Server installed,downloadand install it now. SeeInstallation for SQL Server onMSDN for step-by-step instructions.

SQL Server allows two types of authentication: SQL Server Authentication and Windows Authentication. To make sure Confluence will be able to connect to your database you'll need to set your SQL server to allow Mixed Authentication (both SQL Server and Windows modes). This setup is generally found under Properties > Security > Server Authentication.

2. Create a database and database user

Once you've installed SQL Server, create a database user and database for Confluence as follows:

1. Using your SQL administrator permissions, create a new database (for example`confluence`)
2. Set the default collation for the database to**SQL_Latin1_General_CP1_CS_AS**(case sensitive).

```
ALTER DATABASE <database-name> COLLATE SQL_Latin1_General_CP1_CS_AS
```

If you seea 'database could not be exclusively locked to perform the operation' error, you may need to prevent other connections by setting the mode to single user for the transaction

```
ALTER DATABASE <database-name> SET SINGLE_USER WITH ROLLBACK IMMEDIATE;
<your ALTER DATABASE query>
ALTER DATABASE <database-name> SET MULTI_USER;
```

3. Check the database isolation level of READ_COMMITTED_SNAPSHOT is ON.

```
SELECT is_read_committed_snapshot_on FROM
sys.databases WHERE name= 'database-name'
```

If this query returns **1**, thenREAD_COMMITTED_SNAPSHOT is ON, and you're good to go.

If this query returns **0**,READ_COMMITTED_SNAPSHOT option is OFF and you will need to turn it on as follows:

```
ALTER DATABASE <database-name>
   SET READ_COMMITTED_SNAPSHOT ON
   WITH ROLLBACK IMMEDIATE;
```

4. Using your SQL administrator permissions,create a new SQL user account for Confluence(for example,`confluenceuser`).
5. Give this user the default schema as follows:

```
ALTER USER <confluenceuser> WITH DEFAULT_SCHEMA = dbo
```

6. Make sure this user has full create, read and write permissions for the database tables. Confluence must be able to create its own schema, and have the ability to create/drop triggers and functions.Refer to the SQL Server documentation for more information.

3. Install Confluence

Check out theConfluence Installation Guidefor step-by-step instructions on how to install Confluence on your operating system.

4. Enter your database details

The Confluence setup wizard will guide you through the process of connecting Confluence to your database.

Use a JDBC connection (default)

JDBC is the recommended method for connecting to your database.

The Confluence setup wizard will provide you with two setup options:

- **Simple**- this is the most straightforward way to connect to your database.
- **By connection string**- use this option if you want to specify additional parameters and are comfortable constructing a database URL.

Depending on the setup type, you'll be prompted for the following information.

Setup type	Field	Description
Simple	Hostname	This is the hostname or IP address of your database server.
Simple	Port	This is the SQL Server port. If you didn't change the port when you installed SQL Server, it will default to `1433`.
Simple	Database name	This is the name of your confluence database. In the example above, this is c`onfluence`

Simple	Instance name	To find out your instance name, connect to your database and run one of the following:
		`select @@SERVICENAME;`
		`SELECT SERVERPROPERTY('InstanceName');`
		If you have a default named instance setup in SQL Server, you won't need to specify this parameter.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

By connection string	Database URL	The database URL is entered in this format:
		`jdbc:sqlserver://<hostname>:<port>;databaseName=<database>`
		
		For example:
		`jdbc:sqlserver://yourserver:1433;databaseName=confluence`

Both	Username	This is the username of your dedicated database user. In the example above, this is `confluenceuser`.
Both	Password	This is the password for your dedicated database user.

5. Test your database connection

In the database setup screen, hit the**Test connection**button to check:

- Confluence can connect to your database server
- the database collation and isolation level is correct
- your database user has appropriate permissions for the database

Once the test is successful, hit**Next**to continue with the Confluence setup process.

Database driver changes

In Confluence 6.6 we replaced the open source jTDS driver for Microsoft SQL Server with the official Microsoft JDBC Driver for SQL Server. You will be automatically migrated to the new driver when you upgrade to 6.6 or later.

If for some reason the automatic migration fails, you'll need to make this change manually. See Migrate from the jTDS driver to the supported Microsoft SQL Server driver in Confluence 6.4 or later.

Troubleshooting

- If you get the following error message, check you've given the `confluenceuser` user all the required database permissions when connecting from `localhost`.

```
Could not successfully test your database: : Server connection failure during transaction. Due to underlying exception: 'java.sql.SQLException: Access denied for user 'confluenceuser'@'localhost' (using password: YES)'
```

- You may need to open additional ports. See thisMicrosoft KBabout the ports required for SQL Server.
- The following page contains common issues encountered when setting up your SQL Server database to work with Confluence:Known Issues for SQL Server.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Database Setup For MySQL

This page provides instructions for configuring Confluence to use a MySQL database.

On this page:

- Before you start
- 1. Install MySQL Server
- 2. Configure MySQL Server
- 3. Create database and database user
- 4. Install Confluence
- 5. Download and install the MySQL driver
- 6. Enter your database details
  - Use a JDBC connection (default)
  - 7. Test your database connection
- Upgrade your database and driver
- Troubleshooting

Related pages:

- Database Configuration
- Database Troubleshooting for MySQL
- Confluence installation and upgrade guide

Before you start

- SeeSupported Platformstocheck your version of MySQL is supported. You may need to upgrade your database before installing Confluence.
- If you're switching from another database, including the embedded evaluation database, readMigrating to Another Databasebefore you begin.

> Confluence will not work on MySQL variants such as MariaDB (CONFSERVER-29060) and Percona Server (CONFSERVER-36471)

1. Install MySQL Server

If you don't already have MySQL installed, downloadand install it now.See the MySQL documentationfor step-by-step instructions.

2. Configure MySQL Server

In this step, you will configure your MySQL database server.

***Note:** If you intend to connect Confluence to an existing MySQL database server, we strongly recommend that you reconfigure this database server by running through the configuration steps in the MySQL installation wizard as described below .*

> These instructions apply to Confluence 7.3 and later. Using an earlier version? See Database Setup For MySQL in Confluence 7.2 and earlier.

To configure MySQL Server:

1. Run the MySQL installation wizard:
   a. If you are connecting Confluence to your existing MySQL server, choose **Reconfigure Instance**.
   b. Choose **Advanced Configuration**.
   c. Choose the**type of MySQL Server** that best suits your hardware requirements. This will affect the MySQL Server's usage of memory, disk and CPU resources. Refer to the MySQL documentation for further information.
   d. Choose **Transactional Database Only** to ensure that your MySQL database will use**InnoDB** as its default storage engine.
   You must use the InnoDB storage engine with Confluence. Using the MyISAM storage enginec an lead to data corruption in Confluence.
   e. Set the **InnoDB Tablespace** settings to your requirements. (The default settings are acceptable.)
   f. Set the approximate**number of concurrent connections** permitted to suit your Confluence usage requirements. You can use one of the presets or enter a number manually. Refer to the MySQL documentation for further information.

g. For the **networking options**, ensure the **Enable TCP/IP Networking** and **Enable Strict Mode** options are selected (default). Refer to the MySQL documentation on setting the networking and server SQL modes for further information.
h. For the MySQL server's **default character set**, choose **Best Support For Multilingualism** (in other words, utf8mb4). This will ensure Confluence's support for internationalization. For more information, see Configuring Database Character Encoding.
i. For the Windows configuration option, choose whether or not to install the MySQL Server as a Windows service. If your hardware is going to be used as a dedicated MySQL Server, you may wish to choose the options to **Install As Windows Service** (and **Launch the MySQL Server automatically**). Refer to the MySQL documentation for further information.
Note: *If you choose not to install the MySQL Server as a Windows Service, you will need to ensure that the database service has been started before running Confluence.*
j. Select **Modify Security Settings** to enter and set your MySQL Server (root) access password.

2. Edit the `my.cnf` file (`my.ini` on Windows operating systems) in your MySQL server. Locate the `[mysqld]` section in the file, and add or modify the following parameters:
(Refer to MySQL Option Files for detailed instructions on editing `my.cnf` and `my.ini`.)
Locate the `[mysqld]` section in the file, and add or modify the following parameters:
   - Specify the default character set to be utf8mb4:

```
[mysqld]
...
character-set-server=utf8mb4
collation-server=utf8mb4_bin
...
```

   - Set the default storage engine to InnoDB:

```
[mysqld]
...
default-storage-engine=INNODB
...
```

   - Specify the value of `max_allowed_packet` to be at least 256M:

```
[mysqld]
...
max_allowed_packet=256M
...
```

   - Specify the value of `innodb_log_file_size` to be at least 2GB:

```
[mysqld]
...
innodb_log_file_size=2GB
...
```

   - Ensure the sql_mode parameter does not specify NO_AUTO_VALUE_ON_ZERO

```
// remove this if it exists
sql_mode = NO_AUTO_VALUE_ON_ZERO
```

   - Ensure that the global transaction isolation level of your Database had been set to READ-COMMITTED.

```
[mysqld]
...
transaction-isolation=READ-COMMITTED
...
```

   - Check that the binary logging format is configured to use 'row-based' binary logging, and that your database user can create and alter stored functions.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

```
[mysqld]
...
binlog_format=row
log_bin_trust_function_creators = 1
...
```

- If you're using MySQL 5.7, turn off the 'derived merge' optimizer switch, as this can cause the dashboard to load slowly.

```
optimizer_switch = derived_merge=off
```

3. Restart your MySQL server for the changes to take effect:
   - On Windows, use the Windows Services manager to restart the service.
   - On Linux:
     - Run one of the following commands, depending on your setup: '`/etc/init.d /mysqld stop`' or '`/etc/init.d/mysql stop`' or '`service mysqld stop`'.
     - Then run the same command again, replacing '`stop`' with '`start`'.
   - On Mac OS X, run '`sudo /Library/StartupItems/MySQLCOM/MySQLCOM restart`'.

3. Create database and database user

Once you've installed and configured MySQL, create a database user and database for Confluence as follows:

1. Run the '`mysql`' command as a MySQL super user. The default user is 'root' with a blank password.
2. Create an empty Confluence database schema (for example `confluence`):

```
CREATE DATABASE <database-name> CHARACTER SET utf8mb4 COLLATE utf8mb4_bin;
```

3. Create a Confluence database user (for example `confluenceuser`):

```
GRANT ALL PRIVILEGES ON <database-name>.* TO '<confluenceuser>'@'localhost' IDENTIFIED BY
'<password>';
```

   If Confluence is not running on the same server, replace localhost with the hostname or IP address of the Confluence server.

4. Install Confluence

Check out the Confluence Installation Guide for step-by-step instructions on how to install Confluence on your operating system.

5. Download and install the MySQL driver

Due to licensing restrictions, we're not able to bundle the MySQL driver with Confluence. To make your database driver available to Confluence follow the steps below for your MySQL version.

MySQL 5.7

1. Stop Confluence.
2. Head to Database JDBC Drivers and download the appropriate driver. The driver file will be called something like `mysql-connector-java-5.1.xx-bin.jar`
3. Drop the .jar file in your `<installation-directory>/confluence/WEB-INF/lib` directory.
4. Restart Confluence then go to `http://localhost:<port>` in your browser to continue the setup process.

MySQL 8.0

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

You can't use MySQL 8.0 with Confluence 7.1 or earlier.

1. Stop Confluence.
2. Head toDatabase JDBC Driversanddownload the appropriate driver for MySQL 8. The driver file will be called something like`mysql-connector-java-8.0.xx-bin.jar`
3. Drop the .jar file in your`<installation-directory>/confluence/WEB-INF/lib`directory.
4. Restart Confluence then go to`http://localhost:<port>`in your browser to continue the setup process.

6. Enter your database details

The Confluence setup wizard will guide you through the process of connecting Confluence to your database.

Use a JDBC connection (default)

JDBC is the recommended method for connecting to your database.

The Confluence setup wizard will provide you with two setup options:

- **Simple**- this is the most straightforward way to connect to your database.
- **By connection string**- use this option if you want to specify additional parameters and are comfortable constructing a database URL.

Depending on the setup type, you'll be prompted for the following information.

Setup type	Field	Description
Simple	Hostname	This is the hostname or IP address of your database server.
Simple	Port	This is the MySQL port. If you didn't change the port when you installed MySQL, it will default to `3306`.
Simple	Database name	This is the name of your confluence database. In the example above, this is `confluence`

By connection string	Database URL	The database URL is entered in this format:
		`jdbc:mysql://<hostname>:<port>/<database>`
		
		For example:
		`jdbc:mysql://localhost:3306/confluence`

Both	Username	This is the username of your dedicated database user. In the example above, this is `confluenceuser`.
Both	Password	This is the password for your dedicated database user.

7. Test your database connection

In the database setup screen, hit the**Test connection**button to check:

- Confluence can connect to your database server
- the database character encoding, collation, isolation level and storage engine are correct
- your database user has appropriate permissions for the database.

Once the test is successful, hit**Next**to continue with the Confluence setup process.

Upgrade your database and driver

If you upgrade MySQL you may also need to upgrade the database driver Confluence uses to connect to your database. Always use the driver recommended on theDatabase JDBC Drivers page.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Before you begin, back up your database, Confluence installation directory and Confluence home directory. We strongly recommend you test your changes in astaging environment first.

To upgrade your database driver:

1. Stop Confluence.
2. Go to<installation-directory>/confluence/WEB-INF/lib/ and delete your existing driver. It will be called something like`mysql-connector-java-x.x.xx-bin.jar`
3. Drop the new driver .jar file in your`<installation-directory>/confluence/WEB-INF/lib`directory.
4. Upgrade your MySQL server.
5. Restart Confluence.

If you're using a datasource connection, you may need to also update the driver classname in the datasource.

Troubleshooting

- There is a known issue when running Confluence with MySQL 8.0.29 and later due to a change to the UTF8 alias in MySQL. We're working on a fix, but if you have Confluence 7.3 or later, you can change the character set and collation to UTF8MB4 to avoid this issue. See How to Fix the Collation and Character Set of a MySQL Database manually.
- There is a known issue with MySQL 5.7 where parts of the dashboard can take a very long time to load. You may need to turn the "derived merge" optimizer switch off in your database configuration. See CONFSERVER-54984 READY FOR DEVELOPMENT for details of the workaround.
- If Confluence complains that it is missing a class file, you may have placed the JDBC driver in the wrong folder.
- If you get the following error message, verify that you have given the `confluenceuser` user all the required database permissions when connecting from `localhost`.

```
Could not successfully test your database: : Server connection failure during transaction. Due to
underlying exception: 'java.sql.SQLException: Access denied for user 'confluenceuser'@'localhost'
(using password: YES)'
```

- The following page contains common issues encountered when setting up your MySQL database to work with Confluence:Database Troubleshooting for MySQL

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Embedded H2 Database

> ⚠ We ended support for the H2 database in Confluence 8.0. If you wish to continue your development cycle with the H2 database in Confluence 8.0 or later, you will need to upgrade to AMPS 8.6.0.

From Feb 2, 2021 (server end of sale date) you will only be able to generate Data Center evaluation licenses. This means the Confluence Setup Wizard wont include an option to use an embedded H2 database.

On this page:

- Connect to the embedded H2 database using DB Visualizer
- Connect to the embedded H2 database using the H2 console
- Remote connections
- Migrate to a supported external database

Related pages:

- Confluence Home and other important directories
- Database Configuration

The embedded H2 database is**only**fortesting and app development purposesonnon-clustered (single node) Confluence Data Center installations.

To find out if you are still using the embedded database, go to > **General Configuration** > **Troubleshooting and support tools**.

The embedded database files are stored in your Confluence home directory`<confluence-home>/database`.

Connect to the embedded H2 database using DB Visualizer

If you need to make changes directly in the database, and you're using the H2 database, here's how you can connect to it using DBVisualizer.

DBVisualizer is just one database administration tool. You can use any administration tool that supports embedded H2 databases. The steps will be similar.

1. Shut down Confluence.
2. Back up your`<confluence-home>/database`directory.
3. LaunchDBVisualizer.
4. Choose**Create new database connection**and follow the prompts to set up the connection.
   The information you'll need is:
   - **Database driver:**H2 embedded
   - **Database Userid:**sa
   - **Database password:**leave this field blank
   - **Database filename:**`<confluence-home>/database/h2db`
     ⚠leave off the`.h2.db`file extension.
5. Connect to the database.

Refer to theDBVisualizer documentationfor help using DBVisualizer.

Connectto the embedded H2 databaseusing the H2 console

Alternatively you can connect using the browser based H2 console.The easiest way to access the console is to double click the H2 database jar file at`<installation-directory>\confluence\WEB-INF\lib\h2-x.x.x.jar`.

Remote connections

Remote connections to the embedded H2 database are not permitted. You can only connect to H2 from the server on which Confluence is installed.

Plugin vendors can connect remotely when Confluence is running in dev mode, but admins should not use this as a workaround, and instead should migrate to a supported external database.

Note:The H2 database doesnt work on amulti-node Confluence cluster. A shared database is required for a multi-node cluster.

Migrate to a supported external database

If you're using the H2 database, but running Confluence as a production system, you should migrate to a supported database as soon as possible.

To migrate to a supported external database:

1. CheckSupported Platforms to find out which databases and versions are supported.
2. Head toMigrating to Another Database for a step-by-step guide.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Migrating to Another Database

This document describes how to migrate your Confluence data from your existing database to another database. The instructions are designed primarily for migrating from an evaluation to a production database.

Large data sets will require third party database migration tools.

This page covers the following scenarios:

- Moving from the embedded, trial database to a supported external database.
- Moving from one external database to another, for example from Oracle to PostgreSQL (provided your dataset is not large)
- Upgrading to a new version of the same external database. Note: you don't need to migrate your data if you're upgrading the database in place.

If you are moving your database from one server to another you can change the JDBC URL in `<confluence-home>/confluence.cfg.xml` (if you are using a direct JDBC connection) or in the definition of your datasource (if you are connecting via a datasource).

Please see the note at the bottom of this page, if you are migrating **MySQL** database from one server to another.

On this page:

- Limitations of database migration
- Database migration
- Migrating to an Amazon Aurora database
- Method one standard procedure
  - Step 1: Take note of your Marketplace apps
  - Step 2: Back up your data
  - Step 3: Set up the new database
  - Step 4. Install Confluence (same version number) in a new location
  - Step 5. Download and install the database driver if necessary
  - Step 6. Run the Confluence setup wizard and copy your data to your new database
  - Step 7. Re-install your Marketplace apps
  - Step 8. Check settings for new machine
- Method two for installations with a large volume of attachments
  - Before you start
  - Step 1: Take note of your Marketplace apps
  - Step 2: Back up your data
  - Step 3: Set up the new database
  - Step 4. Install Confluence (same version number) in a new location
  - Step 5. Download and install the database driver if necessary
  - Step 6. Run the Confluence setup wizard and copy your data to your new database
  - Step 7: Copy your attachments across
  - Step 8. Re-install your Marketplace apps
  - Step 9. Check settings for new machine
- A note about case sensitivity in your database
  - Setting up a new Confluence instance
  - Migrating an existing Confluence instance to a different database
- Migrating MySQL database between servers
- Troubleshooting

Related pages:

- Database Configuration
- Confluence Home and other important directories

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.

Migrating to an Amazon Aurora database

If you plan to migrate to an Amazon Aurora database, seeConfiguring Confluence Data Center to work with Amazon Aurora. This guide explains how to migrate to an Amazon Aurora cluster and connect it to Confluence Data Center.

Method one standard procedure

Step 1: Take note of your Marketplace apps

Take note of the apps (also knowns as plugins or add-ons) currently installed and enabled in Confluence, so that you can reinstate them later. Make a note of the following for each app:

- App name and vendor
- Version
- Enabled or disabled status. This is useful if you have enabled or disabled modules yourself, making your configuration differ from the default.

Step 2: Back up your data

1. Create an XML backup of your existing data. SeeBack up a Site. Make a note of the location where you put the XML file. You will need it later to restore your Confluence data into your new database.
2. Stop 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. Make a separate backup using the utilities that were installed with your external 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.)

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

- 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:

- Select **Production Installation** as the installation type.
- Enter your license key, as usual.
- The next screens will depend on your license type:
  - If you are using a Confluence Data Center license, under **Choose your deployment type,** sele ct either non-clustered (single node) or clustered.
  - If you are using a Confluence Server license, select **My own database**then select your particular database from the**Database type** from the dropdown menu.
- Enter your database details. Use **test connection**to check your database is set up correctly.
- On the load content step, select **Empty Site**. You will need to restore from backup after you've completed the setup wizard.
- After you've completed the setup wizard, go to > **General Configuration** > **Backup and restore** and follow the steps on Restore a Site.

Step 7. Re-install your Marketplace apps

Re-install any apps (also known as plugins or add-ons) that are not bundled with Confluence.

- Use the same version of the app as on your old Confluence site.
- The data created by the app will already exist in your new Confluence site, because it is included in the XML backup.

Step 8. Check settings for new machine

If you are moving Confluence to a different machine, you need to check the following settings:

- Configure your new base URL. See Configuring the Server Base URL.
- Check your application links. See Linking to Another Application.
- Update any gadget subscriptions from external sites pointing to this Confluence site. For example, if your Jira site subscribes to Confluence gadgets, you will need to update your Jira site.
- Review any other resources that other systems are consuming from Confluence.

Method two for installations with a large volume of attachments

Before you start

These instructions only apply to attachments stored in the file system. If you store attachments in the database seeAttachment Storage Configurationto find out how to migrate between different attachment storage methods.

Step 1: Take note of your Marketplace apps

Take note of the apps (also knowns as plugins or add-ons) currently installed and enabled in Confluence, so that you can reinstate them later. Make a note of the following for each app:

- App name and vendor
- Version
- Enabled or disabled status. This is useful if you have enabled or disabled modules yourself, making your configuration differ from the default.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Step 2: Back up your data

1. Create an XML backup of your existing data. SeeBack up a Site. Make a note of the location where you put the XML file. You will need it later to restore your Confluence data into your new database.
2. Stop Confluence.
3. Make a copy of the attachments directory (`<CONFLUENCE-HOME-DIRECTORY>\attachments`) in your Confluence Home directory. You will need it later to copy your Confluence attachments data into your new Confluence installation.
4. Make a separate backup using the utilities that were installed with your external database. This is also a precautionary measure.

Step 3: Set up the new database

Choose the database setup instructions for your new database, and follow those instructions to do the following:

- Install the database server.
- Perform any required configuration of the database server, as instructed.
- Add the Confluence database and user. Make a note of the username and password that you define in this step. You will need them later, when running the Confluence Setup Wizard.

Step 4. Install Confluence (same version number) in a new location

Now you will install Confluence again, with a different home directory path and installation path.

Note: You must use the same version of Confluence as the existing installation. (If you want to upgrade Confluence, you must do it as a separate step.) For example, if your current site is running Confluence 5.1.2, your new installation must also be Confluence 5.1.2.

When running the Confluence installer:

- Choose **Custom Install**. (Do **not** choose to upgrade your existing installation.)
- Choose a **new destination directory**. This is the installation directory for your new Confluence. It must not be the same as the existing Confluence installation.
- Choose a **new home directory**. This is the data directory for your new Confluence. It must not be the same as the existing Confluence installation.

Step 5. Download and install the database driver if necessary

Note that Confluence bundles some database drivers, but you'll need to install the driver yourself if it is not bundled. Follow the database setup instructions for your new database, to download and install the database driver if necessary.

Step 6. Run the Confluence setup wizardand copy your data to your new database

When running the Confluence setup wizard:

- Select **Production Installation** as the installation type.
- Enter your license key, as usual.
- The next screens will depend on your license type:
  - If you are using a Confluence Data Center license, under **Choose your deployment type,** sele ct either non-clustered (single node) or clustered.
  - If you are using a Confluence Server license, select **My own database**then select your particular database from the**Database type** from the dropdown menu.
- Enter your database details. Use **test connection**to check your database is set up correctly.
- On the load content step, select **Empty Site**. You will need to restore from backup after you've completed the setup wizard.
- After you've completed the setup wizard, go to > **General Configuration** > **Backup and restore** and follow the steps on Restore a Site.

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.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Step 8. Re-install your Marketplace apps

Re-install any apps (also known as plugins or add-ons) that are not bundled with Confluence.

- Use the same version of the app as on your old Confluence site.
- The data created by the app will already exist in your new Confluence site, because it is included in the XML backup.

Step 9. Check settings for new machine

If you are moving Confluence to a different machine, you need to check the following settings:

- Configure your new base URL. See Configuring the Server Base URL.
- Check your application links. See Linking to Another Application.
- Update any gadget subscriptions from external sites pointing to this Confluence site. For example, if your Jira site subscribes to Confluence gadgets, you will need to update your Jira site.
- Review any other resources that other systems are consuming from Confluence.

A note about case sensitivity in your database

'Collation' refers to a set of rules that determine how data is sorted and compared. Case sensitivity is one aspect of collation. Other aspects include sensitivity to kana (Japanese script) and to width (single versus double byte characters).

Setting up a new Confluence instance

For new Confluence instances, we recommend using **case sensitive** collation for your Confluence database. This is the default collation type used by many database systems.

Note: Even if the database is configured for case sensitive collation, Confluence reduces all usernames to lower case characters before storing them in the database. For example, this means that 'joebloggs', 'joeBloggs' and 'JoeBloggs' will be treated as the same username.

Migrating an existing Confluence instance to a different database

The default Confluence configuration uses case sensitive database collation. This is typical of databases created under default conditions. If you are migrating from this type of configuration to a new database, we recommend that the new database uses case sensitive collation. If you use case insensitive collation, you may encounter data integrity problems after migration (for example, via an XML import) if data stored within your original Confluence site required case sensitive distinctions.

Migrating MySQL database between servers

Confluence 7.11 and higher versions introduced Database triggers and procedures. If you use `mysqldump` for migration, you need to add additional parameters to your `mysqldump` command. For more details see Confluence MySQL database migration causes content_procedure_for_denormalised_permissions does not exist error.

In addition, Confluence uses`DEFINER` clauses for its procedures that have hardcoded user information and hostname/IP of the server. When a database dump is generated that will export procedures having username /account in their DDL. Errors will happen if that dump is imported into a Database without the same username/account and privileges granted. For more information, see MySQL error 1449: The user specified as a definer does not exist.

Troubleshooting

See our troubleshooting guideif you're unable to restore your XML backup.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Configuring Database Character Encoding

Confluence and your database must be configured to use the same character encoding.

Confluence uses UTF-8 character encoding, so your database will also need to be configured to use UTF-8 (or the equivalent for your database, for example, AL32UTF8 for Oracle databases, or UTF8MB4 for MySQL databases).

On this page:

- New installations
- Existing installations

Related pages:

- Troubleshooting Character Encodings
- Configuring Character Encoding
- Database Troubleshooting for MySQL

New installations

When installing Confluence for the first time you will need to consider character encoding:

- when creating your database
- when connecting to the database via a JDBC connection string or datasource (if you use the simple setup method in the Confluence setup wizard, we'll take care of this for you).

The Confluence setup wizard will alert you if there is a problem with your character encoding, this will make sure you don't experience problems down the track. It is much easier to solve problems now, than later when you have Confluence data in your database.

The setup guide for each of our supported databases outlines how to configure character encoding correctly when creating your database:

- Database Setup for PostgreSQL
- Database Setup For MySQL
- Database Setup for SQL Server
- Database Setup for Oracle

Existing installations

For existing Confluence sites, where the first version of Confluence installed was 6.4 or earlier, we many not have checked the collation or character encoding of your database during the initial setup.

If your database is not correctly configured to use UTF-8 character encoding (or the equivalent for your database, for example AL32UTF8 for Oracle databases, or UTF8MB4 for MySQL databases):

- you may see a health check warning while using Confluence
- you may not be able to start Confluence after an upgrade.

If this happens, you'll need to change the character encoding for your existing database. The way you do this will depend on your database.

Also seeTroubleshooting Character Encodings for help diagnosing character encoding problems.

MySQL

SeeHow to Fix the Collation and Character Set of a MySQL Database manuallyfor details of what you'll need to do to fix the character encoding in your database. You should also make sure the collation is correct.

Microsoft SQL Server

SeeHow to fix the collation of a Microsoft SQL Server Confluence databasefor details of what you'll need to do to fix the character encoding in your database.

PostgreSQL

If you use PostgreSQL, the best option is to recreate your database.

SeeDatabase Setup for PostgreSQL for how to create your database using the correct character encoding, then follow the steps inMigrating to Another Database.

Oracle

If you use Oracle, the best option is to recreate your database.

SeeDatabase Setup for Oraclefor how to create your database using the correct character encoding, then follow the steps inMigrating to Another Database.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

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 `databaseSubsystemContext.xml` file changes from version to version.

Surviving Database Connection Closures

When a database server reboots or a network failure has occurred, all connections in the database connection pool are broken. To overcome this issue, Confluence would normally need to be restarted. To avoid this situation Confluence uses a validation query to check a database connection is alive before attempting to use it. If a broken connection is detected in the pool, a new one is created to replace it.

This validation query is**enabled by default on new installations**(from Confluence 6.5 and later), but if you've upgraded from an older Confluence version you can enable this manually by following the steps below.

While there are several different ways to perform this validation query, we recommend letting the database driver choose how to validate if a connection is still alive, rather than overriding the driver configuration with a specific validation query.

Enable validation query with a direct JDBC connection

To ensure Confluence validates database connections in the database connection pool:

1. Stop Confluence.
2. Edit the `<home-directory>confluence.cfg.xml`file.
3. Insert the following propertyin the `<properties>` block.

```
<property name="hibernate.c3p0.validate">true</property>
```

4. Save`confluence.cfg.xml`
5. Restart Confluence.

You should now be able to recover from a complete loss of all connections in the database connection pool without the need to restart Confluence.

Enable validation query with a datasource connection

⚠ We ended support for datasource connections in Confluence 8.0. If you are currently using a JNDI datasource connection, we recommend you use a direct JDBC connection to your database. This will also make upgrading to future versions of Confluence easier.

To ensure Confluence validates database connections in the database connection pool:

1. Stop Confluence.
2. Edit the`<installation-directory>/conf/server.xml`file (or wherever you have configured your datasource).
3. Find the Resource element for your data source, and add the "testOnBorrow" parameter as in the example for PostgreSQL below. Remember to give it the appropriate value for your database type.

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="60"
          maxIdle="20"
          testOnBorrow="true" />
...
```

4. Save`conf/server.xml`
5. Restart Confluence.

You should now be able to recover from a complete loss of all connections in the database connection pool without the need to restart Confluence.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Configuring a datasource connection

> We ended support for datasource connections in Confluence 8.0. If you are currently using a JNDI datasource connection, we recommend you switch to a direct JDBC connection. This will also make upgrading to future versions of Confluence easier.
>
> How to convert a datasource to a direct JDBC connection

On this page:

- New Confluence installation
- Existing Confluence installation
- Upgrading Confluence with a datasource
- Known issues

Related pages:

- Database JDBC Drivers

This guide covers how to configure a JNDI datasource connection to your database. With this type of connection, Confluence asks the application server (Tomcat) for your database connection information.

If you'd prefer to use a JDBC connection see the guide for your database:

- Database Setup for PostgreSQL
- Database Setup for SQL Server
- Database Setup For MySQL
- Database Setup for Oracle

Direct JDBC is the most common way to connect Confluence to your database and is the easiest method when it comes time to upgrade Confluence.

New Confluence installation

The Confluence setup wizard will only provide an option to use a datasource if it detects one in your Tomcat configuration. If you want to use a datasource, follow the steps below.

1. Stop Confluence

In the Confluence setup wizard, you'll be prompted to choose your database. At this point, you should:

1. Stop Confluence.
2. Back up the following files, in case you need to revert your changes:
   - `<installation-directory>/conf/server.xml`
   - `<installation-directory>/confluence/WEB-INF/web.xml`
   - `<home-directory>/confluence.cfg.xml`

2. Add your database driver

Copy your database driver into the `<installation-directory>/lib`directory.

Here's where to find the driver for your database:

- **PostgreSQL**: bundled with Confluence at`<installation-directory>/confluence/WEB-INF/lib/postgresql-x.x.x.jar`
- **Microsoft SQL Server**: bundled with Confluence at`<installation-directory>/confluence/WEB-INF/lib/mssql-jdbc-x.x.x.x.jar`
- **MySQL**: head to Database JDBC Drivers todownload the driver
- **Oracle**:head toDatabase JDBC Drivers todownload the driver

3. Configure the datasource in Tomcat

Next, add the datasource configuration to Tomcat.

1. Edit `<installation-directory>/conf/server.xml`
2. Find the following lines:

```
<Context path="" docBase="../confluence" debug="0" reloadable="true">
<!-- Logger is deprecated in Tomcat 5.5. Logging configuration for Confluence is
specified in confluence/WEB-INF/classes/log4j.properties -->
```

3. Insert the following DataSource `Resource` element for your specific database directly after the lines above (inside the `Context` element, directly after the opening `<Context .../>` line, before `Manager`).

```
<Resource name="jdbc/confluence" auth="Container" type="javax.sql.DataSource"
    username="<database-user>"
    password="<password>"
    driverClassName="org.postgresql.Driver"
    url="jdbc:postgresql://<host>:5432/<database-name>"
    maxTotal="60"
    maxIdle="20"
    testOnBorrow="true"/>
```

```
<Resource name="jdbc/confluence" auth="Container" type="javax.sql.DataSource"
    username="<database-user>"
    password="<password>"
   driverClassName="com.microsoft.sqlserver.jdbc.SQLServerDriver"
    url="jdbc:sqlserver://<host>:1433;database=<database-name>"
    maxTotal="60"
    maxIdle="20"
    testOnBorrow="true"/>
```

If you are using the 5.1.x driver (for MySQL 5.7):

```
<Resource name="jdbc/confluence" auth="Container" type="javax.sql.DataSource"
    username="<database-user>"
    password="<password>"
    driverClassName="com.mysql.jdbc.Driver"
    url="jdbc:mysql://<host>:3306/<database-name>?useUnicode=true&amp;characterEncoding=utf8"
    maxTotal="60"
    maxIdle="20"
    defaultTransactionIsolation="READ_COMMITTED"
    testOnBorrow="true"/>
```

If you're using the 8.0.x driver (for MySQL 8):

```
<Resource name="jdbc/confluence" auth="Container" type="javax.sql.DataSource"
    username="<database-user>"
    password="<password>"
    driverClassName="com.mysql.cj.jdbc.Driver"
    url="jdbc:mysql://<host>:3306/<database-name>?useUnicode=true&amp;characterEncoding=utf8"
    maxTotal="60"
    maxIdle="20"
    defaultTransactionIsolation="READ_COMMITTED"
    testOnBorrow="true"/>
```

```
<Resource name="jdbc/confluence" auth="Container" type="javax.sql.DataSource"
    driverClassName="oracle.jdbc.OracleDriver"
    url="jdbc:oracle:thin:@<host>:1521:<SID>"
    username="<database-user>"
    password="<password>"
    connectionProperties="SetBigStringTryClob=true"
        accessToUnderlyingConnectionAllowed="true"
    maxTotal="60"
    maxIdle="20"
    maxWaitMillis="10000"
    testOnBorrow="true"/>
```


Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Seehow to find your Oracle URL.

Replace `<database-user>`, `<password>`, `<host>`and `<database-name>`(or `<SID>`for Oracle) with details of your own database. You may also need to change the port, if your database server is not running on the default port.

4. Configure the connection pool and other properties.See theApacheTomcat 9 Datasource documentationfor more information.

   Here are the configuration properties for Tomcat's standard data source resource factory (`org.apache.tomcat.dbcp.dbcp.BasicDataSourceFactory`):

   - `driverClassName` - Fully qualified Java class name of the JDBC driver to be used.
   - `maxTotal` - The maximum number of active instances that can be allocated from this pool at the same time.
   - `maxIdle` - The maximum number of connections that can sit idle in this pool at the same time.
   - `maxWaitMillis` - The maximum number of milliseconds that the pool will wait (when there are no available connections) for a connection to be returned before throwing an exception.
   - `password` - Database password to be passed to the JDBC driver.
   - `url` - Connection URL to be passed to the JDBC driver. (For backwards compatibility, the property driverName is also recognized.)
   - `user` - Database username to be passed to the JDBC driver.
   - `validationQuery` - We don't recommend you set a validation query explicitly. Instead, we recommend you set `testOnBorrow`, which will use the validation query defined by your database driver.See Surviving Database Connection Closures for more information.

5. If you plan to use collaborative editing, you'll need to make sure:
   - You're using a supported database driver. Collaborative editing will fail if you're using an unsupported or custom JDBC driver or`driverClassName`in your datasource. SeeDatabase JDBC Driversfor the list of drivers we support.
   - Your database connection pool allows enough connections to support both Confluence and Synchrony (which defaults to a maximum pool size of 15)
   - You're using simple username and password authentication for your database.

4. Configure the Confluence web application

Configure Confluence to use this datasource:

1. Edit`<CONFLUENCE_INSTALLATION>/confluence/WEB-INF/web.xml`.
2. Insert the following element just before`</web-app>`near the end of the file:

```
<resource-ref>
  <description>Connection Pool</description>
  <res-ref-name>jdbc/confluence</res-ref-name>
  <res-type>javax.sql.DataSource</res-type>
  <res-auth>Container</res-auth>
</resource-ref>
```

5. Restart Confluence and continue setup process

Now that your datasource is configured, you can continue with the setup wizard.

1. Start Confluence.
2. Go to`http://localhost:8090` to return to the setup wizard.
3. When prompted choose **My own database(datasource)**.
4. Enterthe JNDI name of your datasource, for example,`java:comp/env/jdbc/confluence`
5. Follow the prompts to finish setting up Confluence.

6. Update your datasource to turn off auto commit

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Once you've confirmed that Confluence is up and running, you'll need to make a final change to your datasource to avoid a known issuewitheditting pages. See ~~CONFSERVER-59524~~ CLOSED

1. Stop Confluence.
2. Edit `<installation-directory>/conf/server.xml`
3. Add the following parameter in your datasource Resource element.

```
defaultAutoCommit="false"
```

4. Start Confluence.
5. Repeat this for all cluster nodes.

Existing Confluence installation

If you want to switch from using a direct JDBC connection to a datasource:

- Stop Confluence.
- Back up the following files, in case you need to revert your changes:
  - `<installation-directory>/conf/server.xml`
  - `<installation-directory>/confluence/WEB-INF/web.xml`
  - `<home-directory>/confluence.cfg.xml`
- Copy your database driver into the `<installation-directory>/lib`directory, as described in the steps above. You can find the details of your current database connection in `<home-directory>/confluence.cfg.xml`.
- Edit `<installation-directory>/conf/server.xml` and insert the following `DataSource` Resource element for your specific database (inside the `Context` element, directly after the opening `<Context.../>` line before `Manager`)

```
<Resource name="jdbc/confluence" auth="Container" type="javax.sql.DataSource"
    username="<database-user>"
    password="<password>"
    driverClassName="org.postgresql.Driver"
    url="jdbc:postgresql://<host>:5432/<database-name>"
    maxTotal="60"
    maxIdle="20"
    testOnBorrow="true"
    defaultAutoCommit="false"/>
```

```
<Resource name="jdbc/confluence" auth="Container" type="javax.sql.DataSource"
    username="<database-user>"
    password="<password>"
   driverClassName="com.microsoft.sqlserver.jdbc.SQLServerDriver"
    url="jdbc:sqlserver://<host>:1433;database=<database-name>"
    maxTotal="60"
    maxIdle="20"
    testOnBorrow="true"
    defaultAutoCommit="false"/>
```

If you are using the 5.1.x driver (for MySQL 5.7):

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

```
<Resource name="jdbc/confluence" auth="Container" type="javax.sql.DataSource"
    username="<database-user>"
    password="<password>"
    driverClassName="com.mysql.jdbc.Driver"
    url="jdbc:mysql://<host>:3306/<database-name>?useUnicode=true&amp;characterEncoding=utf8"
    maxTotal="60"
    maxIdle="20"
    defaultTransactionIsolation="READ_COMMITTED"
    testOnBorrow="true"
    defaultAutoCommit="false"/>
```

If you're using the 8.0.x driver (for MySQL 8):

```
<Resource name="jdbc/confluence" auth="Container" type="javax.sql.DataSource"
    username="<database-user>"
    password="<password>"
    driverClassName="com.mysql.cj.jdbc.Driver"
    url="jdbc:mysql://<host>:3306/<database-name>?useUnicode=true&amp;characterEncoding=utf8"
    maxTotal="60"
    maxIdle="20"
    defaultTransactionIsolation="READ_COMMITTED"
    testOnBorrow="true"
    defaultAutoCommit="false"/>
```

```
<Resource name="jdbc/confluence" auth="Container" type="javax.sql.DataSource"
    driverClassName="oracle.jdbc.OracleDriver"
    url="jdbc:oracle:thin:@<host>:1521:<SID>"
    username="<database-user>"
    password="<password>"
    connectionProperties="SetBigStringTryClob=true"
        accessToUnderlyingConnectionAllowed="true"
    maxTotal="60"
    maxIdle="20"
    maxWaitMillis="10000"
    testOnBorrow="true"
    defaultAutoCommit="false"/>
```

See how to find your Oracle URL.

Replace `<database-user>`, `<password>`, `<host>` and `<database-name>` (or `<SID>` for Oracle) with details of your own database. You may also need to change the port, if your database server is not running on the default port.

- Edit the `<home-directory>/confluence.cfg.xml` file and remove any line that contains a property that begins with `hibernate`.
- Insert the following at the start of the `<properties>` section.

```
<property name="hibernate.setup"><![CDATA[true]]></property>
<property name="hibernate.dialect"><![CDATA[net.sf.hibernate.dialect.PostgreSQLDialect]]><
/property>
<property name="hibernate.connection.datasource"><![CDATA[java:comp/env/jdbc/confluence]]><
/property>
```

```
<property name="hibernate.setup"><![CDATA[true]]></property>
<property name="hibernate.dialect"><![CDATA[net.sf.hibernate.dialect.SQLServerIntlDialect]]><
/property>
<property name="hibernate.connection.datasource"><![CDATA[java:comp/env/jdbc/confluence]]><
/property>
```


Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

```
<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>
```

```
<property name="hibernate.setup"><![CDATA[true]]></property>
<property name="hibernate.dialect"><![CDATA[com.atlassian.confluence.impl.hibernate.dialect.
OracleDialect]]></property>
<property name="hibernate.connection.datasource"><![CDATA[java:comp/env/jdbc/confluence]]><
/property>
```

- Start Confluence.

Upgrading Confluence with a datasource

If you're upgrading Confluence (manually or using the installer) you will need to:

- Stop Confluence (if you have attempted to start it).
- Copy your database driver into the `<installation-directory>/lib`directory.
- Edit `<installation-directory>/conf/server.xml`and add your datasource resource.
- Edit `<installation-directory>/confluence/WEB-INF/web.xml`to configure Confluence to use this datasource.

If you forget to do these steps, Confluence will not start up after upgrade and you'll see the following error:

```
HTTP Status 500 - Confluence is vacant, a call to tenanted [public abstract org.hibernate.Session org.
hibernate.SessionFactory.getCurrentSession() throws org.hibernate.HibernateException] is not allowed.
```

Known issues

- If you experience a lot of`ContentUpdatedEvent` errors in the logs, you may need to add adddefau ltAutoCommit="false" to the datasource in the `server.xml` file. See
  ~~CONFSERVER-59524~~ CLOSED for more information and full details of the workaround.
- There's a known issue whereSynchrony does not start if Confluence connects to the database using a datasource. See
  ~~CONFSERVER-60120~~ - Synchrony not starting with datasource after upgrade to Confluence 7.5.2 , 7.6.0 , 7.6.1 & 7.6.2 CLOSED
  for more information and a workaround.
- There's a known issue when running Oracle with Native Network Encryption that can cause Confluence to become unresponsive. See ~~CONFSERVER-60152~~ CLOSED for more detailsandC onfluence Unresponsive Due to High Database Connection Latencyfor some suggested mitigation strategies.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Configuring Confluence Data Center to work with Amazon Aurora

On this page:

- Deploying Confluence Data Center with Amazon Aurora
- Connecting an existing Quick Start deployment to Amazon Aurora
- Manually setting up an Amazon Aurora Database
  - AWS documentation
- Connecting Confluence Data Center to a new Amazon Aurora database
  - Step 1: Shut down Confluence Data Center
  - Step 2: Update the database URL each Confluence node uses
  - Step 3: Configure collaborative editing
  - Step 4: Restart Confluence

Confluence Data Center supports the use of a single-writer, PostgreSQL-compatible Amazon Aurora clustered database. A typical production-grade cluster includes one or more readers in a different availability zone. If the writer fails, Amazon Aurora will automatically promote one of the readers to take its place. For more information, see Amazon Aurora Features: PostgreSQL-Compatible Edition.

Deploying Confluence Data Center with Amazon Aurora

To create a new Confluence Data Center deployment with Amazon Aurora,we recommend that you use the AWS Quick Start for Confluence .This Quick Start lets you configure a PostgreSQL-compatible Amazon Aurora cluster with one writer and two readers in separate availability zones. See Running Confluence Data Center in AWS for more information.

Connecting an existing Quick Start deployment to Amazon Aurora

If you deployed Confluence Data Center using the Quick Start before 11 June 2019, you wont be able to connect it to a new Amazon Aurora cluster. The Quick Start version prior to that date applied some settings that are incompatible with Aurora.

Instead, youll have to migrate your existing data to a new Confluence Data Center deployment:

1. Use the latest AWS Quick Start for Confluence to create a new Confluence Data Center deployment.
2. Shut down Confluence on the application nodes of both old and new deployments. If you use a standalone Synchrony cluster, shut down all the nodes in that cluster too.
3. Migrate your data from the old deployment to the new one:
   - **EFS**: EFS-to-EFS Backup explains how you can use an easy-to-deploy backup solution to perform a backup of your old EFS and restore it in the new deployment.
   - **Database**: Migrating Data to Amazon Aurora PostgreSQL contains instructions for migrating from Amazon RDS to a PostgreSQL-compatible Amazon Aurora cluster.

Once you finish the migration, re-start Confluence on all application nodes in the new deployment. If you use a standalone Synchrony cluster, re-start all its nodes.

We strongly recommend yourebuild your content index after performing a migration, to ensure Confluence search works as expected.

Manually setting up an Amazon Aurora Database

Confluence Data Center specifically supports the use of an Amazon Aurora cluster with the following configuration:

- It must have only one writer, replicating to one or more readers.
- Your PostgreSQL engine must be version 9.6 or higher.

CheckSupported Platformsfor more details.

AWS documentation

AWS has some helpful guides for setting up an Aurora database and migrating to it:

- Modular Architecture for Amazon Aurora PostgreSQL: a Quick Start that guides you through the deployment of a PostgreSQL-compatible Aurora Database cluster. This cluster has one writer and two readers, preferably in different availability zones.
- Upgrading the PostgreSQL DB Engine for Amazon RDS: shows you how upgrade your database engine to a supported version before migrating it to Amazon Aurora.
- Migrating Data to Amazon Aurora PostgreSQL: contains instructions for migrating from Amazon RDS to a PostgreSQL-compatibleAmazon Aurora cluster.
- Best Practices with Amazon Aurora PostgreSQL: contains additional information about best practices and options for migrating data to a PostgreSQL-compatible Amazon Aurora cluster.

Amazon also offers anAWS Database Migration Serviceto facilitate a managed migration. This service offers minimal downtime, and supports migrations to Aurora from a wide variety of source databases.

If you deployed Confluence Data Center through our AWS Quick Start before 11 June 2019, you cant connect it to a new Amazon Aurora cluster. Rather, youll need to re-deploy Confluence Data Center using our updated Quick Start and migrate your data across. See **Connecting an existing Quick Start deployment to Amazon Aurora** for more information.

Connecting Confluence Data Center to a new Amazon Aurora database

After deploying an Aurora cluster and migrating your database to it, youll need to properly connect it to Confluence. This will involve updating the database URL used by Confluence Data Center.

Confluence Data Center should point to the the Aurora cluster writer endpoint URL, and include the`targetServerType`parameter. This parameter allows Confluence to target the writer database instance, which ensures the application can reconnect to it after a failover.

Your database URL will look something like this:

```
jdbc:postgresql://<CLUSTER_WRITER_ENDPOINT>:<CLUSTER_WRITER_PORT>/<DATABASE_NAME>?targetServerType=master
```

If you deployed your Aurora cluster through the Modular Architecture for Amazon Aurora PostgreSQL Quick Start, you can find then the cluster writer details from the **Outputs** tab in AWS. The `RDSEndpointAddress` and `RDSEndpointAddressPort` values will be your `CLUSTER_WRITER_ENDPOINT` and `CLUSTER_WRITER_PORT` , respectively.

The following steps will walk you through the process of connecting Confluence and Aurora.

Step 1: Shut down Confluence Data Center

To safely reconfigure Confluence Data Centers database connection, we recommend a full outage. To do this, stop Confluence on all application nodes.

If you have a standalone Synchrony cluster, stop Synchrony on each node there.

Step 2: Update the database URL each Confluence node uses

How you perform this step depends on how Confluence currently connects to your database.

If you use a direct JDBC connection

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

1. On the first node, edit the `<local-home>/confluence-cfg.xml` file.
2. Update the `hibernate.connection.url` property with your cluster writer endpoint URL as follows:

```
<property name="hibernate.connection.url">jdbc:postgresql://<CLUSTER_WRITER_ENDPOINT>:
<CLUSTER_WRITER_PORT>/<DATABASE_NAME>?targetServerType=master
```

3. Repeat this change on all other nodes.
4. Start Confluence, one node at a time.

⚠ This change must be made in the **local home** directory on each node, not in the copy of the `confluence-cfg.xml` that can be found in the shared home.

If you use a datasource connection

1. Stop Confluence on all nodes.
2. On the first node, edit the `<install-directory>/conf/server.xml` file.
3. Update the url parameter in the datasource `Resource` element with your cluster writer endpoint URL as follows:

```
<Resource name="jdbc/confluence" auth="Container" type="javax.sql.DataSource"
username="<database-user>"
password="<password>"
driverClassName="org.postgresql.Driver"
url="jdbc:postgresql://<CLUSTER_WRITER_ENDPOINT>:<CLUSTER_WRITER_PORT>/<DATABASE_NAME>?
targetServerType=master"
maxTotal="60"
maxIdle="20"
validationQuery="select 1"/>
```

4. Repeat this change on all other nodes.
5. Start Confluence, one node at a time.

Step 3: Configure collaborative editing

Synchrony, the engine that powers collaborative editing, can be deployed in two different ways, which affects how you pass the database URL:

1. **Managed by Confluence** - Confluence will automatically launch a Synchrony process on the same node, and manage it for you.
2. **Standalone Synchrony cluster** - You deploy and manage Synchrony standalone in its own cluster with as many nodes as you need. This is the default method when you deploy Confluence in AWS using our Quick Start.

If Synchrony is managed by Confluence, you don't need to do anything. Confluence will pass the URL to Synchrony for you.

If you run a Standalone Synchrony cluster, you will need to provide the cluster writer endpoint URL in your startup script. This script will be either `<synchrony-home>/start-synchrony.sh` or `start-synchrony.bat`, depending on your operating system. Edit your script as follows:

start-synchrony.sh (Linux)

```
DATABASE_URL="jdbc:postgresql://<CLUSTER_WRITER_ENDPOINT>:<CLUSTER_WRITER_PORT>/<DATABASE_NAME>?
targetServerType=master"
```

start-synchrony.bat (Windows)

```
set DATABASE_URL=jdbc:postgresql://<CLUSTER_WRITER_ENDPOINT>:<CLUSTER_WRITER_PORT>/<DATABASE_NAME>?
targetServerType=master
```


Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

SeeSet up a Synchrony cluster for Confluence Data Centerfor more information about setting up Synchrony standalone cluster.

⚠ If you run Synchrony as a Linux service, you'll need to reinstall the service.

Step 4: Restart Confluence

After making the necessary database URL updates, you can now restart Confluence on each application node, one node at a time.

If you have a standalone Synchrony cluster, restart it on each of the clusters nodes.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Encrypt database password

To add extra security to yourConfluence Data Center and Server instance, you can encrypt the database password that is stored in the `confluence.cfg.xml` file. You can choose to use a Base64 cipher, Algorithm cipher, or create your own cipher.

YourConfluence instance needs a plain-text password to connect to the database. The configuration file will contain all the information needed to decrypt the password, so it's necessary to protect access toConfluence and database servers.

You can only encrypt your password if you're using a JDBC connection (not a JNDI datasource).

On this page:

- Basic encryption
- Advanced encryption
- Custom encryption

Basic encryption

This method uses a Base64 cipher, which is simple obfuscation. It is a straightforward solution for users who dont want to store database passwords in plain-text.

Basic encryption

Advanced encryption

This method allows you to choose an algorithm to encrypt a database password. It provides more security as you don't have to store the encrypted password anywhere in the configuration file, which makes it difficult for unauthorised parties to find and decrypt it.

Advanced encryption

Custom encryption

If you have special requirements for database password encryption, you can create your own encryption mechanism based on our examples.

Custom encryption

Basic encryption

To add extra security to your Confluence site, you can encrypt the database password that is stored in the confluence.cfg.xml file.

For this method, we'll use Base64Cipher, which is a way to achieve simple obfuscation of the database password.

On this page:

- Step 1. Encrypt the password
- Step 2. Add the encrypted password to the properties file
- Decrypt the password

Step 1. Encrypt the password

To encrypt the database password:

1. Go to `<install-directory>/bin`.
2. Run the following command to encrypt your password.

```
java -cp "./*" com.atlassian.db.config.password.tools.CipherTool
```

   - **-- silent -s**: limits logging to a minimum
   - **-- help -h**: prints a help message with all parameters
   - **-- mode -m**: defines what to do with the password, either**encrypt**or**decrypt**. If omitted, 'encrypt' will be used.
   - **-- password -p**: plain text password. If omitted, you'll be asked to enter it. We recommend that you omit this parameter so that your password is not stored in the history.

Step 2. Add the encrypted password to the properties file

To add the encrypted password:

1. Back up the `<home-directory>/confluence.cfg.xml` file. Move the backup to a safe place outside of your instance.
2. In the `confluence.cfg.xml` file, replace the `hibernate.connection.password` property with the password generated by the CLI. Also, specify the `jdbc.password.decrypter.classname` property. For example:

```
<jdbc.password.decrypter.classname>com.atlassian.db.config.password.ciphers.base64.Base64Cipher<
/jdbc.password.decrypter.classname>
<hibernate.connection.password>your-obfuscated-password</hibernate.connection.password>
```

3. Restart Confluence.

Decrypt the password

To decrypt the password, extend the command with the `-m decrypt` parameter:

```
java -cp "./*" com.atlassian.db.config.password.tools.CipherTool -m decrypt
```

When asked for a password, provide the encrypted one from your `confluence.cfg.xml` file.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Advanced encryption

To add extra security to your Confluence site, you can encrypt the database password that is stored in the confluence.cfg.xml file.

In this method, we'll use `AlgorithmCipher` that allows you to choose the algorithm used to encrypt your password.

On this page:

- Before you begin
- Step 1. Encrypt the password
- Step 2. Secure the generated files
- Step 3. Add the encrypted password to confluence.cfg.xml
- Decrypt the password
- Troubleshooting

Before you begin

Prepare the JSON object

Youll need to provide all arguments required to encrypt your password in a JSON object. Prepare beforehand by using the information and examples below.

Field	Description
plainTextPassword	Password in plain text.
algorithm	You can choose one of the following algorithms: • AES/CBC/PKCS5Padding • DES/CBC/PKCS5Padding • DESede/CBC/PKCS5Padding
algorithmKey	The algorithm key must correspond with the algorithm chosen above: • AES • DES • DESede

```
{"plainTextPassword":"yourPassword","algorithm":"AES/CBC/PKCS5PADDING","algorithmKey":"AES"}
```

Note that the JSON object must be entered as a single line.

Step 1. Encrypt the password

To encrypt the database password:

1. Go to `<install-directory>/bin`.
2. Run the following command to encrypt your password:

```
java -cp "./*" com.atlassian.db.config.password.tools.CipherTool -c com.atlassian.db.config.
password.ciphers.algorithm.AlgorithmCipher
```

- **-- silent -s**: limits logging to a minimum
- **-- help -h**: prints a help message with all parameters

- **-- mode -m**: defines what to do with the password, either**encrypt**or**decrypt**. If omitted, 'encrypt' will be used.
- **-- password -p**: JSON object with required arguments. If omitted, you'll be asked to enter a password. We recommend that you omit this parameter so that your password is not stored in the terminal history.

3. After running the command, you'll be asked to provide the required arguments in a JSON object in a single line.

When encrypting your password, the encryption tool generates three files and prints the output JSON object that you'll later add to the `confluence.cfg.xml` file. The next step discusses how to secure those files.

Step 2. Secure the generated files

Move the files generated by the tool to a secure place. Change them to read-only and accessible only to the user running Confluence. If you're running Confluence in a cluster, the files should be available to all nodes via the same path. Confluence needs to be able to access and read those files to decrypt your password and connect to the database.

The following files are generated:

- `javax.crypto.SealedObject_[timestamp]`
  File with the encrypted password.
- `javax.crypto.spec.SecretKeySpec_[timestamp]`
  Key used to encrypt your password. You will need this file to decrypt your password.
- `java.security.AlgorithmParameters_[timestamp]`
  Algorithm parameters used to encrypt your password. You will only need this file if you want to recreate the password.

Store paths as environment variables

This step is optional.

You can choose to store paths to the generated files as environment variables. If the paths aren't present in the `confluence.cfg.xml` file and the `hibernate.connection.password` is an empty JSON object (`{}`), `AlgorithmCipher` will look for them in the environment. This way, file paths are not stored in the file, making it difficult to locate the files used for encryption.

To store the paths to the generated files as environment variables:

1. Store the two generated files as environment variables. You don't need to add the file with algorithm parameters, because `AlgorithmCipher` does not use it to decrypt the password. You must set the following environment variables to the correct values in any of the scripts used for launching your Confluence instance.

```
com_atlassian_db_config_password_ciphers_algorithm_javax_crypto_spec_SecretKeySpec
com_atlassian_db_config_password_ciphers_algorithm_javax_crypto_SealedObject
```

2. Edit the output from the first step, Encrypt the password, and remove paths to the files. The decrypter class name and password should match the following:

```
<jdbc.password.decrypter.classname>=com.atlassian.db.config.password.ciphers.algorithm.
AlgorithmCipher
<hibernate.connection.password>={}
```

Step 3. Add the encrypted password to confluence.cfg.xml

To add the encrypted password:

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

1. Back up the `<local-home>/confluence.cfg.xml` file. Move the backup to a safe place, ideally outside your instance.
2. In the `confluence.cfg.xml` file, replace the hibernate.connection.password property with the output JSON object. Depending on whether youre using environment variables or not, adjust the JSON object to one of the following examples:

   If youre storing file paths as environment variables, remove the paths from the output. The properties should look like the following:

   When using environment variables

   ```
   <property name="jdbc.password.decrypter.classname">com.atlassian.db.config.password.ciphers.
   algorithm.AlgorithmCipher</property>
   <property name="hibernate.connection.password">{}</property>
   ```

   If youre not using environment variables and want to stick to file paths in the `confluence.cfg.xml` file, make sure you update their paths after moving them to a secure place. The properties should look like the following:

   When not using environment variables

   ```
   <property name="jdbc.password.decrypter.classname">com.atlassian.db.config.password.ciphers.
   algorithm.AlgorithmCipher</property>
   <property name="hibernate.connection.password">{"sealedObjectFilePath":"/home/confluence/javax.
   crypto.SealedObject_123456789","keyFilePath":"/home/confluence/javax.crypto.spec.
   SecretKeySpec_123456789"}</property>
   ```

   If you're running Confluence on Windows, you should avoid backslashes in the path to prevent JSON parsing errors. The paths should look like the following example:

   If using Windows

   ```
   <property name="jdbc.password.decrypter.classname">com.atlassian.db.config.password.ciphers.
   algorithm.AlgorithmCipher</property>
   <property name="hibernate.connection.password">{"sealedObjectFilePath":"C:/confluence/javax.
   crypto.SealedObject_123456789","keyFilePath":"C:/confluence/javax.crypto.spec.
   SecretKeySpec_123456789"}</property>
   ```

3. Restart Confluence.

Decrypt the password

To decrypt an encrypted password, extend the command used earlier with the `-m decrypt` parameter:

```
java -cp "./*" com.atlassian.db.config.password.tools.CipherTool -c com.atlassian.db.config.password.
ciphers.algorithm.AlgorithmCipher -m decrypt
```

When asked for the password, provide the JSON object from your `confluence.cfg.xml` file.

Sample JSON object when using environment variables

```
{}
```

Sample JSON object when using files

```
{"sealedObjectFilePath":"/home/confluence/javax.crypto.SealedObject_123456789","keyFilePath":"/home
/confluence/javax.crypto.spec.SecretKeySpec_123456789"}
```


Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Recreate an encrypted password

If you lose an encrypted password and try to encrypt the plain text password once again, the new encrypted password will look different. This is not an issue, as it will still represent the same plain text password. However, in some cases, you might want to keep it consistent, for example by having the same encrypted password when a Confluence instance is migrated to another server.

To encrypt the password in the exact same way as you did before, you will need the key used to encrypt the original password and the algorithm parameters. Both of these were generated by the encryption tool and saved in the following files:

- **Key** - `javax.crypto.spec.SecretKeySpec_[timestamp]`
- **Algorithm parameters**-`java.security.AlgorithmParameters_[timestamp]`

Once you've located these files, you can point the encryption tool to their location by using two extra fields in the JSON object.

Field	Description
keyFilePath	Path to a file that contains the key used to encrypt your original password, e.g. `javax.crypto.spec.SecretKeySpec_[timestamp]`. If you stored the file path as environment variable, you can omit this parameter.
algorithmParametersFilePath	Path to a file that contains the algorithm parameters used to encrypt your original password, e.g. `java.security.AlgorithmParameters_[timestamp]`.

```
{"plainTextPassword":"yourPassword", "algorithm":"AES/CBC/PKCS5PADDING","algorithmKey":"AES","
algorithmParametersFilePath":"java.security.AlgorithmParameters_123456789","keyFilePath":"javax.
crypto.spec.SecretKeySpec_123456789"}
```

Note that the JSON object must be entered as a single line.

To encrypt the password, follow the steps in the first step, Encrypt the password and use the JSON object with the key and algorithm parameters.

Troubleshooting

To revert the changes, remove the `<atlassian-password-cipher-provider>` tag from the `confluence.cfg.xml` file, and change the encrypted password to a plain text one.

The setup screen means that Confluence couldn't connect to the database to access your configuration, most probably because of an error with decrypting your password.

To solve this problem, open `<home-directory>/logs/atlassian-confluence.log`, and check the lines after "Reading database configuration from".

Youll probably see the following message:

```
[c.a.c.config.database.DatabaseConfigHandler] Trying to get encrypted password from xml and decrypt
it
[c.a.d.c.p.ciphers.algorithm.AlgorithmCipher] Runtime Exception thrown when decrypting:
```

If thats the case, read the message, as it contains details about the error and a possible solution.

If the error is related to missing files, there might be a problem with your environment variables. They may have been deleted, or are no longer available if you changed the environment from staging to production. To verify, try adding file paths to the JSON object in the `confluence.cfg.xml` file.

If youre seeing "Bouncy Castle" errors, you will need encrypt the password again.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

To investigate this problem, go to `<home-directory>/logs/atlassian-confluence.log`, and check the lines after: "Reading database configuration from".

Youll probably see the following messages:

```
[c.a.c.config.database.DatabaseConfigHandler] Trying to get encrypted password from xml and decrypt
it
[c.a.c.config.database.DatabaseConfigHandler] Database password decryption success!
[c.a.config.bootstrap.DefaultAtlassianBootstrapManager] Could not successfully test your database:

[c.a.c.health.HealthChecks] Confluence couldn't connect to your database
[c.a.c.health.HealthChecks] Confluence failed to establish a connection to your database.
```

This means that Confluence decrypted the password successfully, but the password itself is incorrect.

To verify this:

1. Open the `confluence.cfg.xml` file, and copy the encrypted password.
2. Decrypt the password.
3. Check if the decrypted password is the same as the one in your backup `confluence.cfg.xml` file.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Custom encryption

To add extra security to your Confluence site, you can encrypt the database password that is stored in the `confluence.cfg.xml` file.

If you don't want to use the basic or advanced encryption methods provided by Confluence you can choose to create your own cipher. This may be especially useful if:

- you're required to use a specific vault to store the password
- you want to use a different encryption algorithm.

This procedure assumes you are familiar with Java and Maven.

On this page:

- Step 1. Create a Maven project and get API dependencies
- Step 2. Implement the Cipher interface
- Step 3. Test your implementation

Step 1. Create a Maven project and get API dependencies

To create a maven project and get API dependencies:

1. Get password-cipher-api and password-cipher-base dependencies.
    a. Go to `<install-directory>/confluence/WEB-INF/lib`and copy the following jar files:
    b. `password-cipher-api-<version>.jar`
    This file contains the API
    c. (`password-cipher-base-<version>.jar`
    (optional) This file contains some sample implementations
2. Create a Maven project.
3. Go to `resources` and create a new folder, named `libs`.
4. Copy the jar files to the libs folder.
5. Next, use the following pom:

```
<?xml version="1.0" encoding="UTF-8"?>
<project xmlns="http://maven.apache.org/POM/4.0.0"
xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
xsi:schemaLocation="http://maven.apache.org/POM/4.0.0http://maven.apache.org/xsd/maven-4.0.0.xsd">
<modelVersion>4.0.0</modelVersion>

  <groupId><your_group_ID></groupId>
  <artifactId><your_artifact_ID></artifactId>
  <version><your_version></version>

  <properties>
    <maven.compiler.source>1.8</maven.compiler.source>
    <maven.compiler.target>1.8</maven.compiler.target>
  </properties>

  <repositories>
    <repository>
      <id>local-maven-repo</id>
      <url>file:///${project.basedir}/libs</url>
    </repository>
  </repositories>

  <build>
    <resources>
      <resource>
        <directory>src/main/resources/libs</directory>
        <excludes>
          <exclude>*</exclude>
        </excludes>
        <filtering>false</filtering>
      </resource>
    </resources>
  </build>

  <dependencies>
    <dependency>
      <groupId>com.atlassian.db.config</groupId>
      <artifactId>password-cipher-api</artifactId>
      <version><api_version></version>
      <scope>provided</scope>
    </dependency>
    <dependency>
      <groupId>com.atlassian.db.config</groupId>
      <artifactId>password-cipher-base</artifactId>
      <version><base_version></version>
      <scope>provided</scope>
    </dependency>
  </dependencies>
</project>
```

Step 2. Implement the Cipher interface

The Cipher interface contains two methods that you need to implement according to your requirements; `encrypt` and `decrypt`. `decrypt` is called during Confluence startup, which means that long-running tasks can affect the startup time. encrypt is not called by Confluence, as it's only used in the encryption tool.

You can use the `Base64Cipher` and `AlgorithmCipher` as examples.

Step 3. Test your implementation

The encryption tool described in Basic encryption and Advanced encryption uses the same code as Confluence to decrypt the password. You can use it to test your implementation.

Assuming that the CLI and your jar is in the same folder:

```
java -cp "./*" com.atlassian.db.config.password.tools.CipherTool -c your.package.here.ClassName
```


Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Step 4. Make your library available

Confluence must be able to access your library. Your class will be instantiated using reflection.

Put the library in the`<install-directory>/confluence/WEB-INF/lib` directory.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Backup and Restore

When setting up your Confluence site, it's important to consider how you will back up your data, and restore it, if things go wrong.

Recommended backup strategy

Having a robust backup strategy for your Confluence site is essential. You should back up your database, installation directory, and home directories (including attachments) on a regular basis using the database administration or backup tool of your choice.

See Production Backup Strategy

Manually XML backup

You can export your entire site or selected spaces at any time. The backup will be generated as a zipped XML file.

See Backup a Site

See Backup a Space or multiple Spaces

Scheduled XML backup

Since Confluence 8.3, we have changed the way we do backup and restore. Learn more about these changes in the Confluence 8.3 Release Notes.

Scheduled backups don't use this new approach; it still uses our legacy system. We recommend you stop using scheduled backups unless required because it contains many of the issues resolved in the new system.

Confluence provides a scheduled XML backup option, which backs up your site by performing a full site XML export each day. This method can be useful for small sites, test sites, or in addition to your database and directory backups. We don't recommend you rely solely on this backup method for your production site.

There are a number of reasons why XML site backups are unsuitable for large Confluence sites:

- As the number of pages in your site increases, the XML backup takes progressively longer to complete, and in extreme cases the process of generating the export can cause an outage.
- XML backups can consume a lot of disk space rapidly. For example, a 1GB Confluence site will create 30GB worth of backups in a month, if unattended.
- If the XML export file is very large, restoring your site can take a long time, or may time out.
- Marketplace and other user-installed apps are not included in the XML backup. After importing your backup into a new Confluence site, you will need to re-install all user installed apps.

On Confluence Data Center, the scheduled XML backup is disabled by default.

To learn how to enable this job or change its frequency, see Scheduling a Backup

Restoring your site from a backup

In the event you need to restore your site from a backup, the way you do this depends on your backup method.

- See Restore a Site to find out how to restore data from an XML backup into an existing Confluence site.
- SeeRestoring data from other backupsfor tips on how to restore Confluence from a database backup.

Version compatibility

Confluence accepts XML backups from Confluence 6.0.5 and newer, but it is only possible to restore a backup to **the same or a newer version**.

For example, a backup created in Confluence 7.19.0 can only be restored to Confluence 7.19.0, 7.19.1, 8.0, 8.1 or newer. It can't be restored to Confluence 7.18.9, 7.18, or any other preceding version. That's because backward compatibility is not supported.

Recent backup and restore activity

View details about backup and restore tasks that have been run on your Confluence site over the last two weeks at > **General Configuration**> **Backup and restore**. This list includes space backups generated in **Space tools** but excludes backups generated byscheduled backups.

The recent activity table provides an overview of:

- task type (space backup, site backup, space restore, or site restore)
- file name of the zip file
- task status
- user who created the task, and
- date the task was created.

Screenshot: Backup and restore admin screen

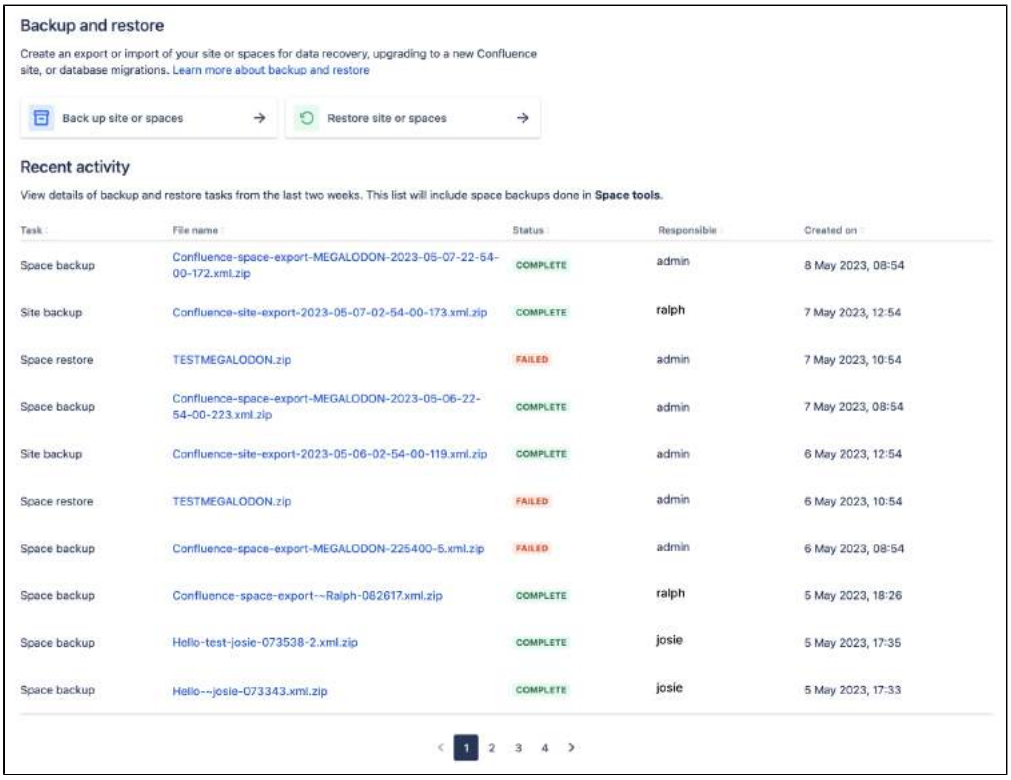

Select the file name to find out more details about a specific job. This will take you to the **Task details** page. The below table describes what information is available on this page.

Task detail	Description
File name	The name of the XML backup or restore file. For backups, you can select the file name to download a copy of the XML within 72 hours of exporting it (this is the default setting). If you chose to save permanently, the download link will not expire.
Status	This could be in progress, queued, canceling, canceled, failed, or completed. See below for more details on the queuedand failed statuses.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Spaces *Space backup and restore only*	It is the number of spaces contained within your XML space backup or restore zip file.
Records processed	This is the number of objects processed. Objects are the different records from the database related to your space or site.
Queue time	This is the amount of time your task was in the queue before it started processing. Learn more about the queued status below.
Time lapsed	This is the amount of time your task took to run until it was completed, canceled, or failed. Time lapsed doesn't include queue time.
Attachments	This specifies whether attachments have been included in your backup. Space backups will always include attachments.
Created on	This is the date and time the backup or restore task was created.
Created by	This is the user who created the backup or restore task.
Delete scheduled *Backup only*	This will only appear for backups that haven't been saved permanently. It is the date the XML backup zip file is scheduled to be deleted from the server. By default, this is 72 hours after the backup was created.
Records skipped *Restore only*	This is the number of records that were skipped while importing your site or space. Records can be skipped for a variety of reasons, and when they're skipped it doesn't always represent an error or data loss. If you suspect you're missing data from the restore, check the main application log to audit objects that were skipped.
Errors	This is a summary of the issues or errors found during the task. If you need more information about the issue or error, refer to the main application log.

Queued status

We only process one task at a time to provide a stable and fast backup and restore experience. That means, your task will be placed in a queue when there is another in progress, or if you created the task during a rolling upgrade.

If your task was queued, and someone restarts Confluence, your task will keep its spot in the queue.

Failedstatus

A back up and restore will fail if someone restarts Confluence while your task is in progress. In this case, you'll need to run the task again.

A restore also fails if you try to import a file that doesn't match your job, for example, importing a site backup when running a space restore.

Migrate to Confluence Cloud

If you're migrating from Confluence Server to Confluence Cloud, you can use theConfluence Cloud Migration Assistantto migrate your content and spaces.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Production Backup Strategy

Although Confluence can provide a scheduled XML backup, this backup method is only suitable for small sites, test sites, or in addition to database and directory backups.

> ⚠ Since Confluence 8.3, we have changed the way we do backup and restore. Learn more about these changes in the Confluence 8.3 Release Notes.
>
> Scheduled backups don't use this new approach; it still uses our legacy system. We recommend you stop using scheduled backups unless required because it contains many of the issues resolved in the new system.

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:

- Backup and Restore

Establishing a production system backup solution

We recommend establishing a robust database backup strategy:

- Create a backup of your database using the tools provided by your database. If your database doesn't support online backups, you will need to stop Confluence while you do this.
- Create a copy of your home directory (both local home and shared home for Data Center).

Once this is in place, you can disable any scheduled backup job you have set.

Having a backup of your database and home directories is more reliable and easier to restore than a large XML backup.

Which files need to be backed up?

Backing up the whole home directory is the safest option, however most files and directories are populated on startup and can be ignored. At minimum, these files/directories*must*be backed up:

- **<conf-home>/confluence.cfg.xml**
- **<conf-home>/attachments**(you can exclude extracted text files if space is an issue)

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

- **<conf-home>/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 onMigrating Confluence Between Servershas instructions on restoring a backup using this technique.

Other processes

XML site backups can be used for other processes in Confluence, for example moving servers or switching to a different database. Using the backup strategy described above will work for those processes too.

- Ourmigrate server procedure, which is used to set up a test server, can use a SQL dump as well.
- Thedatabase migration procedure uses the XML backup. You could also use third-party database migration tools.

If you would like help selecting the right migration tools, or help with the migration itself, reach out to one of our Atlassian Solution Partners.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Scheduling a Backup

Since Confluence 8.3, we have changed the way we perform backup and restore. Learn more about these changes in the Confluence 8.3 Release Notes.

Scheduled backups don't use this new approach; it still uses our legacy system. We recommend you stop using scheduled backups unless required because it contains many of the issues resolved in the new system.

See Production Backup Strategy for recommended methods.

On this page:

- Configure automated backups
- Perform manual backups

Confluence can automatically back up your data by performing a full site export at a scheduled time each day.

- If you have a **Confluence Data Center** license, the scheduled backup is disabled by default. It is known to cause outages in large sites.
- If you have a **Confluence Server** license, the scheduled XML backup happens every day at 2:00 am by default.

The zipped XML backup file will be named'`backup-yyyy_MM_dd`', and stored in the`backups`directory of yourConfluence Home directory. For example, `<shared-home>/backups`

This page describes how you can:

- enable or disable the scheduled backups (in Confluence Data Center it is disabled by default)
- change the naming convention
- include or exclude attachments
- schedule the backup at a different time
- store the backup files in a different location

You'll need System Administrator permissions to do this.

Configure automated backups

To configure these daily scheduled XML site backups:

1. Go to > **General Configuration**> **Backup administration**
2. Select **Edit** to:
    - change the backup file name prefix
    - use a different date format (uses the syntax described in simple date format)
    - choose whether to include or exclude attachments from backups (attachments are included by default)
    - choose to store backup files in a custom location (this is disabled by default - see Enabling backup path configuration below)
3. **Save** your changes

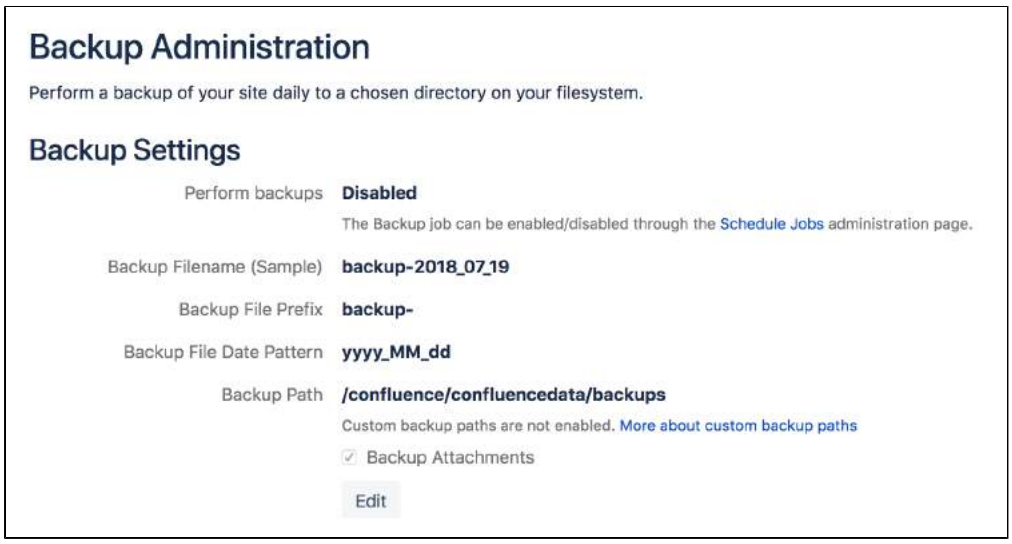

Enable Backup Path Configuration

For security reasons, the ability to change the backup file location **Backup administration** screen is disabled by default.

To enable custom backup paths:

1. Stop Confluence
2. Edit the `<confluence-home>/confluence.cfg.xml` file
3. Change the value of the following property to true:

```
<property name="admin.ui.allow.daily.backup.custom.location">true</property>
```

4. Restart Confluence to pick up the change
5. Go to > **General Configuration**>**Backup administration** to enter the new path

The directory must be on either a local drive or amounted network drive. Make sure the mounted drive is on a physical server and not a Virtual Machine image.

If you migrate Confluence to a new server or change your architecture, you will need to update this path. Changing your home directory location will not automatically update your backup file path if you've enabled a custom path.

Enable/disable scheduled backups

If you have an appropriateProduction Backup Strategy, you should disable automatic backups to save on disk space.

To turn on/off scheduled backups:

1. Go to > **General Configuration**>**Scheduled jobs**
2. Choose**Disable**next to the**Back up Confluence** job

Change the backup schedule

To change the frequency of backups, or to change the time the backup runs each day:

1. Go to > **General Configuration**>**Scheduled jobs**
2. Choose**Edit**next to the**Back up Confluence**job
3. Enter the new schedule using acron expression

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

The time zone used for the scheduled job is taken from the server on which Confluence is running. Go to > **General Configuration** > **System Information** to look up the **System Time**

Perform manual backups

If you need a one-off XML backup, you can manually perform a site export. See Manually Backing Up the Sitefor more information.

These backup files are not saved to the same location as the scheduled backups, they are saved in the restore directory of your Confluence Home directory, for example `<shared-home>/restore/site`.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

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
```


Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Back up a Site

You can back up Confluence at any time by performing a full site export.

You'll needSystem Administrator permissions to do this.

Related pages:

- Restore a Site
- Scheduling a Backup
- Production Backup Strategy

Good to know:

- We don't recommend you rely on XML backups as your main backup method. Instead, you should regularly back up your database, installation directory, and home directories. See Production Backup Strategy for recommended methods.
- Marketplace and user-installed apps are not included in the XML backup, however their data is. After restoring your site export file into a new Confluence site, you'll need to re-install all apps that are not bundled with Confluence as the `plugindata` table is not exported in this kind of backup.
- You can't restore (import) an XML backup file into **an earlier version** of Confluence. Backward feature compatibility is not supported.
- You can't restore an XML site backup file into Confluence Cloud either. Use the Confluence Cloud Migration Assistant app to do this. This is pre-installed with your Confluence installation.

Whats included in a site backup (export)?

A site backup includes, but is not limited to:

- spaces
- pages
- blog posts
- comments
- attachments (optional)
- draft pages and blogs
- user and groups
- app data

Basically, its everything on your site except user-installed apps.

Create the site backup (export)

To create an XML export of your site:

1. Go to > **General Configuration**>**Backup and restore**
2. Select**Back up site or spaces**
3. In the **Create a backup** screen:
    a. Select **Site**
    b. Give your XML export a file name prefix. Confluence will add a time stamp and unique identifier to the end of this file name. If this field is left blank, the file will be assigned a default prefix and time stamp, for example, `Confluence-site-export-2015-04-14-11-07-36-639.xml.zip`
    c. Select**Save permanently** if you want your file to remain in the `<confluence-home>/restore/site` folder. Otherwise, the backup will be deleted in 72 hours by default to maintain storage capacity. You can change the storage time by configuring the system property `confluence.backuprestore.backup.ttl-in-hours`. See Configuring System Properties to learn how.
    d. Select**Include attachments** to include attachments in your backup
4. Select **Back up**
5. You will receive a confirmation message, select**Back up now**to start the task

This process can take some time for large sites.

Retrieve the site backup

Method 1

Once the back up is complete, you can download the XML backup zip file by selecting the file name. This download link will remain active for 72 hours by default, unless you have chosen to save permanently.

Method 2

Confluence will create your XML backup zip file in your `<home-directory>/restore/site` folder, where your `<home-directory>` is `shared-home-directory` for Data Center, and `home-directory` for Server. See Confluence Home and other important directories for more information.

You'll need access to the Confluence server in order to retrieve this file this way.

Restore (import) the site backup

There are some restrictions on which Confluence versions you will be able to import this file into. The most important is that you **can't import into an earlier version of Confluence or to Confluence Cloud**. See Restore a Site for more information and troubleshooting tips.

Scheduled backups

Confluence can also be configured to automatically back up your data by performing a full site export at a scheduled time each day.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Back up a Space or multiple Spaces

You can export a single space or multiple spaces to an XML backup in the Confluence administration console.

Multi-space backup is helpful when you want to export a large site in smaller batches or generate a backup based on a team, project, or topic.

You'll needSystem Administrator permissions to do this.

A single space XML backup can also be generated by Space Admins in Space tools.

Related pages:

- Restore a Space or multiple Spaces
- Production Backup Strategy
- Export Content to Word, PDF, HTML and XML

Good to know:

- You can't restore a multi-space XML backup into an earlier version of Confluence.
- You cant restore a multi-space XML backup to Confluence Cloud either. Use the Confluence Cloud Migration Assistant app to do this. This is pre-installed with your Confluence installation.

Whats included in a space backup (export)?

A space backup includes, but is not limited to:

- pages
- blogs
- comments
- attachments
- drafts

Basically, its everything in your space.

Create a single or multi-space backup (export)

To create an XML backup of one or more spaces:

1. Go to > **General Configuration**>**Backup and restore**
2. Select**Back up site or spaces**
3. In the **Create a backup** screen:
    a. Select **Spaces**
    b. Select the spaces you want to include in your XML export file. You can type the name of a space, or select recent spaces from the dropdown menu.
    c. Give your XML export a file name prefix. Confluence will add a time stamp and unique identifier to the end of this file name. If this field is left blank, the file will be assigned a default prefix and time stamp, for example, for single spaces `Confluence-space-export-<SPACE KEY>-2015-04-14-11-07-36-639.xml.zip` or for multi-space backups, `Confluence-space-export-<NUM SPACES>-spaces-2015-04-14-11-07-36-639.xml.zip`
    d. Select**Save permanently** if you want your file to remain in the `<confluence-home>/restore/space`folder. Otherwise, the backup will be deleted in 72 hours by default to maintain storage capacity. You can change the storage time by configuring the system property `confluence.backuprestore.backup.ttl-in-hours`. See Configuring System Properties to learn how.
4. Select **Back up**
5. You will receive a confirmation message, select**Back up now**to start the task

This process can take some time if you have many large spaces selected.

Retrieve the space backup

Method 1

Once the back up is complete, you can download the XML backup zip file by selecting the file name. This download link will remain active for 72 hours, unless you have chosen to save permanently.

Method 2

Confluence will create the backup as a zipped XML file in your `<home-directory>/restore/space` folder, where `<home directory>` is `shared-home-directory` for Data Center, and `home-directory` for Server. See Confluence Home and other important directories for more information.

You'll need access to the Confluence server in order to retrieve the file this way.

Restore (import) the space backup

There are some restrictions on which Confluence versions you will be able to import this file into. The most important is that you **can't import into an earlier version of Confluence or to Confluence Cloud**. See Restore a Space or multiple Spaces for more information and troubleshooting tips.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Restore a Site

This page describes how to restore data from an XML backup into a new or existing Confluence site.

You needSystem Administrator permissions to do this.

On this page:

- Before you start
- Check your backup is compatible
- Restore (import) a site from an XML backup
- Restore from a scheduled backup
- Troubleshooting

Related pages:

- Production Backup Strategy
- Exporting a site
- Importing a Space

Restoring a site backup will:

- Overwrite all existing Confluence data in your database. Back up your database before you start, see Production Backup Strategy for recommended methods.
- Log you out of Confluence. Make sure you know the login details contained in the file you're about to import.

Before you start

- **All content replaced.** Importing a site will replace all your content and users. Make sure you have a copy of your database.
- **Selective space restoration not possible.** You can't select a single space to restore from the entire site backup. Instead, you can manually back up a space or multiple spaces, and then restore that XML.
- **XML export files should not be used to upgrade Confluence.** Upgrade Confluence by followingUpgrading Confluence.
- **If you are migrating to Confluence Cloud**, use the Confluence Cloud Migration Assistant app that is pre-installed with Confluence Data Center and Server.
- **Stop your Synchrony standalone cluster**. If you use Confluence Data Center, you'll need to stop your Synchrony standalone cluster completely before you restore a site. Once the restore is complete, you can restart your Synchrony cluster. This is not required if you allow Confluence to manage Synchrony for you.

Check your backup is compatible

You can't restore a backup into **an earlier version** of Confluence.

For example, if your XML backup was generated from Confluence 8.3, you can't import it into Confluence 7.19.

To check whether your backup can be successfully restored:

- Check which Confluence version you are using in > **General Configuration> System Information.** The version will be listed next to **Confluence Version**.

- Check which Confluence version your XML backup was generated from. See How to Determine XML Backup Confluence Version.

If you are restoring a backup to a later version, it can be restored successfully.

If you are restoring a backup to an earlier version, this is not supported and your import may fail.

Restore (import) a site from an XML backup

There are two ways to restore a site - by uploading a file (local drive), or from a directory on your Confluence server (restore directory).

Uploading a file is only suitable for small sites. For best results, we recommend importing from the restore directory.

Upload a file

To upload and restore a small site:

1. Go to > **General Configuration**>**Backup and restore**
2. Select **Restore site or spaces**
3. Under **What do you want to restore?**, select **Site**
4. Under **Restore from**, select **Local drive**
5. Under**Select file,** select **select XML zip** and browse for your XML site export file
6. Select **Upload and restore**
7. You will see a confirmation message, read this carefully as your site will not be accessible and its contents will be completely overwritten once the task starts
8. Select**Upload and restore** once more

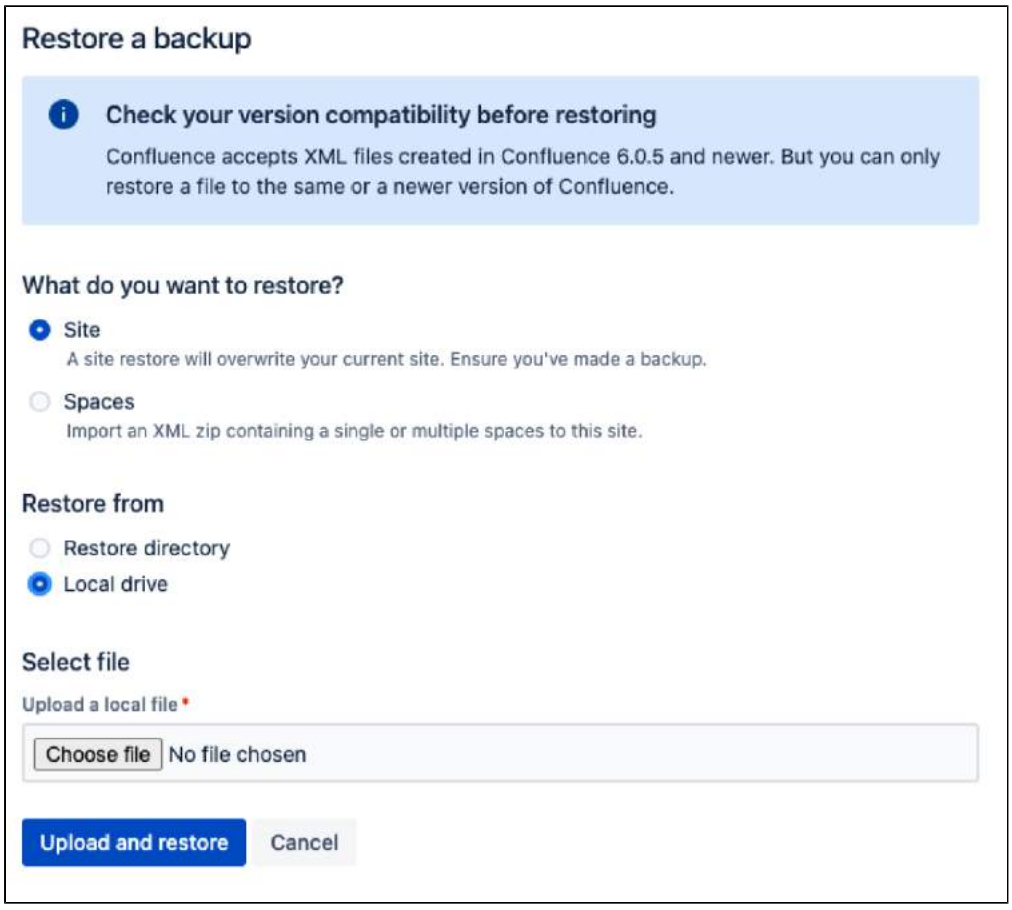

Restore from server

To restore a site from the home directory:

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

1. Copy the XML backup to `<confluence-home>/restore/site`. If you're not sure where this directory is located, see Confluence Home and other important directories.
2. Go to > **General Configuration**>**Backup and restore**.
3. Select **Restore site or spaces**
4. Under **What do you want to restore?**, select **Site**
5. Under**Restore from**, select**Restore directory**
6. Under **Select file**, type or browse for your site export file in the dropdown menu
7. Select **Restore**
8. You will see a confirmation message, read this carefully as your site will not be accessible and its contents will be completely overwritten once the task starts
9. Select**Restore now**

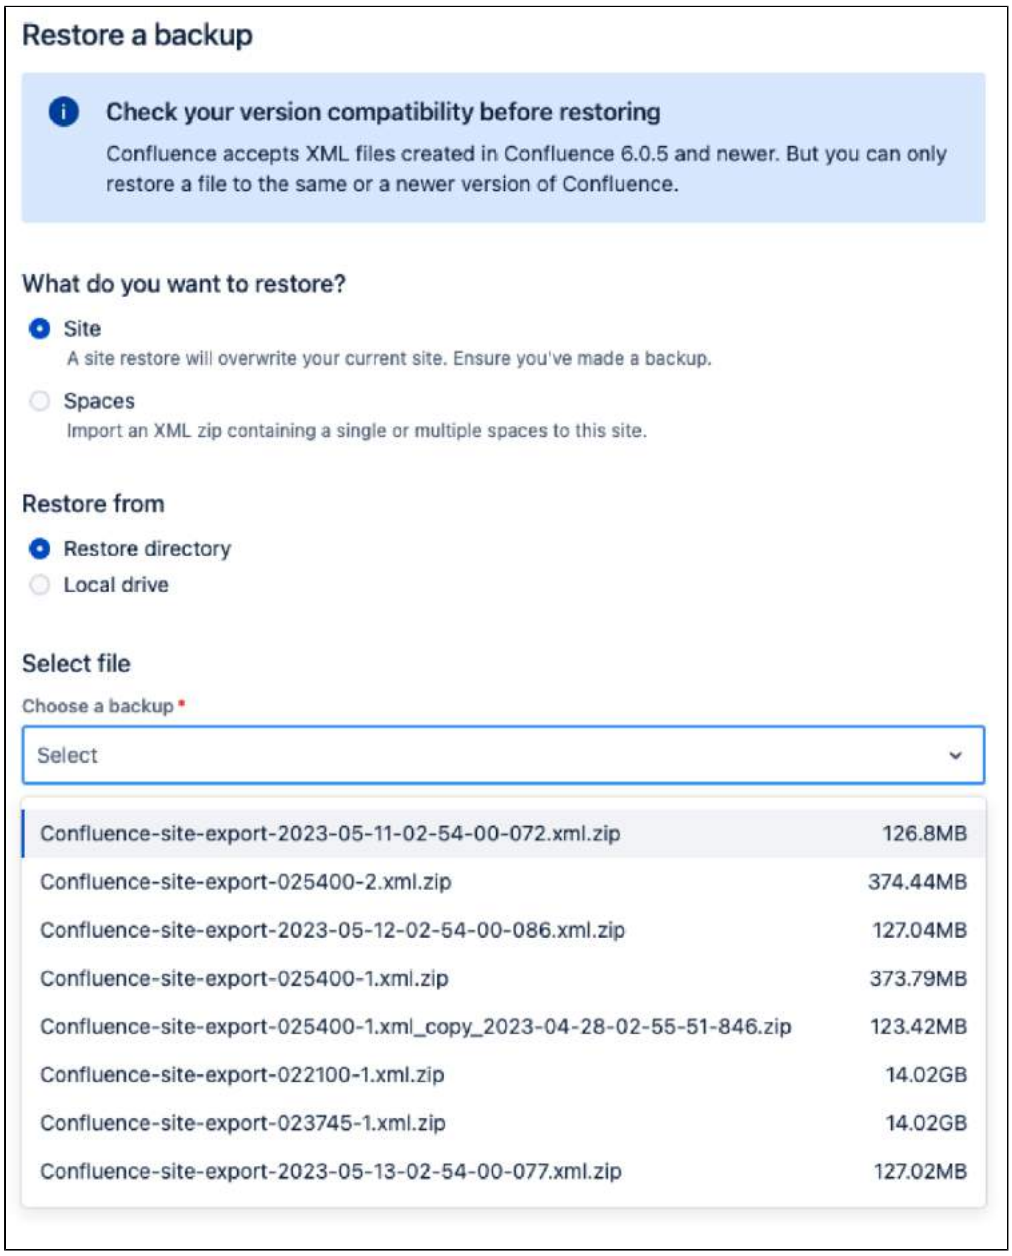

Site reindexing

A full site re-index will start when the site has been restored. The site restore task isn't finished until this re-indexing job is complete.

Canceled site restore

If you cancel a site restore while its status is queued, the task will simply be canceled and you will be shown the task details at that point.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

If you cancel the restore while it is "in progress" and your site is inoperable, this could corrupt your instance since the overwrite will have already started. You will be required to restore from a database backup in order to revert the site back to its previous state.

Restore from a scheduled backup

Scheduled backups are saved in `<confluence-home>/backups` so you will need to retrieve the XML backup from the correct folder, then copy it to the restore directory, and follow the steps listed above in order to restore the site from the restore directory.

We dont recommend uploading scheduled backups of large sites from your local drive to restore it.

Troubleshooting

If you have problems restoring a site, check out these hints.

- **Is your file too large to upload?**
  This is a very common problem. It happens when the file can't be uploaded to the server in time. To avoid this problem, drop your export file into the`<confluence-home>/restore/site` directory and restore it from there.
- **Are you trying to import into an earlier version of Confluence?**
  This is not possible. You can only restore a site into the same version or a later compatible version.
- **Is the import timing-out or causing out of memory errors?**
  If the site to be restored is large, you may need to temporarily increase the memory available to Confluence. See How to fix out of memory errors by increasing available memory.
- **Is the import taking a very long time?**
  The size of the`entities.xml` file provides an indication of the amount of entities such as pages, versions, and comments. A very large site may take quite some time to restore. You can check the size of this file by unzipping the export file.
- **Is your username or password not recognized?**
  All user data was overwritten during the restore process. You need to log in with a system administrator account from the site that was exported. If you don't know the password, you'll need to reset it from the database. See Restore Passwords To Recover Admin User Rights.
- **Is your site export from Confluence Cloud?**
  You can only restore into Confluence 6.0 or later. The Cloud export does not include a system administrator account, so you will need to start Confluence in recovery mode, create a new system administrator account, and make it a member of the confluence-administrators group. See Restore Passwords To Recover Admin User Rights for more.
- **Did you download the export file on a Mac?**
  If you get an error saying thatConfluence can't find the`exportDescriptor.properties`file, chances areOS X has unzipped the backup for you and sent the original zipped file to the trash. You need to retrieve the original zip file from the trash and then try the import again.
- **Restoring into a site with a Synchrony standalone cluster?**
  You must stop your Synchrony cluster before commencing the site restore.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Restore a Space or multiple Spaces

This page describes how to restore data from an XML backupinto a new or existing Confluence site.

You can create a space backup in the Confluence administration console or in Space tools for single-space backups only.

However, to restore an XML space backup you'll need System Administrator permissions.

On this page:

- Before you start
- Check your backup is compatible to be restored
- Restore (import) a space from Confluence Cloud
- Restore (import) a space from Confluence Data Center or Server
  - Upload a file
  - Space indexing
  - Groups and permissions
  - Canceled or failed space restore
- Troubleshooting

Related pages:

- Restore a Site

Before you start

- **Identical space keys will cause issues.** Make sure the spaces you plan to restore dont have the same space keys as any spaces in the destination site; this will cause the task to fail.
- **XML export files should not be used to upgrade Confluence.** Upgrade Confluence by followingUpgrading Confluence.
- **If you are migrating to Confluence Cloud**, use the Confluence Cloud Migration Assistant app. This is pre-installed with your Confluence installation.

Check your backup is compatible to be restored

> You can't restore a backup into **an earlier version** of Confluence.
>
> For example, if your XML backup was generated from Confluence 8.3, you can't import it into Confluence 7.19.

To check whether your backup can be successfully restored:

- Check which Confluence version you are using in > **General Configuration> System Information.** The version will be listed next to **Confluence Version**.
- Check which Confluence version your XML backup was generated from. See How to Determine XML Backup Confluence Version.

If you are restoring a backup to a later version, it can be restored successfully.

If you are restoring a backup to an earlier version, this is not supported and your import may fail.

Restore (import) a space from Confluence Cloud

As the way users are managed is different in Confluence Cloud there are a few more considerations when restoring a space from Confluence Cloud into Confluence Data Center or Server.

SeeImport a space from Confluence Cloud for a step-by-step guide.

Restore (import) a space from Confluence Data Center or Server

⚠ We recommend performing a full backup of your database before restoring your XML space backup. Occasionally the space restore task may fail, and a backup will make it easier for you to roll back.

There are two ways to restore a space or multiple spaces by uploading a file (local drive), or from a directory on your Confluence server (restore directory).

Uploading a file is only suitable for small spaces. For best results, we recommend restoring from the restore directory. A very large file may take quite some time to restore.

Upload a file

To upload and restore small spaces:

1. Go to > **General Configuration**>**Backup and restore**
2. Select **Restore site or spaces**
3. Under **What do you want to restore?**, select **Spaces**
4. Under **Restore from**, select **Local drive**
5. Under**Select file,** select **Choose file** and browse for your XML backup
6. Select **Upload and restore**
7. You will see a confirmation message
8. Select**Upload and restore** once more

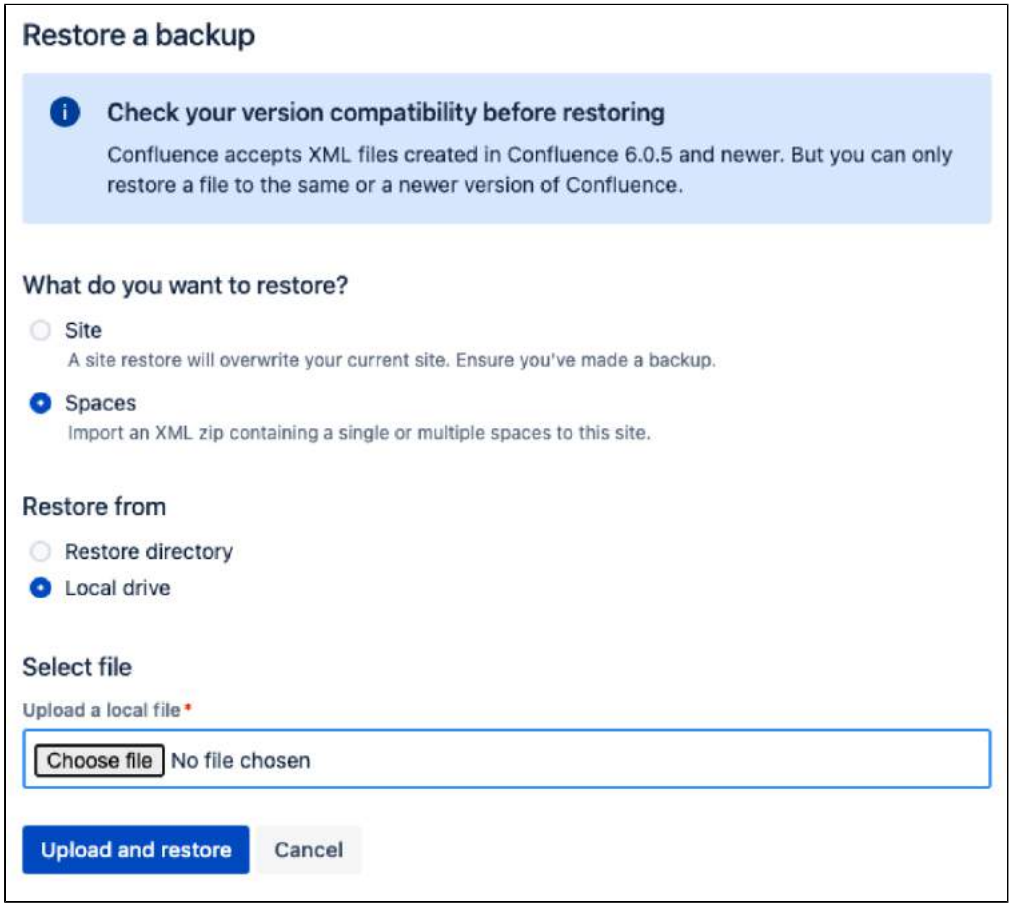

Restore from server

Restoring from the home directory is a reliable alternative for large spaces.

To restore a space from the home directory:

1. Copy the XML backup to `<confluence-home>/restore/space`. If you're not sure where this directory is located, see Confluence Home and other important directories.
2. Go to > **General Configuration**>**Backup and restore**.
3. Select **Restore site or spaces**
4. Under **What do you want to restore?**, select **Spaces**
5. Under**Restore from**, select**Restore directory**
6. Under **Select a file**, type or browse for your space export file in the dropdown menu
7. Select **Restore**
8. You will see a confirmation message
9. Select**Restore now**

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

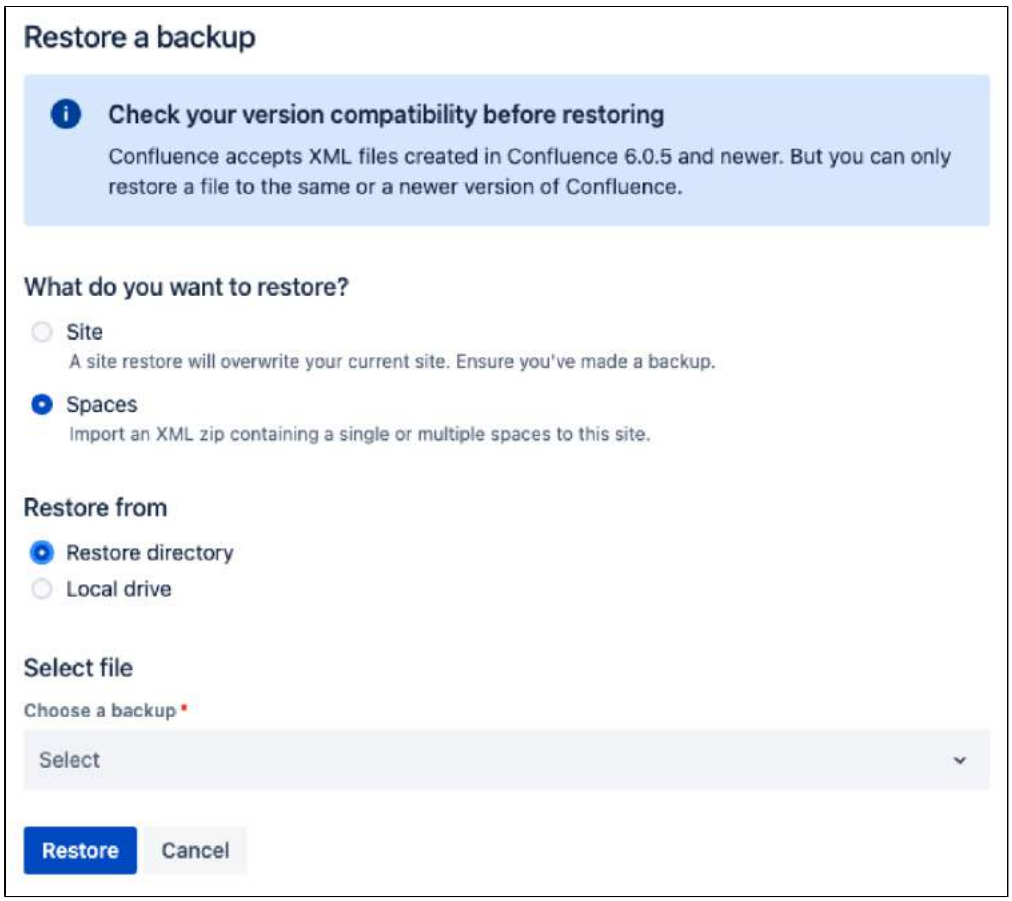

Space indexing

As you restore a space or multiple spaces, they will be indexed in the same way page index usually happens. That is, records will be added to a queue to be indexed and there is a job that checks these records. Indexing will happen when the record is picked up by this job. For this reason, there may be a small delay before a user is able to see the space in the **Space directory** or search for content within the spaces restored.

Confluence doesn't do a full site re-index for space and multi-space restore.

Groups and permissions

Importing a space will not import any users or groups that may have been granted specific space permissions in your source Confluence site. This means that if any pages are restricted to these groups, you may not be able to see them until you recreate these groups in your destination site.

Canceled or failed space restore

If you cancel a space restore while it is queued, the task will simply be canceled and you will be shown the task details at that point.

If you need to cancel the space restore while it is "in progress", the data that has already been processed and imported will need to be cleaned up before you can restore the same spaces to the destination site. You can still restore other spaces or restore a site.

For example, if you attempted to restore Spaces with keys A, B, and C, but the space restore failed or was canceled after it had already started, you will still be able to restore the Space with key D, but you cant restore Spaces with keys A, B, or C again until youve removed any partially imported data. See Clean up leftover space data from a failed space import to learn how to remove this data.

Troubleshooting

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

If you have problems importing a space, check out these hints.

- **Is your file too large to upload?**
  This is a very common problem. It happens when the file can't be uploaded to the server in time. To avoid this problem, drop your backup file into the `<confluence-home>/restore/space` and restore it from there.
- **Are you trying to import into an earlier version of Confluence?**
  This is not possible. You can only restore a space into the same version or a later compatible version.
- **Does a space with the same space key already exist?**
  Space keys are unique, soif you already have a space with the same key, you'll need to delete the existing space before restoring the new one.
- **Is the import timing-out or causing out-of-memory errors?**
  If the space to be restored is very large, you may need to temporarily increase the memory available to Confluence. See How to fix out-of-memory errors by increasing available memory.
- **Did you download the export file on a Mac?**
  If you get an error saying thatConfluence can't find the `exportDescriptor.properties` file, chances areOS X has unzipped the backup for you and sent the original zipped file to the trash. You need to retrieve the original zip file from the trash and then try the restore again.
- **Did your import fail?** Sometimes restoring a space may fail because of invalid data. This can lead to data being left behind in your database. You'll need to clean up leftover space data from a failed space import before you attempt to restore the same spaces again.
- **Do you have adequate disk space?**
  Confluence will need to make copies of the backup file being restored at various points during the restore process. Make sure you have enough disk space to temporarily accommodate multiple copies of the file.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Restore a Test Instance from Production

See Migrating Confluence Between Servers for a more comprehensive explanation.

Many Confluence administrators will have a production instance running the "live" version of Confluence, as well as a test instance for testing upgrades and so on. In this situation, it's quite common that the two instances are running different versions of Confluence. This document describes how to copy the data from a production instance to a test instance, where the production version may be different to the test version.

Before proceeding with this guide, ensure you have read and understood the normal procedure for upgrading Confluence.

Updating a test Confluence instance with production data

Essentially, we are copying both the production home directory and database to the test instance. We then update the database details on the test instance to point to the test database, leaving all other instance metadata (most importantly the Confluence *build number*) the same as production.

1. Shut down your test instance.
2. Restore the production database to the test database server.
3. Create a backup of the `confluence.cfg.xml` file found in the home directory of the test instance.
4. Copy the production confluence-home directory to the test application server.
5. Open the `confluence.cfg.xml` which has been copied in a text editor. Change the database settings to match the test database server. **Ensure you do not point to your production database.** (You can compare with the backup you made in Step 3 if you need to get the database settings. Don't just copy this file you need the build number unchanged from production to indicate the database is from an older version of Confluence.)

Before starting your test instance, you need to do the following steps to ensure no contact with production systems.

Ensuring no contact with production systems

To ensure no contact with external systems, you will need to disable both inbound and outbound mail services.

1. Disable global outbound mail by running the following database query:

   ```
   SELECT * FROM BANDANA WHERE BANDANAKEY = 'atlassian.confluence.smtp.mail.accounts';
   ```

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.

Restore data from other backups

Typically Confluence data is restored in the administration console. See Restore a Site or Restore a Space or multiple Spaces for instructions.

If you are experiencing problems restoring an XML backupfile using these methods, it is still possible to restore the data provided you have:

- a backup of your home directory
- a backup of your database

If youre using an external database like Oracle, MS SQL Server, MySQL or Postgres, 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 file system and the database is available for connection.
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 restore the backup into Confluence. This is useful for recovering attachments that have been deleted by users.

Both scheduled and manual backups allow this, as long as the 'Include attachments' property was set.

Before following the instructions for recovering attachments below, we will review how backups store file and page information.

On this page:

- Backup zip file structure
- Entities.xml Attachment Object
- Entities.xml Page Object

Instructions for recovering attachments

Related pages:

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:

Confluence 8.0 and earlier

```
Listing for test-2006033012_00_00.zip
\attachments\98\10001
\attachments\98\10002
\attachments\99\10001
entities.xml
```

Confluence 8.1 and later

```
Listing for test-2006033012_00_00.zip
\attachments\98\10001\1
\attachments\98\10002\1
\attachments\99\10001\3
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">10001</id>
    <property name="title">myimportantfile.doc</property>
    <property name="lowerTitle">myimportantfile.doc</property>
    <property name="version">1</property>
    ...
    <property name="containerContent" class="Page" package="com.atlassian.confluence.pages"><id name="id"
>98</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:

To recover the latest version of each attachment 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

This option is 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 directory with the attachment id
4. Inside the attachment directory rename the file with the highest number to the original filename and test it
5. Repeat for each attachment directory
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

This option is 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. Rename this directory to the page title
4. Search entities.xml for attachment objects with that page id. Every attachment object for the page will have an attachment id, version and filename
5. For each attachment object find the attachment directory and rename the file with the highest number (latest version) 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

This option is best if you have a small backup but want to restore many or all the attachments inside.

> The following process is applicable to **space** backups only. Site XML backups do not require page id to be updated manually due to the nature of persistent page_ids.

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

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

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

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Troubleshooting failed XML site backups

Since Confluence 8.3, we have changed the way we do backup and restore. Learn more about these changes in the Confluence 8.3 Release Notes.

As a result of this change, many issues with the old system were resolved. That means the recommendations listed below will not be applicable to backup and restore anymore.

XML site backups are only necessary for migrating to a new database. Setting up a test server or Establishing a reliable backup strategy is better done with an SQL dump.

Related pages:

- Enabling detailed SQL logging

Seeing an error when creating or importing a backup?

Problem	Solution
Exception while creating backup	Follow instructions below
Exception while importing backup	Follow Troubleshooting XML backups that fail on restore instead

Common problems

- **Is the export timing out or causing out of memory errors?**
  If your site is large, you may need to temporarily increase the memory available to Confluence. See How to fix out of memory errors by increasing available memory

Resolve errors from manual XML backup

The errors may be caused by a slightly corrupt database. If you're seeing errors such as 'Couldn't backup database data' in your logs, this guide will help you correct the error on your own. We strongly recommend that you backup your database and your Confluence home directory beforehand, so that you can restore your site from those if required. If you are unfamiliar with SQL, we suggest you contact your database administrator for assistance.

Preferable solution

Identify and correct the problem

To work out where the data corruption or problems are, increase the status information reported during backup, then edit the invalid database entry:

1. Stop Confluence.
2. If you have an external database, use a database administration tool to create a manual database backup.
3. Backup your Confluence home directory. You will be able to restore your whole site using this and the database backup.
4. Open the `my_confluence_install/confluence/WEB-INF/classes/log4j.properties`and add this to the bottom and save:

```
log4j.logger.com.atlassian.hibernate.extras.XMLDatabinder=DEBUG, confluencelog
log4j.additivity.com.atlassian.hibernate.extras.XMLDatabinder=false
```

5. Find your application logs. Move or delete all existing Confluence application logs to make it easier to find the relevant logging output. You could also choose to mark the application logs after restarting Confluence, to indicate when you started the export.
6. Restart Confluence and login.
7. Begin a backup so that the error reoccurs.
8. You must now check your log files to find out what object could not be converted into XML format. Open `confluence-home/logs/atlassian-confluence.log`. Scroll to the bottom of the file.
9. Do a search for 'ObjectNotFoundException'. You should see an error similar to this:

```
01 2005-08-24 00:00:33,743 DEBUG [DOCPRIV2:confluence.importexport.impl.XMLDatabinder] Writing
object: com.atlassian.confluence.core.ContentPermission with ID: 5 to XML.
02 2005-08-24 00:00:33,743 DEBUG [DOCPRIV2:confluence.importexport.impl.XMLDatabinder] Writing
property: type
03 2005-08-24 00:00:33,743 DEBUG [DOCPRIV2:confluence.importexport.impl.XMLDatabinder] Writing
property: group
04 2005-08-24 00:00:33,743 DEBUG [DOCPRIV2:confluence.importexport.impl.XMLDatabinder] Writing
property: expiry
05 2005-08-24 00:00:33,743 DEBUG [DOCPRIV2:confluence.importexport.impl.XMLDatabinder] Writing
property: content
06 [DOCPRIV2:ERROR] LazyInitializer - Exception initializing proxy <net.sf.hibernate.
ObjectNotFoundException: No row with the given identifier exists: 2535,
07 of class: com.atlassian.confluence.core.ContentEntityObject>net.sf.hibernate.
ObjectNotFoundException:
08 No row with the given identifier exists: 2535, of class: com.atlassian.confluence.core.
ContentEntityObject
09        at net.sf.hibernate.ObjectNotFoundException.throwIfNull(ObjectNotFoundException.java:24)
10        at net.sf.hibernate.impl.SessionImpl.immediateLoad(SessionImpl.java:1946)
11        at net.sf.hibernate.proxy.LazyInitializer.initialize(LazyInitializer.java:53)
12        at net.sf.hibernate.proxy.LazyInitializer.initializeWrapExceptions(LazyInitializer.java:
60)
13        at net.sf.hibernate.proxy.LazyInitializer.getImplementation(LazyInitializer.java:164)
14        at net.sf.hibernate.proxy.CGLIBLazyInitializer.intercept(CGLIBLazyInitializer.java:108)
15        at com.atlassian.confluence.core.ContentEntityObject$$EnhancerByCGLIB$$cc2f5557.hashCode
(<generated>)
16        at java.util.HashMap.hash(HashMap.java:261)
17        at java.util.HashMap.containsKey(HashMap.java:339)
18        at com.atlassian.confluence.importexport.impl.XMLDatabinder.toGenericXML(XMLDatabinder.
java:155)
```

10. Open a DBA tool such as DbVisualizer and connect to your database instance. Scan the table names in the schema. You will have to modify a row in one of these tables.
11. To work out which table, open `atlassian-confluence.log`, check the first line of the exception. This says there was an error writing the `ContentPermission` object with id 5 into XML. This translates as *the row with primary key 5 in the CONTENTLOCK* table needs fixing. To work out what table an object maps to in the database, here's a rough guide:
    - Pages, blogposts, comments --> CONTENT table
    - attachments --> ATTACHMENTS table
    - More information can be found in the schema documentation
12. Now you must find the primary key of the incorrect row in this table. In this case, you can check the first line and see that the row has a primary key of 5.
13. Each property is written to a column, so the last property that was being written has the incorrect value. The row being written to when the exception was thrown was `CONTENT` (line 5) with a value of 2 `535` (line 6). Now you know the column and value. This value `2535` is the id of an entry that no longer exists.
14. Using a database administrative tool, login tothe 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

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

If you are encountering an error message such as:

```
could not insert: [bucket.user.propertyset.BucketPropertySetItem#bucket.user.propertyset.
BucketPropertySetItem@a70067d3]; SQL []; Violation of PRIMARY KEY constraint
'PK_OS_PROPERTYENTRY314D4EA8'. Cannot insert duplicate key in object 'OS_PROPERTYENTRY'.; nested
exception is java.sql.SQLException: Violation of PRIMARY KEY constraint 'PKOS_PROPERTYENTRY_314D4EA8'.
Cannot insert duplicate key in object 'OS_PROPERTYENTRY'.
```

this indicates that the Primary Key constraint 'PK_OS_PROPERTYENTRY_314D4EA8' has duplicate entries in table 'OS_PROPERTYENTRY'.
You can locate the constraint key referring to 'PK_OS_PROPERTYENTRY_314D4EA8' in your table 'OS_PROPERTYENTRY' and locate any duplicate values in it and remove them, to ensure the "PRIMARY KEY" remains unique. An example query to list duplicate entries in the 'OS_PROPERTYENTRY' table is:

```
SELECT ENTITY_NAME,ENTITY_ID,ENTITY_KEY,COUNT(*) FROM OS_PROPERTYENTRY GROUP BY ENTITY_NAME,ENTITY_ID,
ENTITY_KEY HAVING COUNT(*)>1
```

Prevent this issue from reoccurring

1. If you are using the embedded database, be aware that it is bundled for evaluation purposes and does not offer full transactional integrity in the event of sudden power loss, which is why an external database is recommended for production use. You should migrate to an external database.
2. If you are using an older version of Confluence than the latest, you should consider upgrading at this point.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Troubleshooting XML backups that fail on restore

Since Confluence 8.3, we have changed the way we do backup and restore. Learn more about these changes in the Confluence 8.3 Release Notes.

As a result of this change, many issues with the old system were resolved. That means the recommendations listed below will not be applicable to backup and restore anymore.

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.

On this page:

- Common problems
- Resolve errors when attempting to restore an XML backup
  - Troubleshooting "Duplicate Entry" for key "cp_" or "cps_"
  - Troubleshooting "Duplicate Key" related problems
  - Troubleshooting "net.sf.hibernate.PropertyValueException: not-null" related problems
  - To help prevent this issue from recurring

Related Topics:

- Troubleshooting failed XML site backups

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

Common problems

You can't restore a backup into **an earlier version** of Confluence.

For example, if your XML backup was generated from Confluence 8.3, you can't import it into Confluence 7.19.

To check whether your backup can be successfully restored:

- Check which Confluence version you are using in > **General Configuration> System Information.** The version will be listed next to **Confluence Version**.
- Check which Confluence version your XML backup was generated from. See How to Determine XML Backup Confluence Version.

If you are restoring a backup to a later version, it can be restored successfully.

If you are restoring a backup to an earlier version, this is not supported and your import may fail.

Resolve errors when attempting to restore an XML backup

The errors may be caused by a slightly corrupt database. You will need to find the XML backup file entry that is violating the DB rules, modify the entry and recreate the XML backup:

1. On the instance being restored, follow the instructions to disable batched updates (for simpler debugging), log SQL queries and log SQL queries **with parameters** at Enabling Detailed SQL Logging.
2. Once all three changes have been made, restart Confluence.

3. Attempt another restore.
4. Once the restore fails, check your application log files and the`catalina.<datestamp>.log` (in your installation directory) to find out what object could not be converted into XML format.
5. Scroll to the bottom of the file and identify the last error relating to a violation of the database constraint. For example:

```
2006-07-13 09:32:33,372 ERROR [confluence.importexport.impl.ReverseDatabinder] endElement net.sf.
hibernate.exception.ConstraintViolationException:
  could not insert: [com.atlassian.confluence.pages.Attachment#38]
net.sf.hibernate.exception.ConstraintViolationException: could not insert: [com.atlassian.
confluence.pages.Attachment#38]
...
Caused by: java.sql.SQLException: ORA-01400: cannot insert NULL into ("CONFUSER"."ATTACHMENTS"."
TITLE")
at oracle.jdbc.driver.DatabaseError.throwSqlException(DatabaseError.java:112)
at oracle.jdbc.driver.T4CTTIoer.processError(T4CTTIoer.java: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 the log filescheck the first line of the exception. To work out what table an object maps to in the database, here's a rough guide:
   - Pages, blogposts, comments --> CONTENT table.
   - attachments --> ATTACHMENTS table.
9. To correct the example error, go to the attachment table and find that attachment object with id 38. This will have a a null title. Give a title using the other attachments titles as a guide. You may have a different error and should modify the database accordingly.
10. Once the entry has been corrected, create the XML backup again.
11. Import the backup into the new version.
12. If the import succeeds, revert the changes made in your SQL logging to re-enable disable batched updates and turn off log SQL queries and log SQL queries with parameters.
13. Restart Confluence.

Troubleshooting "Duplicate Entry" for key "cp_" or "cps_"

If you are encountering an error message such as:

```
com.atlassian.confluence.importexport.ImportExportException: Unable to complete import because the data
does not match the constraints in the Confluence schema. Cause:
MySQLIntegrityConstraintViolationException: Duplicate entry '1475804-Edit' for key 'cps_unique_type'
```

This indicates that the XML export came from a version of Confluence with a corrupt permissions database, caused by some 3rd party plugin. This is an issue that was fixed when CONF-22123 was implemented in Confluence 3.5.2. The simplest workaround is to export the space again after upgrading the instance to 3.5.2 or above. If that is not an option, then either the export will need to be edited manually to remove the duplicate permission entries or the source instance will need to have the offending entries removed. The following SQL queries can be used to look for such entries:

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

```
SELECT * FROM CONTENT_PERM WHERE USERNAME IS NULL AND GROUPNAME IS NULL;

SELECT cp.ID, cp.CP_TYPE, cp.USERNAME, cp.GROUPNAME, cp.CPS_ID, cp.CREATOR,
cp.CREATIONDATE, cp.LASTMODIFIER, cp.LASTMODDATE
FROM CONTENT_PERM cp
WHERE cp.USERNAME IS NOT NULL AND cp.GROUPNAME IS NOT NULL;

SELECT cps1.ID, cps1.CONTENT_ID, cps1.CONT_PERM_TYPE FROM CONTENT_PERM_SET cps1, CONTENT_PERM_SET cps2
WHERE cps1.ID <> cps2.ID AND
cps1.CONTENT_ID = cps2.CONTENT_ID AND
cps1.CONT_PERM_TYPE = cps2.CONT_PERM_TYPE
ORDER BY cps1.CONTENT_ID, cps1.CONT_PERM_TYPE, cps1.CREATIONDATE ASC;

SELECT cp.ID, cp.CP_TYPE, cps.CONTENT_ID,
(SELECT scps.ID FROM CONTENT_PERM_SET scps WHERE scps.CONTENT_ID = cps.CONTENT_ID AND scps.
CONT_PERM_TYPE = cp.CP_TYPE) AS suggested_cps_id
FROM CONTENT_PERM cp, CONTENT_PERM_SET cps
WHERE cp.CPS_ID = cps.ID AND
cp.CP_TYPE <> cps.CONT_PERM_TYPE;

SELECT DISTINCT cp1.ID, cp1.CP_TYPE, cp1.USERNAME, cp1.GROUPNAME, cp1.CPS_ID,
cp1.CREATOR, cp1.CREATIONDATE, cp1.LASTMODIFIER, cp1.LASTMODDATE
FROM CONTENT_PERM cp1, CONTENT_PERM_SET cps1, CONTENT_PERM cp2, CONTENT_PERM_SET cps2
WHERE
cp1.CPS_ID = cps1.ID AND
cp2.CPS_ID = cps2.ID AND
cp1.ID <> cp2.ID AND
cps1.CONTENT_ID = cps2.CONTENT_ID AND
cp1.CP_TYPE = cp2.CP_TYPE AND
cp1.USERNAME = cp2.USERNAME
ORDER BY cp1.CPS_ID, cp1.CP_TYPE, cp1.USERNAME, cp1.CREATIONDATE;

SELECT DISTINCT cp1.ID, cp1.CP_TYPE, cp1.USERNAME, cp1.GROUPNAME, cp1.CPS_ID,
cp1.CREATOR, cp1.CREATIONDATE, cp1.LASTMODIFIER, cp1.LASTMODDATE
FROM CONTENT_PERM cp1, CONTENT_PERM_SET cps1, CONTENT_PERM cp2, CONTENT_PERM_SET cps2
WHERE
cp1.CPS_ID = cps1.ID AND
cp2.CPS_ID = cps2.ID AND
cp1.ID <> cp2.ID AND
cps1.CONTENT_ID = cps2.CONTENT_ID AND
cp1.CP_TYPE = cp2.CP_TYPE AND
cp1.GROUPNAME = cp2.GROUPNAME
ORDER BY cp1.CPS_ID, cp1.CP_TYPE, cp1.GROUPNAME, cp1.CREATIONDATE;

SELECT * FROM CONTENT_PERM_SET
WHERE ID NOT IN (SELECT DISTINCT CPS_ID FROM CONTENT_PERM);
```

Remove all matching entries and perform the export again.

Troubleshooting "Duplicate Key" related problems

If you are encountering an error message such as:

```
could not insert: [bucket.user.propertyset.BucketPropertySetItem#bucket.user.propertyset.
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
```


Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

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

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Import a space from Confluence Cloud

As of 16 July 2019, **usernames are no longer included** in space exports from Confluence Cloud.

Email addresses will be included, regardless of profile visibility settings, if the person performing the export is a **Site Admin**.

Email matching is available in selected Confluence versions. See below for more information.

The user base of your Confluence Cloud and Confluence Data Center or Server sites are separate. Although the same people may have accounts on both sites, the way information is stored about them is different. For example, in Confluence Cloud usernames have been replaced by email addresses, and they have an additional ID (a random string of characters) that acts as a unique identifier.

When you restore a space into Confluence we attempt to attribute content based on username. If the two usernames match, we will attribute content to the user.

In spaces exported from Cloud, where there is no username, we will attempt to match users by their email addresses. To reduce the risk of making restricted pages visible to the wrong person, restored content will be attributed to 'unknown user' if:

- the email address is used by multiple user accounts (with different usernames), or
- the user account doesn't have an email address (for example if it is marked private, and the space was exported by someone who is not a Site Admin, the email address would not be included in the export).

Email matching is only available in the following Confluence Data Center and Server versions:

- 6.6.14 and any later 6.6 Long Term Support release version
- 6.13.5 and any later 6.13 Long Term Support release version
- 6.15.4 and later

In all other versions, content will be attributed to 'unknown user' if we're unable to match by username.

Restore (import) a space from Confluence Cloud

To restore a small space from Confluence Cloud:

1. In Confluence Cloud, export the space to XML.
2. In Confluence Data Center or Server, go to > **General Configuration**>**Backup and restore**.
3. Select**Restore site or spaces**
4. Under**What do you want to restore?**, select**Spaces**
5. Under**Restore from**, select**Local drive**
6. Under**Select file,**select**Choose file**and browse for your export file.
7. Select**Upload and restore**
8. You will see a confirmation message
9. Select**Upload and restore**once more

To import a large space, the steps are the same, however we recommend dropping the export file into your home directory, rather than uploading it via your browser. SeeRestore a Space or multiple Spacesfor more details.

You may need torestore some permissions to the space if any users or groups aren't present if your destination site.

About unknown users

Any cloud user accounts found in the space export, that are not reconciled with an existing Data Center or Server user, will appear in the **Unsynced from directory**list. They may be listed by email address, or by ID (depending on whether the Cloud user has chosen to keep their email address private).

Permissions and restrictions are respected, so if a space or page is restricted to just one of these users, it will not be visible to other people. An administrator will need to restore permissions after the import is complete.

Restore permissions and restrictions

If the content you import is not attributed to existing users, there will be some work to do to restore the correct permissions to the right people. People may not be able to see the space until you do this.

Restore space admin permissions

The first step is to make sure the space has at least one space administrator. To do this:

1. Go to > **General Configuration**>**Space Permissions**
2. Choose**Recover Permissions** beside the newly imported space
3. Choose**Manage Permissions**. This will take you to the**Permissions** page in that space
4. Grant a user or group**Space Admin** permission for the space and save your changes

If you're a member of theconfluence-administratorssuper group, you can skip steps 2 and 3, and navigate directly to the space.

Restore space permissions

Now that the space has at least one space admin, they can restore any other permissions.

1. Go to the space and choose**Space tools**>**Permissions**from the bottom of the sidebar
2. Grant each user the desired permissions. It can be useful to have the space permissions screen in Confluence Cloud open while you do this.

As long as any groups are named the same in Confluence Cloud and Confluence Data Center or Server you shouldn't need to make any changes to groups. If your groups aren't named the same, you can add any relevant groups at this point.

Restore page restrictions

Pages with view restrictions applied in Confluence Cloud may be associated with unknown users in Confluence Data Center or Server. This means the pages won't be visible.

Space admins can remove individual page restrictions

1. Go to the space and choose**Space tools**>**Permissions**from the bottom of the sidebar
2. Go to the**Restricted Pages**tab. Any pages with view or edit restrictions will be listed.
3. Click the padlock icon beside the page to remove one of the**View restrictions**.
4. If the page is still restricted, use your browser back button and click the padlock beside another View restriction.Repeat this process until enough restrictions have been removed that you can see the page (you'll land on**Page Information**).
5. Choose ••• >**Restrictions**.
6. You can now reinstate the view and edit restrictions. It can be useful to have the Confluence Cloud page open to refer to.

Removing restrictions so that you can see the page may mean that the page becomes temporarily visible to others. If this is a concern you can either apply a temporary view restriction to a parent page, or perhaps remove space permissions until you've finished restoring the right view restrictions.

Understanding the risks

When restoring a space from Confluence Cloud, there's a small risk that content is attributed to the wrong user, which would make any restricted pages visible to the wrong person. This is because the only information we can use to match the user is their email address, which can be changed by the user themselves or by an administrator.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

It's essential that email addresses are associated with the correct user accounts. Content may be attributed to the wrong user account if the email address has been changed maliciously, or accidentally, for example if a username and email combination has been reused, so that a former and current employee share the same username and email address.

We mitigate this risk by only associating content to user accounts that have a unique email address.We don't match accounts with no email address, or where the same email address has been used for multiple user accounts with different usernames, even if they exist in different user directories.

However, if the space you are importing is sensitive, you may want to manually check whether there have been any changes to email addresses recently, before importing a space from Confluence Cloud.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Attachment Storage Configuration

By default Confluence stores attachments in the home directory (e.g. in a file system).

If you have upgraded from an earlier Confluence version you may still be storing attachments in your database or WebDAV. These storage methods are no longer supported.

See below for instructions to migrate to a supported storage method.

On this page:

- Attachment storage methods
  - S3 object storage
  - Database (deprecated)
- Migrating your attachment data
  - From WebDAV or your database to the file system
  - From the file system to S3 object storage

Related pages:

- Working with Confluence Logs

Attachment storage methods

Local file system

By default, Confluence stores attachments in the `attachments` directory within the configured Confluence home folder.

S3 object storage

Starting from Confluence 8.1, you can also store your attachment data on Amazon S3 object storage. We recommend this method if your team has large or increasing data needs and requires the ability to scale efficiently. Learn more about configuring S3 object storage.

Database (deprecated)

Confluence 5.4 and earlier gave administrators the option to store attachments in the database that Confluence was configured to use. This is no longer a supported option.

You should migrate your attachments to the file system before proceeding with Confluence upgrades.

Migrating your attachment data

From WebDAV or your database to the file system

If you are still storing attachments on your database, you can migrate to storing attachments in the file system. When migrating attachments from your database to a file system, the attachments are removed from the database after migration.

> When the migration occurs, all other users will be locked out of the Confluence instance. This is to prevent modification of attachments while the migration occurs. Access will be restored as soon as the migration is complete.

To improve logging during the migration,add the package `com.atlassian.confluence.pages.persistence.dao`with levelDEBUG. See Configuring Loggingfor more information.

To migrate, follow the steps below:

1. Go to > **General Configuration**>**Attachment storage**.
2. Click**Edit**to modify the configuration.
3. Select**Locally in Confluence home directory**.
4. Click**Save**to save the changes.

5. A screen will appear, asking you to confirm your changes. Selecting 'Migrate' will take you to a screen that displays the progress of the migration.
    - If you're already storing attachments in a file system, the **Attachment Storage**option won't appear in the admin console - this is because you're already using the only supported storage method, and don't need to migrate.

Screenshot: migration warning

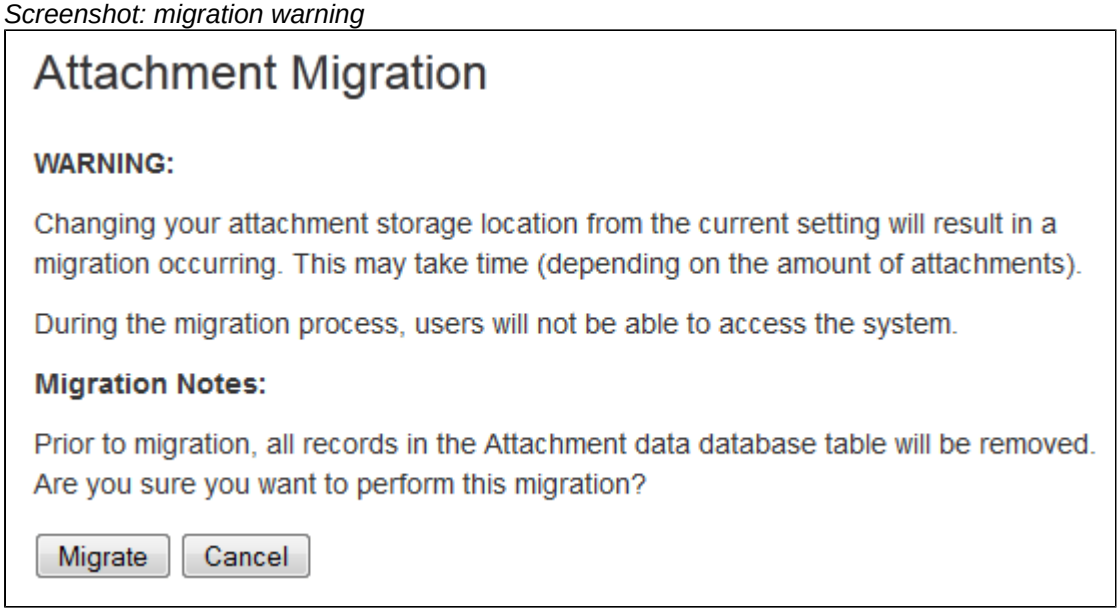

From the file system to S3 object storage

If you have existing attachment data and you want to use Amazon S3, you should migrate attachments to an S3 bucket for Confluence to consume.

To migrate, follow the steps below:

1. Check that you're using Confluence 8.1 or newer.
2. Check that the migration to `v4` hierarchical attachment storage structure is complete. Learn how to do this
3. Create a new Amazon S3 bucket for Confluence (follow step 1 on Configuring S3 object storage).
4. Migrate the `v4` attachment data (`${confluenceHome}/attachments/v4`) from its physical source to the root prefix `confluence/attachments/` in the S3 bucket.
   For example:

File system	`${confluenceHome}/attachments/v4/14/0/327689/327689.1`
S3	`<S3_Bucket>/confluence/attachments/v4/14/0/327689/327689.1`

   The physical location of this data is dependent on your environment. For example, clustered environments typically host this data in a network file system (NFS) as a shared mount. You'll need to consider you setup and the amount of attachment data that needs to be migrated. In general, we recommend using Amazon DataSync for migration. Learn how to do this
5. Wait for the migration to complete.
6. Configure your Confluence node(s) one by one with AWS authentication details and your S3 configuration (follow steps 2 and 3 on Configuring S3 object storage).
    a. Consider putting Confluence into read-only mode to avoid data creation until all node(s) are configured for S3.
    b. After providing the relevant configuration, each node will require a restart.
    c. During this process, if attachments are created on nodes that have yet to be configured for S3, then the attachment data won't be available to those nodes that have been configured for S3.
7. Verify that Confluence is using S3 object storage with the following steps:
    a. Go to > **General Configuration** > **System Information**

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

   b. Next to 'Attachment Storage Type', you'll see 'S3'
   c. Additionally, next to 'Java Runtime Arguments', both the bucket name and region system properties and their respective values will be visible.
8. Re-run the original DataSync job to perform a final sync. This should be done after all the nodes have been configured to ensure all attachment data is migrated.

At this stage, attachment data will be read and written from AWS S3.

> DataSync does not alter or remove the source file system data. So, if you no longer need the attachment data stored on the file system, you'll need to clean this up manually.

Troubleshooting

As the source file system data is not altered or removed by DataSync, Confluence can be reverted back to reading and writing attachment data from the file system. To do this, remove the configuration below from your setenv.shand/or `confluence.cfg.xml`, and restart Confluence:

- `confluence.filestore.attachments.s3.bucket.name`
- `confluence.filestore.attachments.s3.bucket.region`

If you are reverting back to the original file system, any data written to S3 will need to be synced back to the file system manually by the Confluence administrator.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Hierarchical File System Attachment Storage

⚠ Confluence 8.1 introduced a new way to store attachment data in the file system.

When you upgrade to Confluence 8.1 or later, your attachments will be migrated to a new folder structure. More information about this migration task is available on this page.

For information about the previous folder structure, see the version of Hierarchical File System Attachment storagein Confluence 8.0 documentation or earlier.

On this page:

- Directory structure
  - Diagram of the directory
  - Extracted text files
- Migration to version 4
  - Customizations
  - Troubleshooting

Confluence stores attachments, such as files and images, in a file system.

The structure of the attachment storage has been designed to:

1. limit the number of entries at any single level in a directory structure
2. eliminate the need to move attachments between directories when a page is moved to a new location

Directory structure

Attachments in Confluence have a single identifying attribute: the *Content id* of the original version of the attachment.

For example, the original version of the attachment has the *Content id* 12345678 so the attachment file names for versions 1, 2, and 6 will be 12345678.1, 12345678.2, 12345678.6 respectively.

The directory structure consists of 5 levels with the name of each level derived from the following algorithms:

level	Derived From
1 (top)	Always 'v4' indicating the Confluence version 4 storage layout format
2	Calculated as *Content id* modulo 65535, modulo 256
3	Calculated as *Content id* modulo 65535, divided by 256
4	The *Content id* of the attached file
5	These are the file names. They are named the *Content id* and version number of the file, for example: x.1, x.2, x.6.

Themodulo calculation is used to find the remainder after division, for example 800 modulo 250 = 50.

An example in bash showing the path to the first version of the attachment

```
#!/bin/sh

content_id=$1
version=1
domain=`expr $content_id % 65535`
folder_1=`expr $domain % 256`
folder_2=`expr $domain / 256`

echo "v4/$folder_1/$folder_2/$content_id/$content_id.$version"
```

Diagram of the directory

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

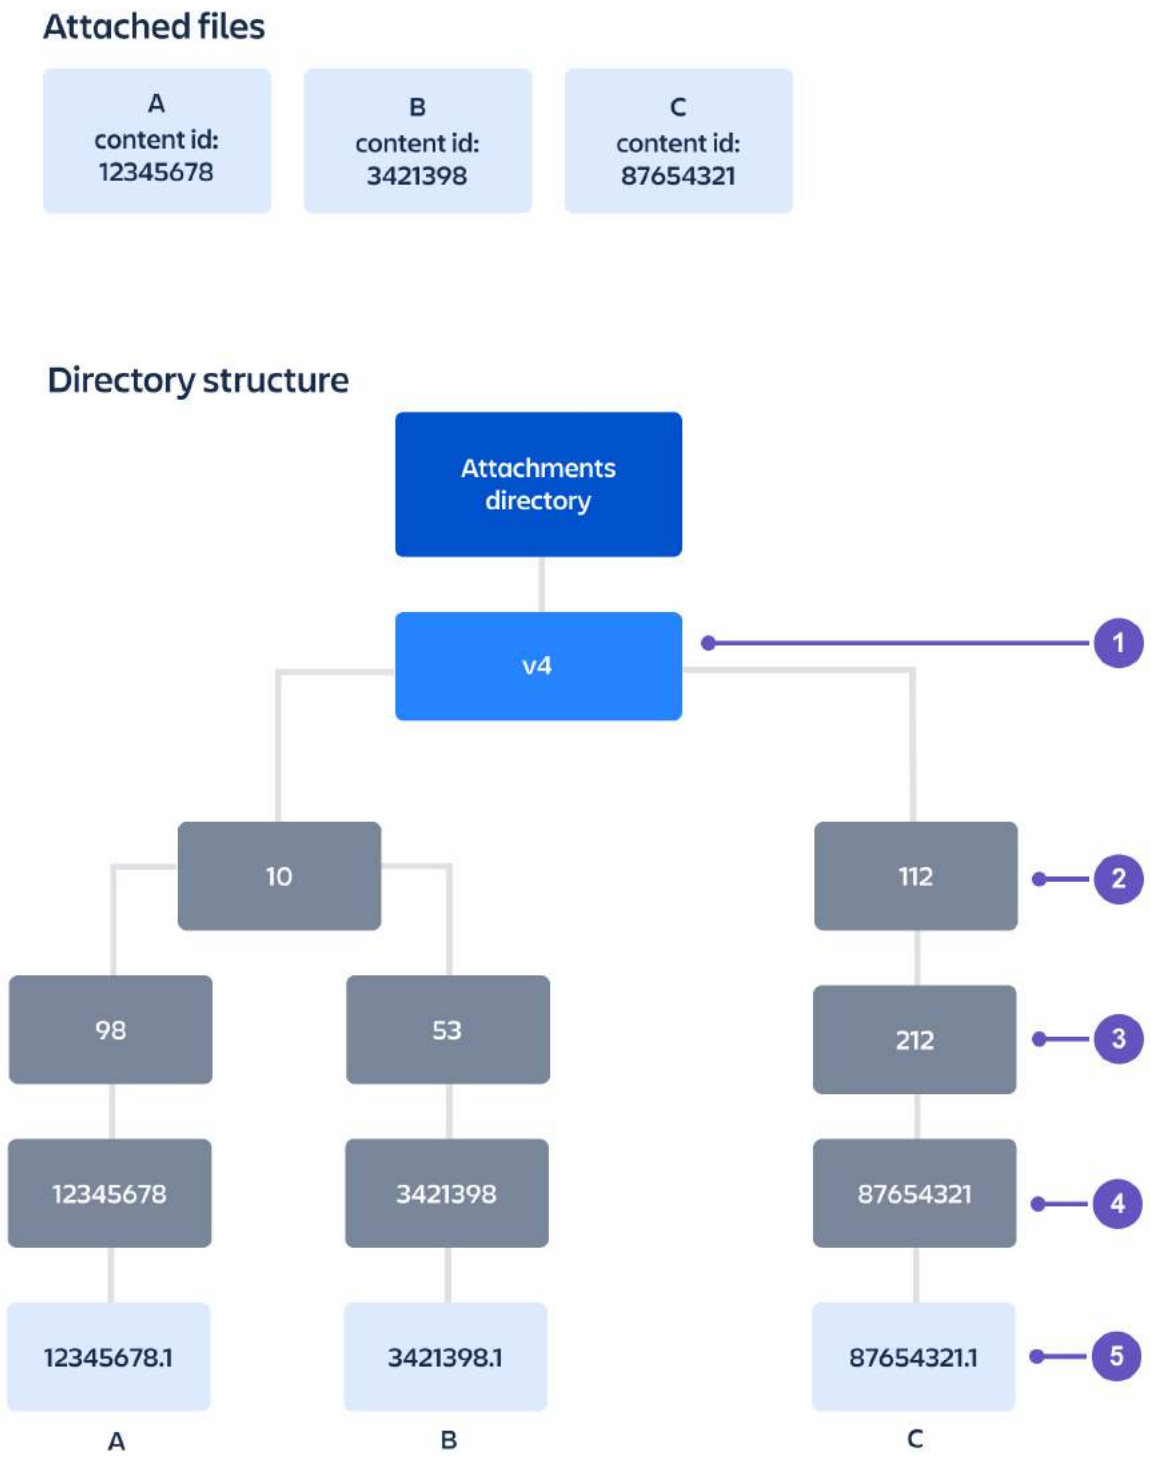

Extracted text files

When a text-based file is uploaded in Confluence (for example Word, PowerPoint, etc), its text is extracted and indexed so that people can search for the content of a file, not just the filename. We store the extracted text so that when that file needs to be reindexed, we don't need to re-extract the content of the file.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

The extracted text file will be named with the content id and version number, for example, `12345.2.extracted_text`, and stored alongside the file versions themselves (within level 5 in the example above). We only keep the extracted text for the latest version, not earlier versions of a file.

Migration to version 4

Confluence 8.1 introduced the v4 layout format for storing attachments. To ensure a smooth transition from the previous version 3 structure, we added an automatic background task to do this migration when you upgrade to Confluence 8.1 or later.

The migration task moves all attachments from the `ver003`directory to the new`v4`directory. Any missing or broken attachments resulting from previous failed page moves that are found in this process will be restored. Duplicate attachments will be saved with the extra extension `.duplicate.X`.

This background task will only run on Confluence startup if a `ver003` directory exists in the attachments folder. During this process, Confluence will work as usual.

When the migration task is finished, a report file `v3-to-v4-report.log` will be available in the attachments directory. A new report is created for each migration run. It contains a list of files with corresponding issues, and the migration status is printed at the bottom for example, completed successfully, completed with warnings, or interrupted. The report file does not print successfully migrated attachments to avoid huge log files.

Customizations

These dark feature customizations are available for your migration. To learn how to configure them, see Configuring System Properties.

System property	Description
-Datlassian.darkfeature.confluence.disable-attachments-ver004=true	Set this property to disable migration to `v4` storage and keep using legacy `ver003`. Any attachments that have already been migrated to `v4` won't be reversed.
-Dconfluence.attachments-ver004-migration-num-of-threads=X	By default, migration uses up to half of the available CPUs on the node. Set the number of CPUs by providing a number instead of X. Reach out to support for assistance with this feature.

Troubleshooting

Issue	Solution
I want to know what logs are available	You can find more about the migration from application logs in `atlassian-confluence.log`. See Working with Confluence Logs There is also a report file `v3-to-v4-report.log` created in the attachments directory.
I want to find out if the migration was successful	The ver003directory will be deleted if migration was successful. You should also review the report log. The final status will be printed at the bottom.
I want to find out if the migration is happening	A report log is created in the attachments directory. Also an entry is printed in a `tlassian-confluence.log` every 50,000 migrated attachments. Search this log file for the string `Attachments migration from ver003 to v4 progressed`

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

I can see multiple report logs in the attachments folder, I want to find out which is the correct one	A report log is created for each migration. If multiple reports are generated, you need to review the latest report and troubleshoot the issue. You may need to delete or move any files that are not real attachments from the ver003 directory. Once all attachments have been migrated successfully, there wont be new report logs in the folder.
The report shows `Duplicate file saved as ...` entries	Due to past issues, some of your attachments may be duplicated in the ver003 directory. See ~~CONFSERVER-62835~~ CLOSED . If the migration task finds duplicate files it will move them to v4 directory and add `.duplicate.X` suffix. After migration, you should review if the duplicates are needed.
The report shows `Failed to migrate. Msg:`	This means a file failed to migrate. Check if permissions allow the move operation, and if the file is not a real attachment you will need to delete or move it from the ver003 directory manually. Then, restart Confluence to trigger another migration task.
I'm running a Confluence DC cluster and can't see any migration logs	The migration runs only on one node, so if you're not in the same node you won't be able to see progress for it in the application log, `atlassian-confluence.log` Every node in the cluster should be able access the report file `v3-to-v4-report.log` in the attachments directory. Check the report for details about your migration.
I want to know if the migration is using up extra disk resources	Migration performs `move` operation only, and on most file systems it will not use any extra disk space. Also, `v4` layout should use less `inodes` than `v3`.
I dont want to do this migration yet	We have a dark feature that allows you to disable migration to v4. See customizations for more info.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Configuring Attachment Size

Related pages:

- Recognized System Properties
- Files

You can limit the size of files that can be uploaded and attached in Confluence.

To configure the maximum file size that can be uploaded:

1. Go to > **General Configuration**.
2. Choose **Edit**.
3. Enter the maximum size next to **Attachment Maximum Size**.
   The default is 100 MB.
4. Choose **Save**.

How attachments are indexed

When a file is uploaded, Confluence will attempt to extract and index its text. This allows people to search for the content of a file, not just the filename. This process is quite memory intensive and can cause out of memory errors when very large files are uploaded. Confluence has a number of safeguards to prevent this happening:

- If the uploaded file is larger than 100 MB, Confluence will not attempt to extract text or index the file contents. Only the filename will be searchable.
- If the uploaded file is one of the following types, Confluence will only extract up to:
  - 1 MB of text from Excel (.xlsx) and PowerPoint (.pptx)
  - 8 MB of text from PDF (.pdf)
  - 10 MB of text from other text files (including .txt, .xml, .html, .rtf etc)
  - 16 MB of text from Word (.docx)

  Note that this is based on the size of the file when it's uncompressed. As .xlsx and .docx files are compressed, text extraction may fail even though the size of the file appears to be under the limit.
- If the text extracted from the file was greater than 1 MB, it will be searchable, but Confluence will not show this text as an excerpt with the search result.

If Confluence stops extracting text, only a portion of the file's content will be searchable.

Confluence will only attempt to extract and index the file once. If it fails, it will not try again.

Some of the values above are configurable via system properties. If you experience out of memory errors when people upload large files, you may want to reduce these limits further, using the following properties:

- `atlassian.indexing.attachment.maxsize`
- `officeconnector.excel.extractor.maxlength`
- `officeconnector.textextract.word.docxmaxsize`
- `atlassian.indexing.contentbody.maxsize`
- `officeconnector.powerpoint.extractor.maxlength`

Configuring S3 object storage

If your team has large or increasing data sets, consider storing your attachments in S3 object storage for greater scalability.

This type of storage is specially designed and optimized for attachment data, unlike traditional file systems.

We currently support **Amazon S3** for attachment object storage. Using this will also mean you get to leverage the out-of-the-box features and benefits of this managed object storage.

On this page:

- Check if object storage is right for you
- Step 1. Create a bucket
- Step 2. Authenticate your Amazon S3 bucket
- Step 3. Connect your S3 bucket with Confluence
- Troubleshooting

Diagram of how object storage works. Attachments uploaded to Confluence are stored in and retrieved from an Amazon S3 bucket.

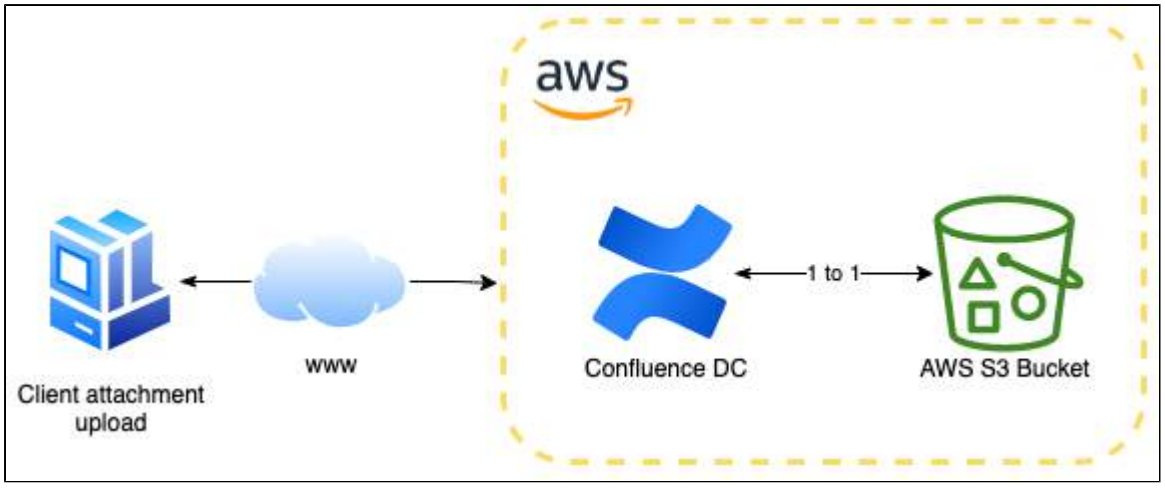

Check if object storage is right for you

We're continuing to build improvements to our object storage solution and we recommend you take some time to read through the requirements and limitations of this version to make sure it's suitable for you.

Requirements

To use Amazon S3 object storage:

- You must be using a Data Center license.
- You should plan to provision Confluence to AWS, or already run Confluence in AWS. This feature isn't supported for on-premise deployments or for any customers not running Confluence in AWS.
- You'll need a dedicated Amazon S3 bucket to hold Confluence attachment data. Learn more about how to create, configure and connect an S3 bucket to Confluence on this page.

For existing customers: You should migrate attachment data to Amazon S3, see Attachment Storage Configuration for instructions on how to do this.

Limitations

- Amazon S3 is currently the only Confluence-supported object storage solution.

- S3 object storage is for attachment data only. You'll still need to use file system storage for other data, for example configuration data.
- There is currently no Atlassian-supported way to migrate attachment data from your file system to Amazon S3, nor from Amazon S3 back to your file system or another storage medium. In general, we'd recommend Amazon DataSync for all migration work.
- Using temporary credentials to authenticate to AWS will require a Confluence restart every time they change. Track this issue at ~~CONFSERVER-81610~~ CLOSED
- There is a known issue when Amazon S3 object storage is configured where performing attachment-related tasks involving more than 50 attachments causes your instance to become temporarily unresponsive or slow. We are actively investigating this bug, and you can track the issue at ~~CONFSERVER-82499~~ CLOSED

Step 1. Create a bucket

Before you can start using Amazon S3 to store your attachments, you'll need an Amazon S3 bucket. Amazon has official guides for how to do this:

- Creating a bucket
- Bucket security
- Bucket restrictions and limitations

Reminder to secure your S3 bucket

Make sure your bucket is correctly secured, and not publicly exposed. You're responsible for your Amazon S3 bucket configuration and security, and Atlassian is unable to provide direct support for issues related to your S3 setup.

Bucket permissions

Make sure you grant Confluence read and write permissions to:

- `s3:ListBucket`
- `s3:PutObject`
- `s3:GetObject`
- `s3:DeleteObject`

Depending on how you authenticate your bucket (see step 2), these permissions can be applied at the bucket level using bucket policies and also via IAM roles for EC2.

Here is an example Identity and Access Management (IAM) policy providing appropriate permissions (based on a least privilege model):

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

```
{
    "Version": "2012-10-17",
    "Id": "PolicyForS3Access",
    "Statement": [
        {
            "Sid": "StatementForS3Access",
            "Effect": "Allow",
            "Principal": {
                "AWS": "arn:aws:iam::123456789012:user/ConfluenceS3"
            },
            "Action": [
                "s3:ListBucket",
                "s3:PutObject",
                "s3:GetObject",
                "s3:DeleteObject"
            ],
            "Resource": [
                "arn:aws:s3:::confluence-attachment-data/*",
                "arn:aws:s3:::confluence-attachment-data"
            ]
        }
    ]
}
```

Supported bucket configurations

Confluence supports these S3 bucket properties and features being enabled via the AWS console or CLI.

⚠ Configuring any property not listed below may result in Confluence not working correctly with your bucket.

Property	Description
Bucket versioning	Keep multiple versions of an object in one bucket by enabling versioning. By default, versioning is disabled for a new S3 bucket. Amazon S3's bucket versioning should be considered complementary to Confluence's way of managing attachment versions. The two approaches are mutually exclusive where one does not affect the other. You should use S3 bucket versioning where you need to preserve, retrieve, and restore every version of every object stored in your bucket, even when purged from Confluence. ⚠ This means when an attachment is purged from Confluence without manually deleting it from the S3 bucket, it will continue to exist in the S3. This may violate certain GDPR compliances if the Confluence admin isn't aware of it. For information about enabling versioning, see Enabling versioning on buckets.
Bucket policies	Control access to the objects stored in the bucket, see Policies and Permissions in Amazon S3.
S3 Intelligent-Tiering	Only those access tiers marked as "automatic" are supported, see S3 Intelligent-Tiering access tiers.

Step 2. Authenticate your Amazon S3 bucket

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Confluence uses the AWS SDK for Java 2.x to communicate with Amazon S3. The SDK will search for credentials in your Confluence environment in this predefined sequence until it can be authenticated:

1. Java system properties
2. Environment variables
3. Web identity token from AWS Security Token Service
4. The shared credentials and `config` files `(~/.aws/credentials)`
5. Amazon ECS container credentials
6. Amazon EC2 instance profile credentials

For information on setting credentials against your environment, Amazon has developer guides on:

- Working with AWS Credentials
- Security Best Practices for Amazon S3

Amazon recommends using IAM roles to authenticate your AWS S3 bucket.

To test your bucket connectivity:

Confirm the authentication mechanism is valid and that the correct permissions are in place using the AWS S3 CLI and the steps below.

1. Create a test file:

```
touch /tmp/test.txt
```

2. Confirm `S3:PutObject` permissions by writing the file to the target bucket:

```
aws s3api put-object --bucket <bucket_name> --key conn-test/test.txt --body /tmp/test.txt
```

3. Confirm `S3:ListBucket` permissions:

```
aws s3api list-objects --bucket <bucket_name> --query 'Contents[].{Key: Key, Size: Size}'
```

4. Confirm `S3:GetObject` permissions:

```
aws s3api get-object --bucket <bucket_name> --key conn-test/test.txt /tmp/test.txt
```

5. Confirm `S3: DeleteObject` permissions:

```
aws s3api delete-object --bucket <bucket_name> --key conn-test/test.txt
```

6. Remove the original test file:

```
rm /tmp/test.txt
```

Step 3. Connect your S3 bucket with Confluence

To connect the Amazon S3 bucket with your Confluence instance:

1. Configure the bucket name and region system properties**:**
   - `confluence.filestore.attachments.s3.bucket.name`
   - `confluence.filestore.attachments.s3.bucket.region`
   - To learn how to do this, see Configuring System Properties.
   - Note: `confluence.cfg.xml` in local home or shared home (if clustering is enabled) will be automatically updated with these properties.
2. Then, start/restart your Confluence instances.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

3. When Confluence starts up, it will check your bucket connectivity, bucket name and region validity, and bucket permissions. If these cant be validated, the startup process will stop and youll receive an error message to tell you why it has failed. See our troubleshooting section below for help with these errors.

To verify that Confluence is using Amazon S3 object storage:

1. Go to > **General Configuration** > **System Information**
2. Next to 'Attachment Storage Type', you'll see 'S3'
3. Additionally, next to 'Java Runtime Arguments', both the bucket name and region system properties and their respective values will be visible.

Note: When using Amazon S3 storage, Confluence ignores the `attachments.dir` property (used for relocating a storage directory). Instead, attachment data is stored in S3 using the root prefix `/confluence/attachments/v4`. In other words, changing the `attachments.dir` property will have no impact on where attachments are stored once Confluence is configured to use Amazon S3.

Troubleshooting

On startup, Confluence will perform a series of health checks to identify any problems. These are listed below with the actions you should take to resolve them.

The main issues will be related to improper S3 configuration, permissions, or authentication.

You can also find more details about the problem by reviewing the health check log at `atlassian-confluence-health-checks.log`. TheWorking with Confluence Logs page explains how to access this and other logs.

Problem

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Missing S3 configuration

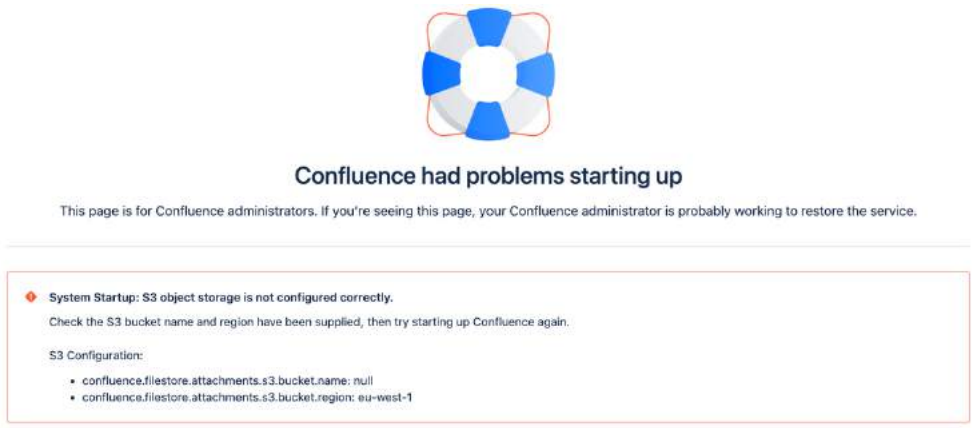

Text version:

⚠System Startup: S3 object storage is not configured correctly.

Check the S3 bucket name and region have been supplied, they try starting up Confluence again.

S3 Configuration:

- confluence.filestore.attachments.s3.bucket.name: null
- confluence.filestore.attachments.s3.bucket.region: eu-west-1

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Bucket region isn't valid

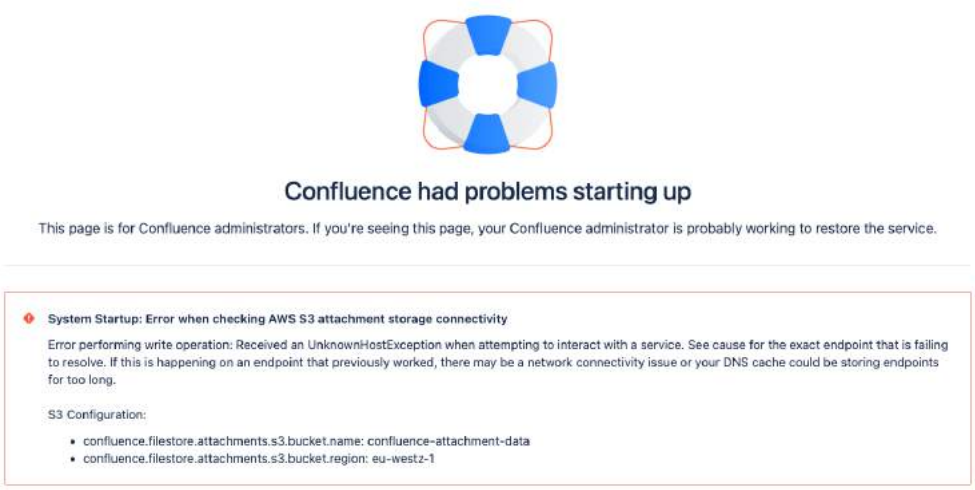

Text version:

⚠️ System Startup: Error when checking AWS S3 attachment storage connectivity

Error performing write operation: Received an UnknownHostException when attempting to interact with a s
be a network connectivity issue or your DNS cache could be storing endpoints for too long.

S3 Configuration:

- confluence.filestore.attachments.s3.bucket.name: confluence-attachment-data
- confluence.filestore.attachments.s3.bucket.region: eu-west-1

Bucket name isn't valid

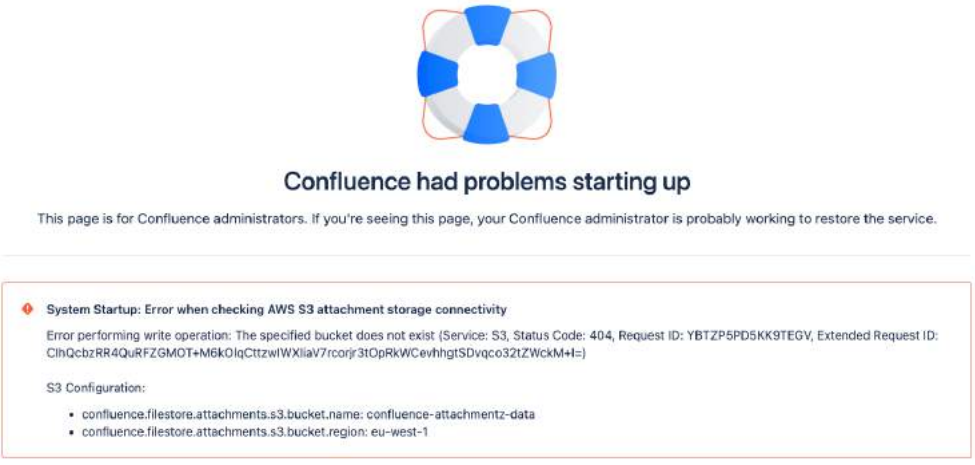

Text version:

⚠️ System Startup: Error when checking AWS S3 attachment storage connectivity

Error performing write operation: The specified bucket does not exist (Service: S3, Status Code: 404, Requ

S3 Configuration:

- confluence.filestore.attachments.s3.bucket.name: confluence-attachment-data
- confluence.filestore.attachments.s3.bucket.region: eu-west-1

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Starting Confluence with bad AWS credentials

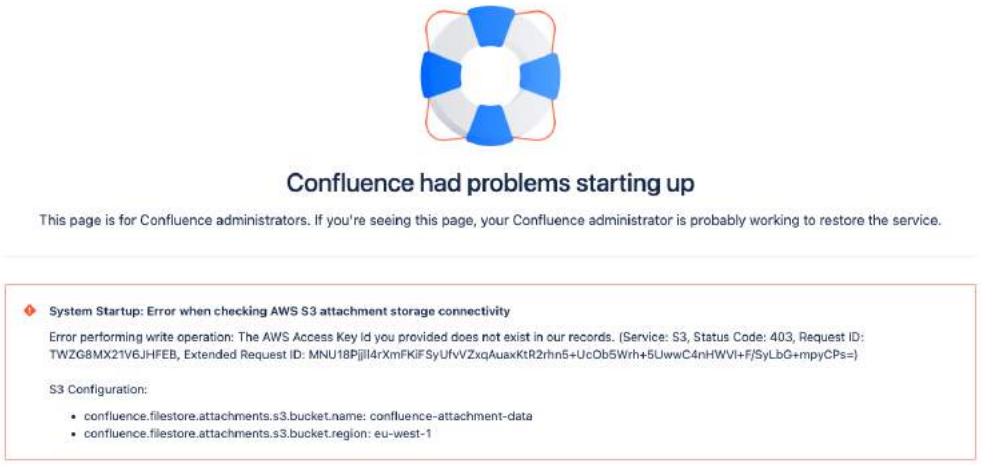

Text version:

⚠️ System Startup: Error when checking AWS S3 attachment storage connectivity

Error performing write operation: The AWS Key Id you provided does not exist in our records. (Service: S3,

S3 Configuration:

- confluence.filestore.attachments.s3.bucket.name: confluence-attachment-data
- confluence.filestore.attachments.s3.bucket.region: eu-west-1

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Bucket has not been configured read permissions (`S3:GetObject`)

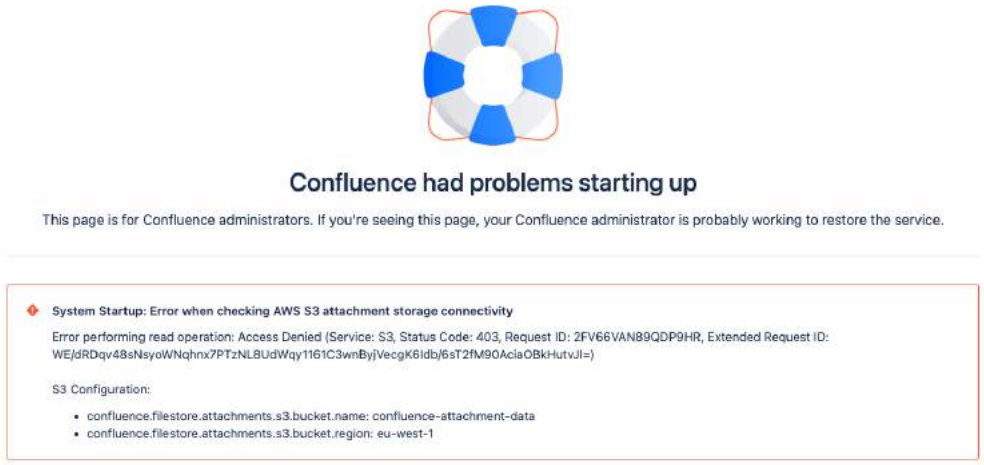

Text version:

⚠️ System Startup: Error when checking AWS S3 attachment storage connectivity

Error performing read operation: Access Denied (Service: S3, Status Code: 403, Request ID: X, Extended

S3 Configuration:

- confluence.filestore.attachments.s3.bucket.name: confluence-attachment-data
- confluence.filestore.attachments.s3.bucket.region: eu-west-1

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Bucket has not been configured write permissions (`S3:PutObject`)

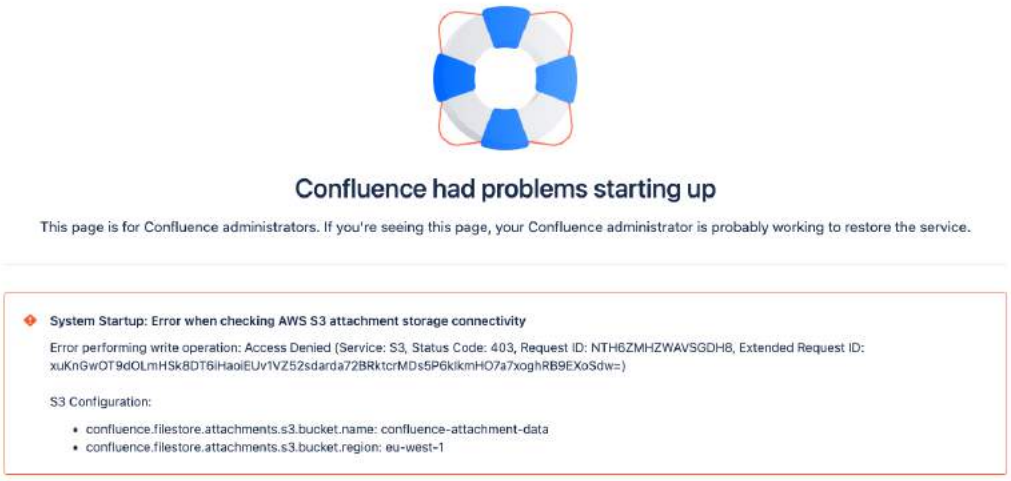

Text version:

⚠ System Startup: Error when checking AWS S3 attachment storage connectivity

Error performing write operation: Access Denied (Service: S3, Status Code: 403, Request ID: X, Extended

S3 Configuration:

- confluence.filestore.attachments.s3.bucket.name: confluence-attachment-data
- confluence.filestore.attachments.s3.bucket.region: eu-west-1

Bucket has not configured delete permissions (`S3:DeleteObject`)

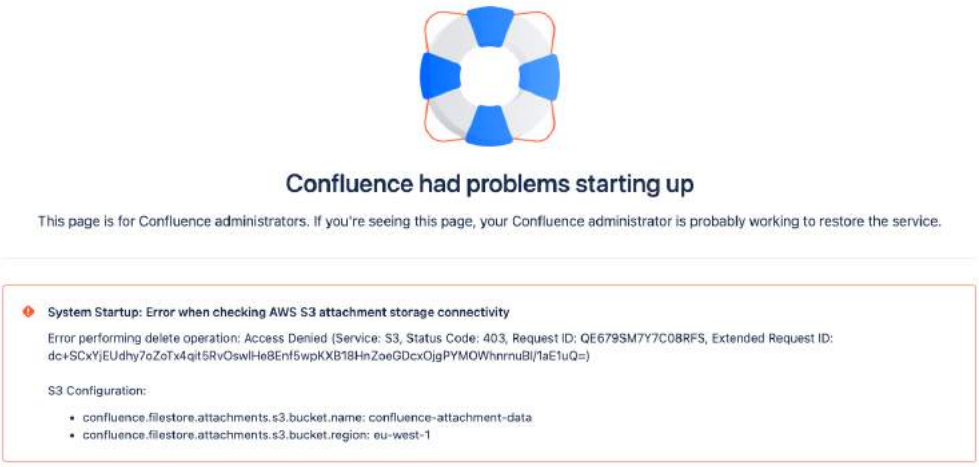

Text version:

⚠ System Startup: Error when checking AWS S3 attachment storage connectivity

Error performing delete operation: Access Denied (Service: S3, Status Code: 403, Request ID: X, Extende

S3 Configuration:

- confluence.filestore.attachments.s3.bucket.name: confluence-attachment-data
- confluence.filestore.attachments.s3.bucket.region: eu-west-1

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Configuring Confluence with a bucket that has no list permissions (`S3:ListBucket`)

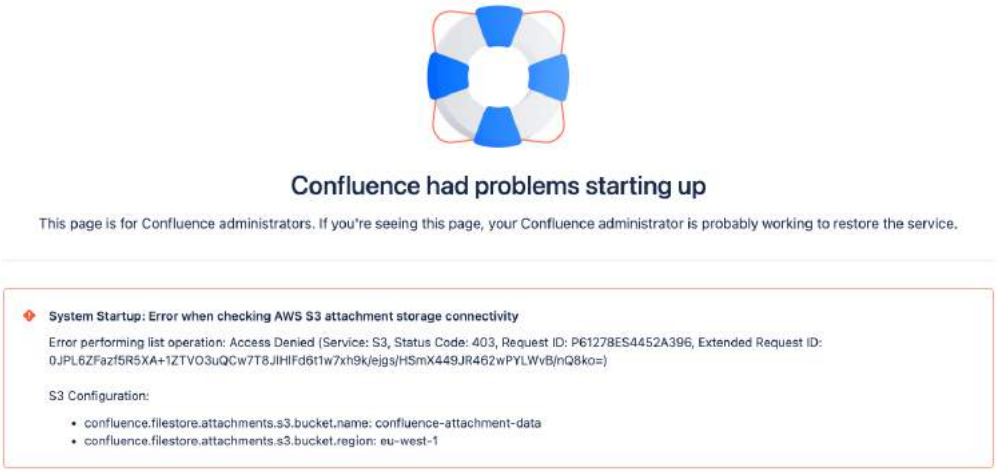

Text version:

⚠️ System Startup: Error when checking AWS S3 attachment storage connectivity

Error performing list operation: Access Denied (Service: S3, Status Code: 403, Request ID: X, Extended R

S3 Configuration:

- confluence.filestore.attachments.s3.bucket.name: confluence-attachment-data
- confluence.filestore.attachments.s3.bucket.region: eu-west-1

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Confluence Data Model

This document provides a diagram of the Confluence schema and a conceptual overview of the data model.

Notes:

- The Hibernate mapping files are the authoritative reference for the Confluence data model. These are the `*.hbm.xml` files which you will find in the main Confluence JAR file (`<CONFLUENCE-INSTALLATION>\confluence\WEB-INF\lib\confluence-x.x.x.jar`).
- The tables, columns and other attributes are likely to change with each major release of Confluence. To find the exact DDL of your Confluence site, please run a query after installation.

On this page:

- Database diagrams
- Database tables and references
- Authentication
- Content
- Clustering
- System information
- Spaces
- Appearance
- Synchrony
- Miscellaneous

Database diagrams

We find that creating your own visualization of the Confluence database can be useful if you want to focus on particular tables or relationships.There are a number of tools you can use to create a visualization. Your own database tool may have options to do this.

View our visualization (excludes some tables, including ActiveObjects tables)

We used DbVisualizer. SeeViewing Table Relationships in the DbVis documentation to find out how it's done.

Database tables and references

Expand the link below to see a table of the primary and foreign keys for each table.

Note that Marketplace apps can also add tables to your database.

Primary Key Table Name	Primary Key Column Name	Foreign Key Table Name	Foreign Key Column Name	Foreign Key Name	Primary Key Name
AUDITRECORD	AUDITRECORDID	AUDIT_AFFECTED_OBJECT	AUDITRECORDID	FK_AFFECTED_OBJECT_RECORD	PRIMARY_KEY_D
AUDITRECORD	AUDITRECORDID	AUDIT_CHANGED_VALUE	AUDITRECORDID	FK_CHANGED_VALUE_RECORD	PRIMARY_KEY_D
CONTENT	CONTENTID	ATTACHMENTDATA	ATTACHMENTID	FKJNH4YVWEN0176QSVH4RPSRY2J	PRIMARY_KEY_6
CONTENT	CONTENTID	BODYCONTENT	CONTENTID	FKMBYIAYESFP1EIQ6GMOL3MK3YL	PRIMARY_KEY_6
CONTENT	CONTENTID	CONFANCESTORS	DESCENDENTID	FKLMHSIPSWOL8IMEQSG906IH62X	PRIMARY_KEY_6
CONTENT	CONTENTID	CONFANCESTORS	ANCESTORID	FKSQB1AF9H7FVQTGY73O8JDCUUE	PRIMARY_KEY_6
CONTENT	CONTENTID	CONTENT	PARENTCOMMENTID	FKAL6O8XWYPD4MDGID9B9NW1Q51	PRIMARY_KEY_6

CONTENT	CONTENTID	CONTENT	PARENTCCID	FKFIYHKA48C7E776QJ90KLBPM9Q	PRIMARY_KEY_6
CONTENT	CONTENTID	CONTENT	PREVVER	FKK6KBB7SUQELOJ82NX7XDCD803	PRIMARY_KEY_6
CONTENT	CONTENTID	CONTENT	PARENTID	FKOXTT893WEUJKYH0IICOXSM37V	PRIMARY_KEY_6
CONTENT	CONTENTID	CONTENT	PAGEID	FKWJYN6091Q3L1GL7BH143MA2A	PRIMARY_KEY_6
CONTENT	CONTENTID	CONTENT_LABEL	CONTENTID	FKI8CVAHSU6D2Y285VTRP4NHC3W	PRIMARY_KEY_6
CONTENT	CONTENTID	CONTENT_PERM_SET	CONTENT_ID	FK2BUUNK1HOR0I3K0M3NT03HW1W	PRIMARY_KEY_6
CONTENT	CONTENTID	CONTENT_RELATION	SOURCECONTENTID	FKE2A00URQYXMYAJ3JOP48UB8QD	PRIMARY_KEY_6
CONTENT	CONTENTID	CONTENT_RELATION	TARGETCONTENTID	FKIPR00838MKLN699CIMD7RG17X	PRIMARY_KEY_6
CONTENT	CONTENTID	CONTENTPROPERTIES	CONTENTID	FK3FLY21XFK13RQH63TXW2T6K2V	PRIMARY_KEY_6
CONTENT	CONTENTID	EXTRNLNKS	CONTENTID	FK5V5LW9X88VM27RVUBSC130NJY	PRIMARY_KEY_6
CONTENT	CONTENTID	IMAGEDETAILS	ATTACHMENTID	FK2301QICIUQ6SC32JAJ8TYSG3S	PRIMARY_KEY_6
CONTENT	CONTENTID	LIKES	CONTENTID	FKBDOIWI70I7O3TC7HPBU4VNLMY	PRIMARY_KEY_6
CONTENT	CONTENTID	LINKS	CONTENTID	FKN8MYCKO8FRERNE7BRH5NR1CSR	PRIMARY_KEY_6
CONTENT	CONTENTID	NOTIFICATIONS	CONTENTID	FK_NOTIFICATIONS_CONTENT	PRIMARY_KEY_6
CONTENT	CONTENTID	SPACES	SPACEDESCID	FK7NDEWMRL3HQCPWC8EYDN9MV8J	PRIMARY_KEY_6
CONTENT	CONTENTID	SPACES	HOMEPAGE	FKJ4CU5838AQCBW57WY7CKT0T7O	PRIMARY_KEY_6
CONTENT	CONTENTID	TRACKBACKLINKS	CONTENTID	FK1TO6OMJL8NHEVCJBVPT3ED7NT	PRIMARY_KEY_6
CONTENT	CONTENTID	USERCONTENT_RELATION	TARGETCONTENTID	FKPWGL85A266IIE5I0ADU8BDBCV	PRIMARY_KEY_6
CONTENT_PERM_SET	ID	CONTENT_PERM	CPS_ID	FKDE5WL1CUR1SE9281GC0DSAWTB	PRIMARY_KEY_BF
CWD_APP_DIR_MAPPING	ID	CWD_APP_DIR_GROUP_MAPPING	APP_DIR_MAPPING_ID	FK_APP_DIR_GROUP_MAPPING	PRIMARY_KEY_2A
CWD_APP_DIR_MAPPING	ID	CWD_APP_DIR_OPERATION	APP_DIR_MAPPING_ID	FK_APP_DIR_MAPPING	PRIMARY_KEY_2A

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

CWD_APPLICATION	ID	CWD_APP_DIR_GROUP_MAPPING	APPLICATION_ID	FK_APP_DIR_GROUP_APP	PRIMARY_KEY_5
CWD_APPLICATION	ID	CWD_APP_DIR_MAPPING	APPLICATION_ID	FKSTEKJ41875RGSW8OTFFRAYHPL	PRIMARY_KEY_5
CWD_APPLICATION	ID	CWD_APPLICATION_ADDRESS	APPLICATION_ID	FK_APPLICATION_ADDRESS	PRIMARY_KEY_5
CWD_APPLICATION	ID	CWD_APPLICATION_ATTRIBUTE	APPLICATION_ID	FK_APPLICATION_ATTRIBUTE	PRIMARY_KEY_5
CWD_DIRECTORY	ID	CWD_APP_DIR_GROUP_MAPPING	DIRECTORY_ID	FK_APP_DIR_GROUP_DIR	PRIMARY_KEY_AA
CWD_DIRECTORY	ID	CWD_APP_DIR_MAPPING	DIRECTORY_ID	FK_APP_DIR_DIR	PRIMARY_KEY_AA
CWD_DIRECTORY	ID	CWD_DIRECTORY_ATTRIBUTE	DIRECTORY_ID	FK_DIRECTORY_ATTRIBUTE	PRIMARY_KEY_AA
CWD_DIRECTORY	ID	CWD_DIRECTORY_OPERATION	DIRECTORY_ID	FK_DIRECTORY_OPERATION	PRIMARY_KEY_AA
CWD_DIRECTORY	ID	CWD_GROUP	DIRECTORY_ID	FK_DIRECTORY_ID	PRIMARY_KEY_AA
CWD_DIRECTORY	ID	CWD_GROUP_ATTRIBUTE	DIRECTORY_ID	FK_GROUP_ATTR_DIR_ID	PRIMARY_KEY_AA
CWD_DIRECTORY	ID	CWD_USER	DIRECTORY_ID	FK_USER_DIR_ID	PRIMARY_KEY_AA
CWD_DIRECTORY	ID	CWD_USER_ATTRIBUTE	DIRECTORY_ID	FK_USER_ATTR_DIR_ID	PRIMARY_KEY_AA
CWD_GROUP	ID	CWD_GROUP_ATTRIBUTE	GROUP_ID	FK_GROUP_ATTR_ID_GROUP_ID	PRIMARY_KEY_C
CWD_GROUP	ID	CWD_MEMBERSHIP	CHILD_GROUP_ID	FK_CHILD_GRP	PRIMARY_KEY_C
CWD_GROUP	ID	CWD_MEMBERSHIP	PARENT_ID	FK_PARENT_GRP	PRIMARY_KEY_C
CWD_USER	ID	CWD_MEMBERSHIP	CHILD_USER_ID	FK_CHILD_USER	PRIMARY_KEY_A3
CWD_USER	ID	CWD_USER_ATTRIBUTE	USER_ID	FK_USER_ATTRIBUTE_ID_USER_ID	PRIMARY_KEY_A3
CWD_USER	ID	CWD_USER_CREDENTIAL_RECORD	USER_ID	FK2RFDH2AP00B8MHOLDSY1B785B	PRIMARY_KEY_A3
EXTERNAL_ENTITIES	ID	EXTERNAL_MEMBERS	EXTENTITYID	FKADLKFU6A03U8F8BS82LM4QLG1	PRIMARY_KEY_6D
GROUPS	ID	EXTERNAL_MEMBERS	GROUPID	FK47K0FUDQNBNSBW0YW44UCSU2R	PRIMARY_KEY_7D
GROUPS	ID	LOCAL_MEMBERS	GROUPID	FKI71UOMCF4F9SESIBDHSMFDBGH	PRIMARY_KEY_7D
KEYSTORE	KEYID	TRUSTEDAPP	PUBLIC_KEY_ID	FKM7N581Y7GROA49TYGAPKMNFIV	PRIMARY_KEY_4D

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

LABEL	LABELID	CONTENT_LABEL	LABELID	FK91V3V5NEMR532QQ4GLA9SJ9TF	PRIMARY_KEY_44
LABEL	LABELID	NOTIFICATIONS	LABELID	FK4TCCRJAMRJVMD2AOGL3HKLPFJ	PRIMARY_KEY_44
OS_GROUP	ID	OS_USER_GROUP	GROUP_ID	FKM2O7638OJNKI05I3U0N5OEPOP	PRIMARY_KEY_DB
OS_USER	ID	OS_USER_GROUP	USER_ID	FK6W5BWO7289K947EE5FWEC30JV	PRIMARY_KEY_E6
PAGETEMPLATES	TEMPLATEID	CONTENT_LABEL	PAGETEMPLATEID	FK28KIFOKT21QD9GES0Q0WV0FB9	PRIMARY_KEY_BC
PAGETEMPLATES	TEMPLATEID	PAGETEMPLATES	PREVVER	FK4WGWY1DQCI8RCWAD4TNQBGLT8	PRIMARY_KEY_BC
SPACES	SPACEID	CONTENT	SPACEID	FKLMWEU06NFT59G7MW1I1MYORYS	PRIMARY_KEY_92
SPACES	SPACEID	NOTIFICATIONS	SPACEID	FKMQE1PHE52XWQC4HK4IB8P9EH6	PRIMARY_KEY_92
SPACES	SPACEID	PAGETEMPLATES	SPACEID	FK18A1D37PVQ2O9HU5X3TPS97MX	PRIMARY_KEY_92
SPACES	SPACEID	SPACEPERMISSIONS	SPACEID	FKBI3X723M8FBGOKO3S84F9ODDL	PRIMARY_KEY_92
TRUSTEDAPP	TRUSTEDAPPID	TRUSTEDAPPRESTRICTION	TRUSTEDAPPID	FKJOFK5643721EFTOW3NJWR73AA	PRIMARY_KEY_DDB
USER_MAPPING	USER_KEY	CONTENT	CREATOR	FK_CONTENT_CREATOR	PRIMARY_KEY_13
USER_MAPPING	USER_KEY	CONTENT	LASTMODIFIER	FK_CONTENT_LASTMODIFIER	PRIMARY_KEY_13
USER_MAPPING	USER_KEY	CONTENT	USERNAME	FK_CONTENT_USERNAME	PRIMARY_KEY_13
USER_MAPPING	USER_KEY	CONTENT_LABEL	OWNER	FK_CONTENT_LABEL_OWNER	PRIMARY_KEY_13
USER_MAPPING	USER_KEY	CONTENT_PERM	CREATOR	FK_CONTENT_PERM_CREATOR	PRIMARY_KEY_13
USER_MAPPING	USER_KEY	CONTENT_PERM	LASTMODIFIER	FK_CONTENT_PERM_LASTMODIFIER	PRIMARY_KEY_13
USER_MAPPING	USER_KEY	CONTENT_PERM	USERNAME	FK_CONTENT_PERM_USERNAME	PRIMARY_KEY_13
USER_MAPPING	USER_KEY	CONTENT_RELATION	CREATOR	FK_C2CRELATION_CREATOR	PRIMARY_KEY_13
USER_MAPPING	USER_KEY	CONTENT_RELATION	LASTMODIFIER	FK_C2CRELATION_LASTMODIFIER	PRIMARY_KEY_13
USER_MAPPING	USER_KEY	EXTRNLNKS	CREATOR	FK_EXTRNLNKS_CREATOR	PRIMARY_KEY_13
USER_MAPPING	USER_KEY	EXTRNLNKS	LASTMODIFIER	FK_EXTRNLNKS_LASTMODIFIER	PRIMARY_KEY_13

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

USER_MAPPING	USER_KEY	FOLLOW_CONNECTIONS	FOLLOWEE	FK_FOLLOW_CONNECTIONS_FOLLOWEE	PRIMARY_KEY_13
USER_MAPPING	USER_KEY	FOLLOW_CONNECTIONS	FOLLOWER	FK_FOLLOW_CONNECTIONS_FOLLOWER	PRIMARY_KEY_13
USER_MAPPING	USER_KEY	LABEL	OWNER	FK_LABEL_OWNER	PRIMARY_KEY_13
USER_MAPPING	USER_KEY	LIKES	USERNAME	FK_LIKES_USERNAME	PRIMARY_KEY_13
USER_MAPPING	USER_KEY	LINKS	CREATOR	FK_LINKS_CREATOR	PRIMARY_KEY_13
USER_MAPPING	USER_KEY	LINKS	LASTMODIFIER	FK_LINKS_LASTMODIFIER	PRIMARY_KEY_13
USER_MAPPING	USER_KEY	LOGININFO	USERNAME	FK_LOGININFO_USERNAME	PRIMARY_KEY_13
USER_MAPPING	USER_KEY	NOTIFICATIONS	CREATOR	FK_NOTIFICATIONS_CREATOR	PRIMARY_KEY_13
USER_MAPPING	USER_KEY	NOTIFICATIONS	LASTMODIFIER	FK_NOTIFICATIONS_LASTMODIFIER	PRIMARY_KEY_13
USER_MAPPING	USER_KEY	NOTIFICATIONS	USERNAME	FK_NOTIFICATIONS_USERNAME	PRIMARY_KEY_13
USER_MAPPING	USER_KEY	PAGETEMPLATES	CREATOR	FK_PAGETEMPLATES_CREATOR	PRIMARY_KEY_13
USER_MAPPING	USER_KEY	PAGETEMPLATES	LASTMODIFIER	FK_PAGETEMPLATES_LASTMODIFIER	PRIMARY_KEY_13
USER_MAPPING	USER_KEY	SPACEPERMISSIONS	CREATOR	FK_SPACEPERMISSIONS_CREATOR	PRIMARY_KEY_13
USER_MAPPING	USER_KEY	SPACEPERMISSIONS	LASTMODIFIER	FK_SPACEPERMISSIONS_LASTMODIFI	PRIMARY_KEY_13
USER_MAPPING	USER_KEY	SPACEPERMISSIONS	PERMUSERNAME	FK_SPACEPERMISSIONS_PERMUSERNA	PRIMARY_KEY_13
USER_MAPPING	USER_KEY	SPACES	CREATOR	FK_SPACES_CREATOR	PRIMARY_KEY_13
USER_MAPPING	USER_KEY	SPACES	LASTMODIFIER	FK_SPACES_LASTMODIFIER	PRIMARY_KEY_13
USER_MAPPING	USER_KEY	TRACKBACKLINKS	CREATOR	FK_TRACKBACKLINKS_CREATOR	PRIMARY_KEY_13
USER_MAPPING	USER_KEY	TRACKBACKLINKS	LASTMODIFIER	FK_TRACKBACKLINKS_LASTMODIFIER	PRIMARY_KEY_13
USER_MAPPING	USER_KEY	USER_RELATION	SOURCEUSER	FK_RELATION_U2USUSER	PRIMARY_KEY_13
USER_MAPPING	USER_KEY	USER_RELATION	TARGETUSER	FK_RELATION_U2UTUSER	PRIMARY_KEY_13

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

USER_MAPPING	USER_KEY	USER_RELATION	CREATOR	FK_U2URELATION_CREATOR	PRIMARY_KEY_13
USER_MAPPING	USER_KEY	USER_RELATION	LASTMODIFIER	FK_U2URELATION_LASTMODIFIER	PRIMARY_KEY_13
USER_MAPPING	USER_KEY	USERCONTENT_RELATION	SOURCEUSER	FK_RELATION_U2CUSER	PRIMARY_KEY_13
USER_MAPPING	USER_KEY	USERCONTENT_RELATION	CREATOR	FK_U2CRELATION_CREATOR	PRIMARY_KEY_13
USER_MAPPING	USER_KEY	USERCONTENT_RELATION	LASTMODIFIER	FK_U2CRELATION_LASTMODIFIER	PRIMARY_KEY_13
USERS	ID	LOCAL_MEMBERS	USERID	FKRCUYOPTNAD1POS41GP1B1F3PI	PRIMARY_KEY_4D4

The following sections describe the principal tables involved in each logical area of Confluence authentication, content, system information, and so on.

Authentication

This section describes the tables involved in user authentication, which is implemented via the Atlassian Crowd framework embedded in Confluence.

Table	Description
`cwd_user`	Information for each user in Confluence.
`cwd_group`	The groups to which users can belong.
`cwd_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.
`cwd_app_dir_group_mapping`	Groups assigned to each application.
`cwd_app_dir_mapping`	Directories assigned to each application.
`cwd_app_dir_operation`	Application-level permissions for adding, modifying and removing users, groups and roles from a directory.
`cwd_application_address`	Remote addresses currently assigned to each application.
`cwd_application_attribute`	Attributes for an application.
`cwd_directory_attribute`	Attributes for a directory.
`cwd_directory_operation`	Permissions for adding, modifying and removing users, groups and roles from a directory.
`cwd_group_attribute`	Attributes for a group.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

cwd_synchronisation_status	Stores the status of the current and most recent synchronization for each user directory
cwd_synchronisation_token	Stores information about the synchronization token used in external user directories' incremental synchronization.
cwd_tombstone	Records removed users, groups, memberships and aliases during incremental synchronization for external user directories.
cwd_user_attribute	Attributes for a user.
cwd_user_credential_record	Hashed passwords for each user.
remembermetoken	Stores 'Remember me' token upon successful user login. Remember me feature is enforced by default when Confluence is clustered.

Content

This section describes the tables involved in storing content. Content is the information that Confluence users are storing and sharing.

Table	Description
AO_*	Active Objects (AO) tables - stores app/plugin data.
attachmentdata	The binary data for attached files. Only populated for Confluence sites created prior to Confluence 5.5, where Confluence was configured to store attachments in the database. Otherwise, attachments are stored in the local file system.
attachments	Only present for Confluence sites created prior to Confluence 5.5, where Confluence was configured to store attachments in the database.
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.
content_relation	Stores interactions between users, content and spaces.
contentproperties	Stores metadata of certain types of contents (including apps), as well as Synchrony caches.
user_mapping	Link betweencwd_userand other tables where a username is needed.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

user_relation, usercontent_relation	Stores interactions between users, content and spaces.

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.
journalentry	The journal service keeps the search index in synch across each Confluence node.
scheduler_clustered_jobs	Stores configurations of scheduled jobs in Confluence.
scheduler_run_details	Records run details of scheduled jobs in Confluence.

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.
diagnosticalerts	The diagnostics tool continuously checks for symptoms or behaviours that we know may contribute to problems in your site. Events are stored for a limited amount of time in this table.
confzdu	Used by Confluence to perform arolling upgrade.
diagnostics_alerts	Stores diagnostics alerts to provide information that admins can use when troubleshooting problems with their site.

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.
spacepermissions	Information about permissions and restrictions for spaces.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Appearance

The following table contains information about the look and feel of your Confluence site.

Table	Description
`decorator`	The custom display templates used to customize Velocity layouts.

Synchrony

The following table contains information about Synchrony, which is used for collaborative editing.

Table	Description
`event`	Stores all events that happen in the editor.
secrets	Used for authenticating Synchrony with Confluence.
snapshots	A cache of events that happen in the editor.

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.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

`denormalised-*`	Several tables used by the faster permissions service to denormalise permissions records.
`audit*`	Entries and configurations for theAudit Log.
`imagedetails`	Used to store metadata for images' attachments.
`logininfo`	Records login details of Confluence users.
`mig_*`	Entries and configurations for Confluence Cloud Migration Assistant app.
`most_used_labels_cache`	Used to store adaptive most used label caches, which was implemented to tackle performance issue with labels.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Finding Unused Spaces or Pages

Sometimes, you want to know what is *not* being used. It's great to know what's getting most attention, but what about stagnant pages, or even entire spaces that are no longer active?

If you have Confluence Data Center, you can view detailed analytics for spaces and pages. See Analytics for more information.

If you have Confluence Server, you'll need to query the database directly to find out this information.

The following query identifies the last date on which content was modified in each space within a single Confluence instance:

```
SELECT spaces.spacename, MAX(content.lastmoddate)
FROM content, spaces
WHERE content.spaceid = spaces.spaceid
GROUP BY spaces.spacename;
```

It returns a list of space names, and the last date and time at which any content was added or changed.

Alternatively, this query identifies spaces where the content hasn't changed since a specified date:

```
SELECT spaces.spacename, spaces.spacekey
FROM content, spaces
WHERE content.spaceid = spaces.spaceid
GROUP BY spaces.spacename, spaces.spacekey
HAVING MAX(content.lastmoddate) < '2006-10-10';
```

The result is a simple list of space names and corresponding space keys.

Similarly, to identify pages where content hasn't changed since a specified date:

```
SELECT c.title,
c.creationdate,
c.lastmoddate
FROM content c
WHERE  c.prevver IS NULL
  AND c.contenttype = 'PAGE'
  AND c.content_status = 'current'
GROUP BY c.title, c.creationdate, c.lastmoddate
HAVING MAX(c.lastmoddate) < '2006-10-10';
```

The result will provide a list of page titles, created date, and last modified date. Change the `lastmoddate` value as per your needs.

Data Import and Export

Confluence administrators and users can import data into Confluence from a number of sources. The permissions required differ, depending on the scope of the import. See Import Content Into Confluence.

You can also export Confluence content to various formats. See Export Content to Word, PDF, HTML and XML.

Related pages:

- Managing Confluence Data
- Confluence Administrator's Guide

Import a Text File

Confluence allows you to import text files from a directory on the Confluence server, and convert them into Confluence pages. Each file is imported as a separate Confluence page with the same name as the file.

Related pages:

- Import Content Into Confluence
- Backup and Restore

- The text file may contain plain text, HTML orConfluence 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 seeSpacesfor information about differences between site spaces and personal spaces.

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**.

Screenshot: Importing text files

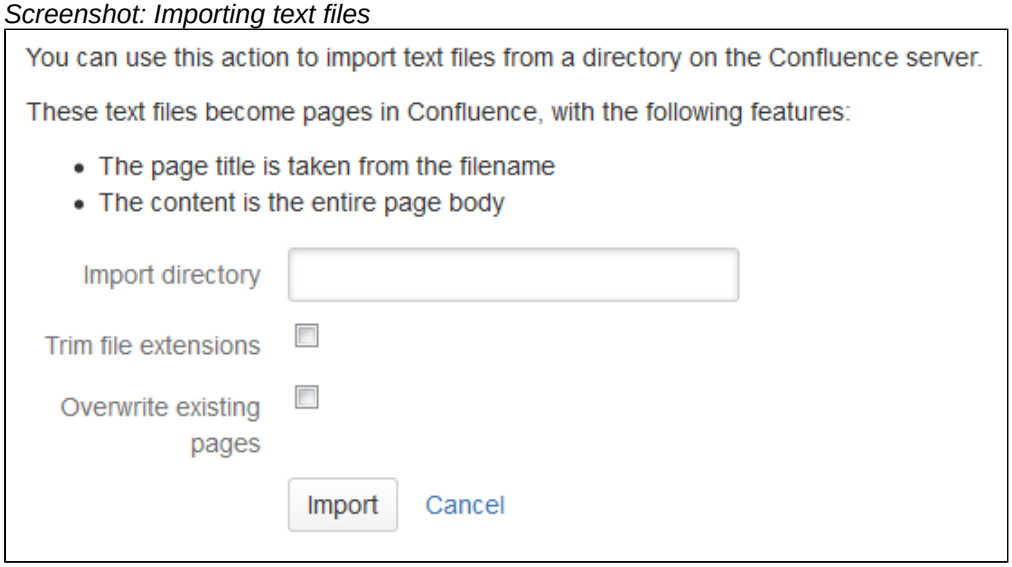

Auditing in Confluence

The audit log allows administrators to look back at changes that have been made in your site. This is useful when you need to troubleshoot a problem or if you need to keep a record of important events, such as changes to global permissions.

The feature is available in both Server and Data Center, however its scope is more extensive in Data Center.

With a Data Center license, Space admins can also view the audit log for their specific space.

On this page:

- Main differences between auditing in Server and Data Center
- View the audit log
- View the space audit log (Data Center)
- Search and filter the audit log
- Edit log settings
- Access the audit log file (Data Center)
- Integrate with external software (Data Center)
- Audit log and migration
- Auditing and the REST API

Main differences between auditing in Server and Data Center

The auditing feature works differently in Confluence Server and Confluence Data Center.

Functionality	Available in Server	Available in Data Center
Coverage areas	Yes (fewer coverage areas than Data Center)	Yes
Selecting coverage areas	No (only **Base** or **Off**)	Yes
Setting database log retention	Yes	Yes
Storing audit logs in two locations	No	Yes
Integrating with 3rd party monitoring tools	No	Yes
Exporting latest 100,000 results	Yes	Yes
Filter by category and summary	No	Yes
Exporting filtered results	Yes	Yes
Space level audit log	No	Yes

View the audit log

To view the global audit log in Confluence:

1. Go to > **General Configuration**
2. Select**Audit log**
3. Select an event to expand it and see details.

Different details will be shown depending on the event itself. These can include:

- **IP address** IP address of the user who performed the action. This is not recorded for system-generated events.
- **Load balancer/proxy IP address** IP address of the load balancer or proxy server that forwarded the request.
- **Node ID** unique ID of the cluster node where the action was performed.
- **Method** depending on how the action was performed, this will be either Browser (end user) or System (system process).

View the space audit log (Data Center)

System admins, Confluence admins and space admins can also access audit logs for a specific space, if they have permission to administer that space.

The space audit log records events related to space permissions and configuration, user actions within the space, and some events related to space security (for example, events related to accessing and granting permissions to restricted pages with a particular space).

To view the audit log for a specific space, go to **Space tools** > **Audit log**.

Search and filter the audit log

You can search the log by keyword, and narrow your results by date, author, and space. If you have a Data Center license you can also filter by category and summary.

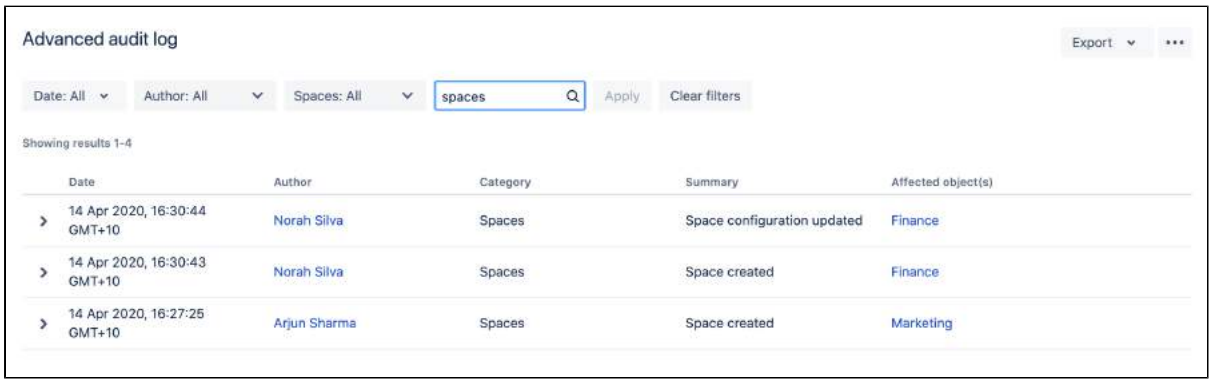

Your query can be up to 100 characters long. To speed up the search, we only search the most recent 1 million events. After this search is performed, you can choose to run a full database search. If you have a large or busy Confluence site, running a full search can take a while.

Can't find a specific event?

Changing coverage level changes the individual events that are logged. If you can't find a specific event, it might be because coverage level was changed, and these events were not logged for a period of time. Check the audit log configuration events to determine if this might be the case.

Edit log settings

In the audit log settings you can decide how long you want to retain the logged events in the database, and the areas from which you want to collect the logs.

Update database retention

The database retention is limited by the retention period, with a maximum of 10 million records.

To update the database retention period:

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

1. Select more options ••• > **Settings**.
2. Enter the period of time. This can be in days, months or years.
3. Select**Save**.

If you choose a long retention period, it can affect the size and performance of your database.Learn more aboutsetting an optimal retention periodfor your Confluence instance.

If you decide to lower the retention period, all the events that exceed the newly set period will be deleted, and disappear from the page. It's a good idea to create a backup before you lower the retention period.

If you migrated from a previous Confluence version, your default retention period is 20 years. If you have a new Confluence installation, its 3 years.

Select events to log

The events that are logged are organized in categories that belong to specific coverage areas.

For example, import and export-related events are logged in the Import/Export category, that belongs to the **L ocal configuration and administration** coverage area. For all coverage areas and events logged in each area, see Audit log events in Confluence.

To adjust the coverage:

1. Go to more options ••• > **Settings**.
2. In the **Coverage level** drop-down, select the level to log the events you need, or **Off** to stop collecting events from a particular area.

Coverage level definitions

Coverage levels reflect the number and frequency of events that are logged. Some coverage levels are only available with a Data Center license.

Coverage level	Definition
Off	Turns off logging for this coverage area.
Base	The lowest level of coverage. Logs only the core events. Base coverage provides a minimum level of insight into your sites activity. If you have a Confluence Server license, this is the only coverage level available.
Advanced (Data Center only)	Logs all the events covered in Base, plus additional events. Advanced coverage provides a more detailed record of your sites activity.
Full (Data Center only)	The highest level of coverage available. Logs all events in Base and Advanced. Depending on your site's activity, setting your coverage level to Full can generate a large volume of events, which can impact your database and disk space.

Export the audit log

You can export up to 100,000 latest or filtered events as a CSV file. If you have more than 100,000 events, only the 100,000 newest events are included in the export.

To export the audit log:

1. Go to **Audit log**,then choose **Export**.
2. Select to export the latest 100,000 or filtered results.
3. Confirm by clicking **Export**.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Space admins can also export from the space level audit log.

Access the audit log file (Data Center)

For Confluence Data Center, each node has its own log, which can be found in the`<local-home>/log/audit`directory. The log is stored as a JSON file.

Confluence creates a new log file every 24 hours, or once the current one reaches 100 MB, whichever occurs first.For more details on log rotation, seeAudit Log Integrations in Confluence.

Change the audit log file retention

You can choose how many audit log files to store in the local home directory on each node. By default we store 100 files. Make sure you've provisioned enough disk space for these files, especially if you have set the logging level to Advanced or Full.

To change the file retention setting:

1. Go to > **General Configuration**> **Audit log**.
2. Select**Settings**.
3. Enter the maximum number of files to be stored and select**Save**.

Once a node reaches the log file retention limit, the oldest one is deleted.If you need to keep these logs, for example for compliance purposes, you may want to manually back up the files in this directory on a regular basis, or send them to a third party logging platform.SeeAudit Log Integrations in Confluence.

Integrate with external software (Data Center)

You can use the log file to integrate with third-party tools such as ELK, Splunk, Sumologic, and Amazon CloudWatch. For more information on integrations, see Audit Log Integrations in Confluence.

Audit log and migration

Migrate database

If you have more that 10 million events stored in your database, and you move to a new database, only the latest 10 million will be migrated, and the remaining data will be removed.

To have access to your older events, you can create a backup before you migrate and access the data in the backup.

Migrate from a previous Confluence version

Migrating audit log records can take a while, depending on the size of the audit log and your database.

Auditing and the REST API

The audit log can also be accessed via the REST API.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Audit Log Events in Confluence

This page outlines the auditing events available in Confluence Server and Data Center, and which events fall into each coverage level.

For more information about how auditing works, see Auditing in Confluence.

On this page

- Definitions
- Global configuration and administration
- User management
- Permissions
- Local configuration and administration
- Security
- End user activity

Definitions

Coverage area

A coverage area is a grouping of events related to a similar theme.

Coverage area	Definition
Global configuration and administration	Logs instance or system admin activity around instance administration or configuration such as platform changes or upgrades to global settings.
User management	Logs activity around users, groups, memberships, and roles, such as adding and removing users and groups.
Permissions	Log activity around local and global permissions and configurations such as changing to anonymous access or updating group permissions.
Local configuration and administration	Logs admin activity around spaces, such as creating or deleting a space.
Security	Logs user actions related to security such as authentication, or granting access to a restricted page.
End user activity (Data Center only)	Logs end user activity on your site, such as user actions on a page (creating, editing, commenting), searching, or viewing pages.

Category

A category is a grouping of related events. Categories can belong to multiple coverage areas.

Category names change over time. You may find your audit log contains some categories not described on this page. These are usually associated with events logged prior to Confluence 7.5.

Coverage level

Coverage levels allow you to control which events are logged. Some levels are only available with a Data Center license.

Coverage level	Definition
Off	Turns off logging for this coverage area.

Base	The lowest level of coverage. Logs only the core events. Base coverage provides a minimum level of insight into your sites activity. If you have a Confluence Server license, this is the only coverage level available.
Advanced (Data Center only)	Logs all the events covered in Base, plus additional events. Advanced coverage provides a more detailed record of your sites activity.
Full (Data Center only)	The highest level of coverage available. Logs all events in Base and Advanced. Depending on your site's activity, setting your coverage level to Full can generate a large volume of events, which can impact your database and disk space.

Global configuration and administration

Category: Global administration

Coverage level	Events logged
Base	Mail server created Mail server deleted Mail server edited Max cache size changed Licence updated Global settings changed Site import Site export Site logo changed Favicon changed Favicon reset to default Custom stylesheet added Custom stylesheet removed Theme changed Custom decorator modified Color scheme modified Color scheme type changed Security configuration updated Global retention rule changed Space retention rule changed Retention rule exemption added Retention rule exemption removed

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Advanced	Mobile apps configuration updated Read-only mode configuration changed Banner configuration changed Collaborative editing mode changed Synchrony restarted Mail queue flushed Mail queue: error queue re-sent Mail queue: error queue deleted Allowlist turned on Allowlist turned off Allowlist URL added Allowlist URL removed Allowlist URL updated Security configuration updated Application navigator link added Application navigator link removed Application navigator link updated Scheduled job enabled Scheduled job disabled Scheduled job edited Scheduled job run manually Application link created Application link edited Application link removed Rate limiting settings updated Rate limiting exemption added Rate limiting exemption removed Rate limiting exemption edited CDN configuration
Full	No events

Category: System

Coverage level	Events logged
Base	No events
Advanced	No events
Full	Monitoring configuration changed

Category: Apps

Coverage level	Events logged
Base	App installed App uninstalled App enabled App disabled App module enabled App module disabled
Advanced	No events
Full	No events

Category: Page templates

Coverage level	Events logged

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Base	Page template updated Page template created Page template deleted
Advanced	No events
Full	No events

User management

Category: Users and groups

Coverage level	Events logged
Base	User created User deleted User renamed User details updated User requested password reset Group created Group deleted User added to group User removed from group User directory created User directory deleted User directory updated
Advanced	User was invited to join site
Full	No events

Permissions

Category: Permissions

Coverage level	Events logged
Base	Space permission removed Space permission added Global permission removed Global permission added
Advanced	No events
Full	No events

Local configuration and administration

Category: Pages and blogs

Coverage level	Events logged
Base	Page hierarchy copy started Page hierarchy delete started
Advanced	No events
Full	No events

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Category: Import / export

Coverage level	Events logged
Base	Space import Space export Space exported to PDF
Advanced	No events
Full	No events

Category: Spaces

Coverage level	Events logged
Base	Space created Space deleted Space archived Space unarchived Space trash emptied Space logo uploaded Space logo enabled Space logo disabled Space logo deleted Unknown update to space logo Space configuration updated
Advanced	No events
Full	No events

Security

Category: Auditing

Coverage level	Events logged
Base	Audit log search performed Audit log exported Audit log configuration updated
Advanced	No events
Full	No events

Category: Authentication

Coverage level	Events logged
Base	No events
Advanced	Secure admin access granted Secure admin access request failed Secure admin access dropped User authorized external application access via OAuth tokens User de-authorized external application access via OAuth token User login failed*
Full	User login successful User logout

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Category: Users and groups

Coverage level	Events logged
Base	No events
Advanced	Forgot password feature triggered Forgot password feature triggered for unknown user
Full	No events

Category: Security

Coverage level	Events logged
Base	No events
Advanced	User tried to access restricted page User requested access to restricted page User requested access to restricted blog Owner of the page authorized access Owner of the blog authorized access
Full	No events

*We only track **User login failed** events if the authentication does not involve a redirect to an external identity provider. If a user tries to log in using SSO and fails, this event will not be logged. Most identity providers track these events in their own audit logs.

End user activity

DATA CENTER ONLY

Category: Pages and blogs

Coverage level	Events logged
Base	No events

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Advanced	Blog post edit restriction added Blog postedit restriction removed Blog postview restriction added Blog postview restriction removed Blog post deleted Blog post restored Blog post purged from trash Blog post version deleted Page view restriction added Page view restriction removed Page edit restriction added Page edit restriction removed Page restored Page deleted Page version deleted Page purged from trash Attachment deleted Attachment version deleted Comment deleted Page moved Blog post moved Page shared Blog post shared
Full	Page created Page edited Blog post created Blog post edited Comment created Comment edited Inline comment created Inline comment edited Attachment downloaded Attachment uploaded Search performed

Note: The **Search performed** event records the search terms entered in search, advanced search, and in search macros (such as Livesearch and Page Tree Search). If you don't want to collect this data you can disable this event using the`audit.log.search.disabled`system property.

Category: Import / export

Coverage level	Events logged
Base	No events
Advanced	Page exported to PDF Blog post exported to PDF Page exported to Word Blog post exported to Word
Full	No events

Category: Data pipeline

Coverage level	Events logged
Base	Full data export cancelled Full data export triggered Unauthorized full data export triggered Full data export failed

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Advanced	No events
Full	No events

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Audit Log Integrations in Confluence

Confluence Data Center writes audit logs to the database anda log file.By itself, the log file saves you the effort of periodically exporting your audit logs from the database for long-term storage. However, the main purpose of the file is to easily integrate ConfluenceData Center to a third-party logging platform.

On this page:

- Event coverage and log retention
- Log file details
- Integrating with logging agents

Event coverage and log retention

TheAudit log settingsmenu controls the coverage of audit logs in both database and log file.However, this menu does not control the log file's retention period.

The log file's retention is ultimately controlled bylog rotation.We use basic log rotation to manage the volume of logs. We automatically archive the audit logfile when:

- the node's time reaches 12:00 midnight, or
- the audit log file reaches100MB.

Once a node reaches the log file retention limit, the oldest one is deleted. By default the limit is 100 log files (the current audit log file + 99 archives). Make sure you allocate enough disk space for these log files on each application node. For the default setting of 100 files, you should allow 10GB.

Log file details

ConfluenceData Center writes audit logs in real time to the home directory. Specifically, these logs are written to the audit logfile. On clusteredConfluenceData Center deployments, each application node will produce its own logfile in itslocalhome directory.

Location

To integrate the audit logfile with a third-party logging platform, you'll need to know its exact location. This may vary, depending on how you configured your home directory. For more information about the local home directory, seeConfluence Home and other important directories).

On a clusteredConfluenceData Center deployment, the audit logfile's directory should be the same on all nodes.

SeeCloudWatch Logs Agent Referencefor more information. If you want to see how we automate this via Ansible, check out our deployment playbooks onhttps://bitbucket.org/atlassian/dc-deployments-automation/src/master/.

File name

The audit log file name uses the following naming convention:

```
YYYYMMDD-XXXXX.audit.log
```

The`XXXXX`portion is a 5-digit number (starting with`00000`) tracking the number of audit log files archived in the same day (`YYYMMDD`). For example, if there are 5 archived log files today (January 1, 2020), then:

- the oldest archived log file is`20200101.00000.audit.log`
- the current audit log file is`20200101.00005.audit.log`

Format

Each audit log is written as a JSON entry to the audit log file.Every line in the filerepresents a single event, allowing you to useregular expressionsto do simple searches if needed.

Integrating with logging agents

Most enterprise environments use a third-party logging platform to aggregate, store, and otherwise manage logs from all hosts. Logging platforms like AWS CloudWatch and Splunkuse*agents*to collect logs from every host in the environment. These agents are installed on each host, collecting local logs and sending them back to a centralized location to be aggregated,analyzed, audited, and/or stored.

If your logging platform uses agents this way, you can configure each node's agent to monitor the audit log file directly. Logging agents from most major platforms (including AWS CloudWatch, Splunk, ELK, and Sumo Logic) are compatible with the audit logfile.

Amazon CloudWatch Agent

We provideQuick Starts for ConfluenceData Centerfor easy deployments on AWS. This Quick Start lets you deploy ConfluenceData Center along with an Amazon CloudWatch instance to monitor it.

To set up Amazon CloudWatch, use the**Enable CloudWatch Integration**parameter's default setting (namely,`Metrics and Logs`).The Quick Start will then configure theAmazon CloudWatch Agentto collect thelogs from each node's audit log files. The agent will send these logs to a separate log group named`confluence-<aws-stack-id>-audit`.

Our Quick Start also sets up a default dashboard to help you read the collected data, including logs from each audit logfile.Refer toWorking With Log Groups and Log Streamsfor related information.

Manual configuration

If needed, you can also manually configure the Amazon CloudWatch agent to collect the audit log files. To do this, set the followingparametersin the Agent Configuration File:

- `file`: set this toto`<local home directory>/log/audit/*`.Don't forget to set the absolute path to thehome directory.
- log_group_nameandlog_stream_name:use these to sendConfluence Data Center's audit logs to a specific log group or stream.

Splunk Universal Forwarder

For Splunk Enterprise or Splunk Cloud, you can use theSplunk Universal Forwarderas your logging agent.This will involve installing the universal forwarder on each application node.

You'll also need to define each node's audit log directoryas one of the forwarder's inputs. This will set the forwarder to send all logs from the audit log directoryto a pre-configuredreceiver.One way to define the forwarder's inputs is through the Splunk CLI. For Linux systems, use the following command on each application node:

```
./splunk add monitor <local home directory>/log/audit/*audit.log
```

Refer to the following links for detailed instructions on configuring the Splunk Universal Forwarder on each node:

- How to forward data to Splunk Enterprise
- How to forward data to Splunk Cloud

Filebeat (for the ELK stack)

Within theELK stack, you can use theFilebeatplugin tocollect logs from each node's audit log files. Each time a log is written to the current audit log file, Filebeat will forward that log to Elasticsearch or Logstash.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

To set this up,install Filebeatfirston each application node. Then, set the audit log filedirectory as aFilebeat input. To do that, add its directory as a`path`inthe`filebeat.inputs`section of each node's`filebeat.yml`configuration file. For example:

```
filebeat.inputs:
- type: log
enabled:true
paths:
- <local home directory>/log/audit/
```

Sumo Logic installed collectors

If you have a Sumo Logic instance, you canuseinstalled collectorsto collect logs from each node's audit log files.To do this,install a collectoron each node first.Then, add`<local home directory>/log/audit/*`as aLocal File Sourceto each node's collector.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Set retention rules to delete unwanted data

This feature is available with a Confluence Data Center license.

When you edit a page or attached file, Confluence stores the previous content so you can restore it if you need to. Over time these historical versions start to add up, increasing the size of your database and attachments directory. It's not uncommon for some pages to have hundreds of historical versions, or for a space to have hundreds of items in the trash.

On this page:

- Retention rule criteria
- How versions are deleted
- Define your retention strategy
- Change the global retention rules
- Add a space exemption
- Remove a space exemption
- Allow space administrators to manage exemptions
- View the retention rules in a space
- Change the retention rules in a space
- Revert back to global retention rules in a space
- Considerations for administrators

Retention rules allow you to automatically delete historical versions of pages and attachments, and purge deleted items from the trash. You can:

- set global rules that will apply to all spaces
- define exemptions for spaces that have special requirements, and need different rules
- allow space administrators to set rules for their spaces.

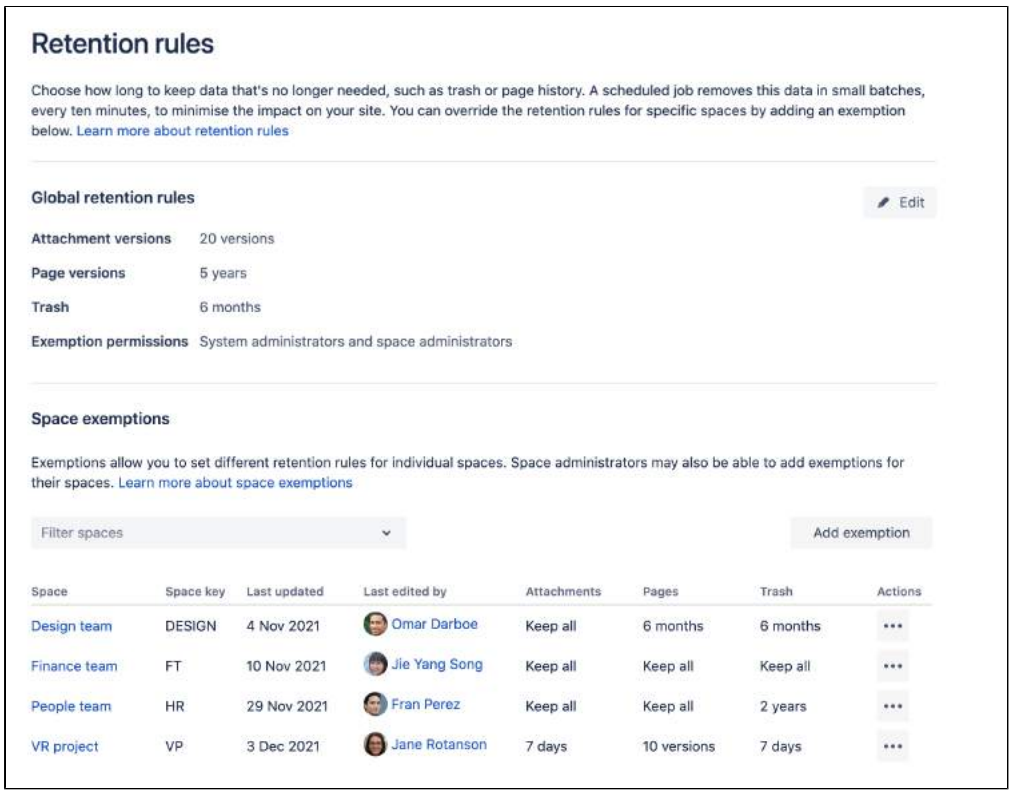

Screenshot showing global retention rules administration screen

Retention rule criteria

You can set a retention rule for:

- historical page versions (pages only, you can't set a retention rule for blog posts)
- historical attachment versions, and
- items in the trash.

The criteria you can use to determine what should be deleted is outlined below.

Rule criteria	Page and Attachment versions	Trash
Keep all (default)	Historical versions will not be automatically deleted.	Items will not be automatically purged from the trash
Keep by age	Keep versions for a specific amount of time. For example you could set it to automatically delete any version older than 2 years.	Keep deleted items in the trash for a specific amount of time. For example you could set it to purge any item deleted more than 3 months ago.
Keep by number	Keep a maximum number of versions. For example you could set it to keep the 5 most recent versions, and automatically delete any earlier versions.	Not applicable for trash.

It's important to note that for page and attachment versions, the latest version is never deleted, only the history. Retention rules never prevent people from creating new versions.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

How versions are deleted

A scheduled job will permanently delete any items that don't the meet retention rules. This "soft" job runs every 10 minutes and deletes items in small batches (of about 3000 items) to ensure there's no performance impact to your site. When you first set a rule,the job may need to run quite a few times before all items that don't meet that rule are deleted.

If you need to delete versions more quickly, you can manually run the "hard" job, which will delete all items that don't meet the rules in one cycle.This can have a performance impact however, so you might want to only run thehard job when Confluence is less busy.

See Scheduled jobs to learn more about these jobs and how to run or disable them.

Each version deleted or item purged from the trash is written to the audit log if the End user activity coverage area is set to **Advanced** or higher.

Define your retention strategy

Before you set any rules, it's important to define your retention strategy. Historical versions and trash that don't meet the retention rule criteria will be permanently deleted, and can't be restored, so it's essential you get it right.

Global rules vs space exemptions

There are two approaches you can take:

- Set your global rule to keep all, and use space exemptions to target individual spaces that can be cleaned up more aggressively.
- Add space exemptions for the spaces where you need to retain history and trash, and set these to keep all, then set a global rule to clean up all remaining spaces.

The order that you set your rules is important - once you set a global rule, it will start deleting almost immediately. Make sure all your exemptions are in place first.

Here's some example scenarios to help you think through your strategy.

Mia is the administrator of a large Confluence site that has been active for about 8 years. Backups and upgrades have become increasingly difficult due to the database and attachment directory now weighing in at over 10 terabytes.

To reduce Confluence's footprint, the Mia:

- Discusses various options with stakeholders, and decides to put in place an aggressive global retention rule.
- Communicates the plan to stakeholders in their organisation, including the date the global rule will be set.
- Makes a note of the current size of the database and attachments directory
- Sets the following global retention rules
  - Page versions - keep by age - 2 years
  - Attachment versions - keep by age - 1 year
  - Trash - keep by age - 6 months

Within a few days, Mia observes a reduction in the size of the database, and a huge 1tb reduction in the size of the attachments directory, both of which contribute to a noticeable decrease in backup time.

Omar administers the Confluence site for an insurance company. Omar would like to clean up the site, but knows that some teams need to keep detailed records of their work, for auditing and compliance purposes.

To clean up unnecessary data Omar:

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

- Discusses various retention rule options with the compliance team. They decide that the best approach is to keep everything by default, but empower individual teams to make decisions about their own spaces.
- Leaves the global retention rules as
  - Page versions - keep all
  - Attachment versions - keep all
  - Trash - keep all
- Sets the exemption permissions to allow system administrators and space administrators to manage retention rules.
- Communicates to team leads that they can set their own rules in Space Tools, and points them to a page prepared by the compliance team which outlines what data must be kept, and for how long.

Within a few days, Omar sees exemptions appearing in the space exemptions list, as several teams add their own retention rules.

Fran administers several different Confluence instances that have sprung up across the organisation as a result of acquisitions over the years. They're in the process of consolidating all their sites, and plan to move some of them to Confluence cloud over the next year or so.

Fran wants this process to be as smooth as possible, and doesn't really want to migrate a huge amount of unnecessary historical data. Fran:

- Uses Analytics to make a short list of spaces that have not been viewed or edited in the last 6 months.
- Informs space owners that much of the version history will soon be removed.
- Leaves the global retention rules as:
  - Page versions - keep all
  - Attachment versions - keep all
  - Trash - keep all
- Adds an exemption for each space that can be cleaned up, and sets a very agressive retention rule:
  - Page versions - keep by number - 5 versions
  - Attachment versions - keep by number - 2 versions
  - Trash - keep by age - 1 month

When the time comes to start migrating spaces to their new cloud instance, Fran prioritises the cleaned up spaces, and finds the migration happens more quickly as there is significantly less data to transfer.

Delegate responsibility to space administrators

You'll need to decide whether to allow space administrators to manage retention rules for their spaces. This allows the administrators most familiar with the content to make decisions about how long to keep historical versions and items in the trash.

See Allow space administrators to manage exemptions.

Change the global retention rules

> ⚠ Check theconsiderations for administratorsbefore changing retention rules.

You need system administrator global permissions to do this.

To change a global retention rule:

1. Go to > **General Configuration** > **Retention rules**.
2. Select **Edit** under **Global retention rules**.
3. Under **Page versions** select**Keep all**, **Keep by number**, or **Keep by age**, and then enter a value (if required).
4. Under **Attachment versions**select**Keep all**, **Keep by number**, or **Keep by age**, and then enter a value (if required).
5. Under **Trash**, select **Keep all**, or **Keep by deleted date**, and then enter a value (if required).
6. Choose whether to allow space administrators to set retention rules for their space.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

7. **Save** your changes.

These rules will apply to all spaces in your site, unless an exemption has been added for a particular space. This includes personal spaces and archived spaces.

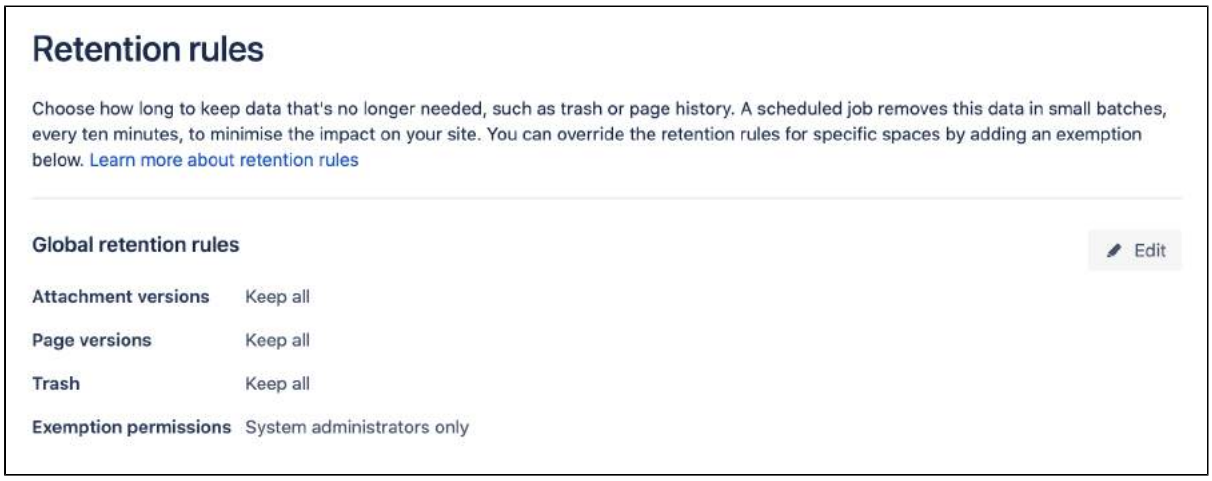

Screenshot showing global retention rules

Add a space exemption

For the situations where you don't want a global retention rule to apply, you can add an exemption to define different rules for a space. For example you may set the global retention rule to delete all page versions older than 2 years, but add an exemption for your Finance and HR spaces, so that all historical versions are kept. Alternatively you could set the global retention rule to keep all, and use exemptions to clean up particular spaces.

You need system administrator global permissions to do this.

To add an exemption for a space:

1. Go to > **General Configuration**> **Retention rules**.
2. Select **Add exemption**.
3. Select a space.
4. Under **Page versions** select**Keep all**, **Keep by number**, or **Keep by age**, and then enter a value (if required).
5. Under **Attachment versions**select**Keep all**, **Keep by number**, or **Keep by age**, and then enter a value (if required).
6. Under **Trash**, select **Keep all**, or **Keep by deleted date**, and then enter a value (if required).
7. Choose whether retention rules for this space can be managed by space administrators.
8. **Save** your changes.

The global retention rules will no longer apply to this space.

Remove a space exemption

To remove an exemption:

1. Go to > **General Configuration**> **Retention rules**.
2. In the space exemptions list, locate your space and choose ... > **Remove**.

The global retention rules will now apply to this space.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

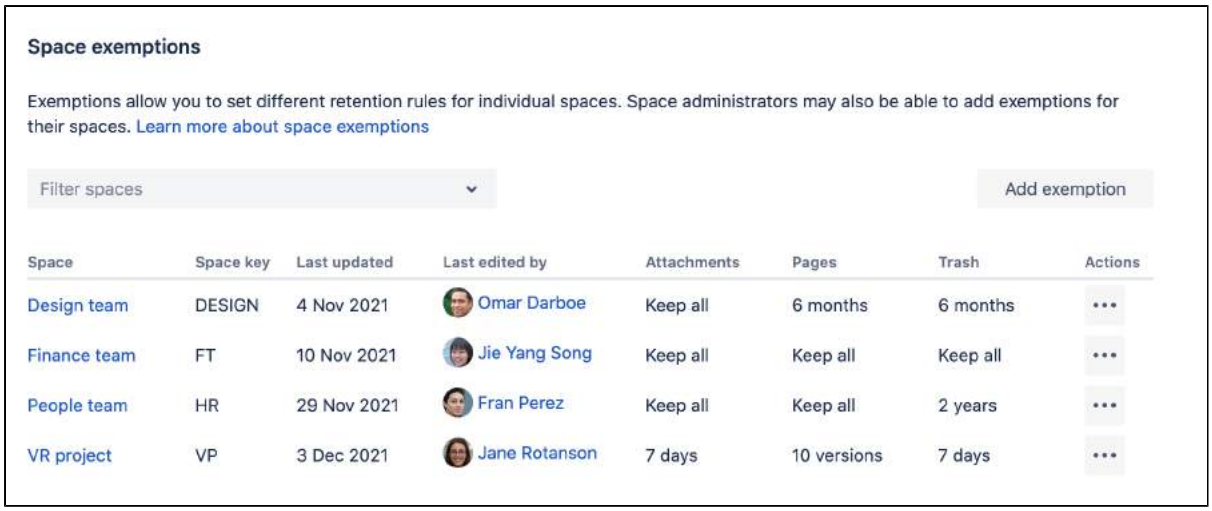

Screenshot showing the global retention rules screen with exemptions listed for particular spaces.

Allow space administrators to manage exemptions

Often system administrators don't have detailed knowledge of the type of content that is stored in each space, so in big sites, it can be useful to allow space administrators to manage the retention rule exemptions for their spaces.

To allow space administrators to manage retention rules:

1. Go to > **General Configuration**> **Retention rules**.
2. Select **Edit** under **Global retention rules**.
3. Under **Exemption permissions**, select **System administrators and space administrators**.
4. **Save** your change.

This applies to all spaces, unless there is an exemption which sets different exemption permissions.

To allow space administrators to manage retention rules for a specific space:

1. Go to > **General Configuration**> **Retention rules**.
2. Add or edit an existing **Exemption**
3. Under **Exemption permissions**, select either **System administrators and space administrators**.
4. **Save** your change.

Space administrators will be able to change the retention rules, as described below.

View the retention rules in a space

Space administrators can always see the retention rules that apply to their space, even if they don't have permission to edit the rules.

To view the retention rules in a particular space:

1. Go to the space and choose**Space tools**>**Content Tools**from the bottom of the sidebar
2. Select the **Retention rules** tab.

Space administrators will only be able to edit the retention rules for their space if a system administrator has allowed space administrators to manage exemptions.

Change the retention rules in a space

You need **space admin** space permissions to do this. The edit button will be disabled if your system administrator has not allowed space administrators to manage exemptions.

To change the retention rules in a particular space:

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

1. Go to the space and choose**Space tools**>**Content Tools**from the bottom of the sidebar
2. Select the **Retention rules** tab
3. Select **Edit**.
4. Select **Use retention rules defined in this space** from the **Exemption permissions** drop down.
5. Under **Page versions** select**Keep all**, **Keep by number**, or **Keep by age**, and then enter a value (if required).
6. Under **Attachment versions**select**Keep all**, **Keep by number**, or **Keep by age**, and then enter a value (if required).
7. Under **Trash**, select **Keep all**, or **Keep by deleted date**, and then enter a value (if required).
8. **Save** your change.

This will add an exemption for this space, or update the existing exemption, if one already exists.

Revert back to global retention rules in a space

To revert back to the global retention rules in a particular space:

1. Go to the space and choose**Space tools**>**Content Tools**from the bottom of the sidebar
2. Select the**Retention rules**tab.
3. Select **Edit**.
4. Select **Inherit global retention rules**from the **Exemption permissions** drop down.
5. **Save** your change.

This will remove the space exemption.

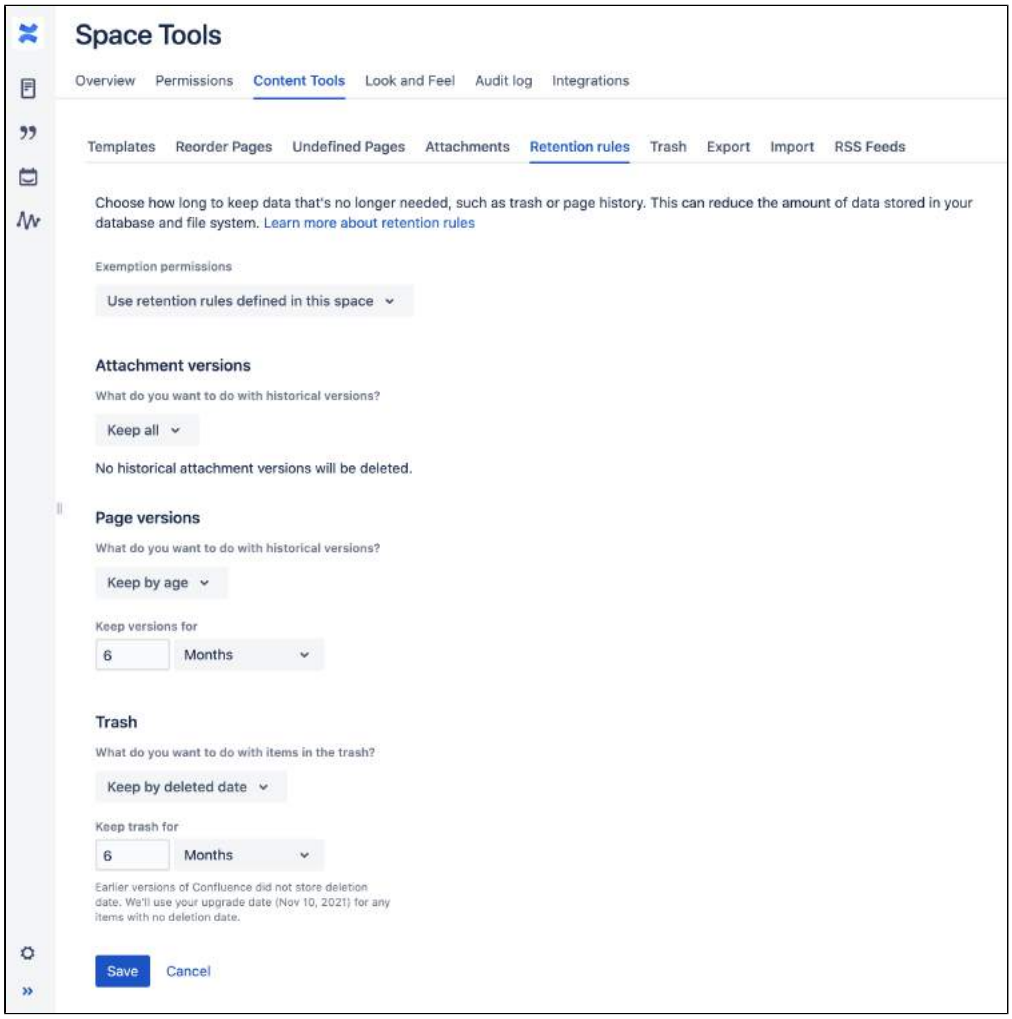

Screenshot showing a space administrator editing retention rules for a space

Considerations for administrators

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

There are a few things you need to consider before changing the retention rules.

Global retention rules apply to all spaces, including archived and personal spaces

The global retention rules apply to all spaces, including archived spaces and personal spaces. If you want to avoid automatically deleting historical versions in archived spaces or personal spaces, you will need to use exemptions either to protect spaces that need to be preserved, or to target spaces that can be cleaned up.

Once you set a global rule, Confluence will start deleting items that don't meet that rule almost immediately, so make sure any exemptions are in place before setting a global rule.

Deleted versions cannot be restored

Deleted versions do not go to the trash, and cannot be restored once deleted.If you need to retain data for regulatory or compliance reasons, you may want to only allow retention rules to be added and updated in the global administration, by system administrators. This is the default.

As always, we recommend you have a robust backup strategy, and a plan for how you will restore data from your backups if required.

Versions can increment very quickly

It's not unusual for many page or file versions to be created within a short space of time.

- A new version of a page is created each someone clicks **Publish** (or **Save** if you have collaborative editing disabled).
- A new version of a file is created each time a file with the same name is attached to the page, or a new version uploaded via the **Upload** button in the preview.

You should factor this in when determining the rule criteria to use, so versions aren't deleted too aggressively . Deleting versions by date rather than number may be more appropriate.

Versions are not renumbered and people who contributed to deleted versions will not be listed in page history

The page history screen will show all remaining versions. Page versions that have been deleted, either by a user, or automatically, will not appear.This is intentional, to indicate that any earlier collaboration (including by the original page creator) is no longer available to view or restore.

In this example version 6 was deleted manually by a user (1), and versions 3 and earlier (2) were deleted because they didn't meet the retention rule criteria.

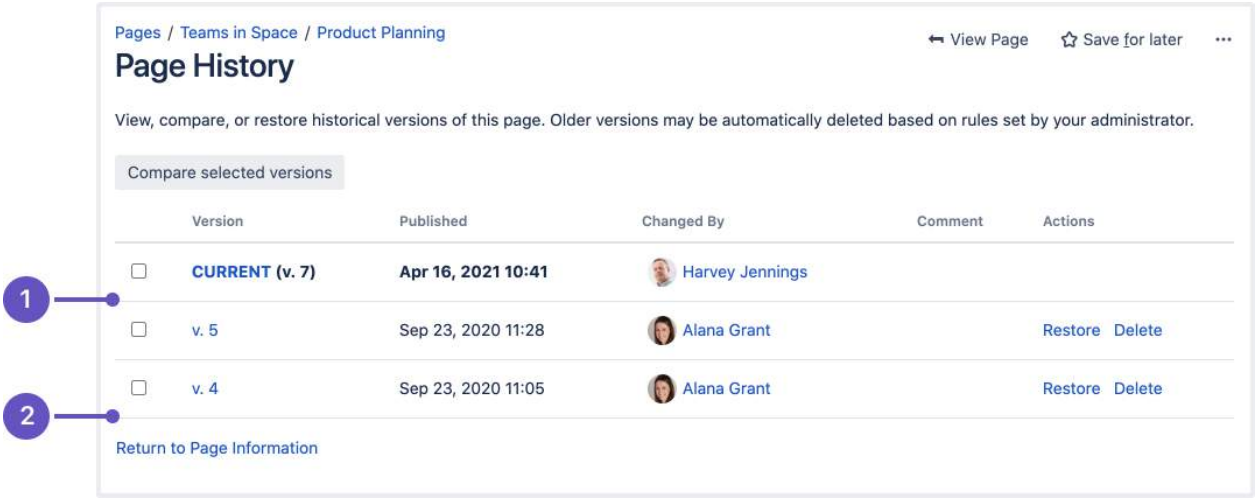

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

The page creator (shown in the byline) will still be the original page creator, and all previous contributors will be recognised if you search by contributor, but we don't list the people who contributed to versions that were deleted on the page history or attachments versions page.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Data pipeline

On this page:

- Requirements
- Considerations
- Access the data pipeline
- Schedule regular exports
- Check the status of an export
- Cancel an export
- Exclude projects from the export
- Configuring the data export
- Use the data pipeline REST API
- Output files
- Troubleshooting issues with data exports

This feature is available with a Confluence Data Center license.

Data pipeline provides an easy way to export data from Jira, Confluence, or Bitbucket, and feed it into your existing data platform (likeTableau orPowerBI). This allows you to:

- generate richer reports and visualizations of site activity
- better understand how your teams are using your application
- make better decisions on optimizing the use of Jira or Confluence in your organization

You can triggera data export in your applications admin console or through the RESTAPI. Data will be exported in CSV format. You can only perform one data export at a time.

For a detailed reference of the exported data's schema, seeData pipeline export schema.

Data pipeline is available inData Centereditions of:

- Jira 8.14 and later
- Confluence 7.12 and later
- Bitbucket 7.13 and later

Requirements

To trigger data exports through the REST API, youll need:

- A valid Confluence Data Center license
- Systems Administratorglobal permissions

Considerations

There are a number of security and performance impacts youll need to consider before getting started.

Security

The export will include all data, including PII (Personally Identifiable Information) and restricted content. This is to provide you with as much data as possible, so you can filter and transform to generate the insights youre after.

If you need tofilter out data based on security and confidentiality, this must be done after the data is exported.

Exported files are saved in your shared home directory, so youll also want to check this is secured appropriately.

Export performance

Exporting data can take a long time in large instances. We intentionally export data at a limited rate to keep any performance impact to your site under a 5% threshold. Its important to note that there is no impact to performance unless an export is in progress.

When scheduling your exports, we recommend that you:

- Limit the amount of data exported using the`fromDate`parameter, as a date further in the past willexport more data, resulting in a longer dataexport.
- Schedule exports during hours of low activity, or on a node with no activity, if you do observe any performance degradation during the export.

	Number	Approximate export duration
Users	100,000	8 minutes
Spaces	15,000	12 minutes
Pages	25 million	12 hours
Comments	15 million	1 hour
Analytics events	20 million	2 hours

The total export time was around 16 hours.

Test performance VS production

The data presented here is based on our own internal testing. The actual duration and impact of dataexport on your own environment will likely differ depending on your infrastructure, configuration, and load.

Our tests were conducted on a single node Data Center instance in AWS:

- EC2 instance type: `c5.4xlarge`
- RDS instance type: `db.m5.4xlarge`

Access the data pipeline

To access the data pipeline go to > **General Configuration**>**Data pipeline**.

Schedule regular exports

The way to get the most value out of the data pipeline is to schedule regular exports. The data pipeline performs a full export every time, so if you have a large site, you may want to only export once a week.

To set the export schedule:

1. From the Data pipeline screen, select **Schedule settings**.
2. Select the **Schedule regular exports** checkbox.
3. Select the date to include data from. Data from before this date wont be included. This is usually set to 12 months or less.
4. Choose how often to repeat the export.
5. Select a time to start the export. You may want to schedule the export to happen outside working hours.
6. Select the **Schema version** to use (if more than one schema is available).
7. **Save** your schedule.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

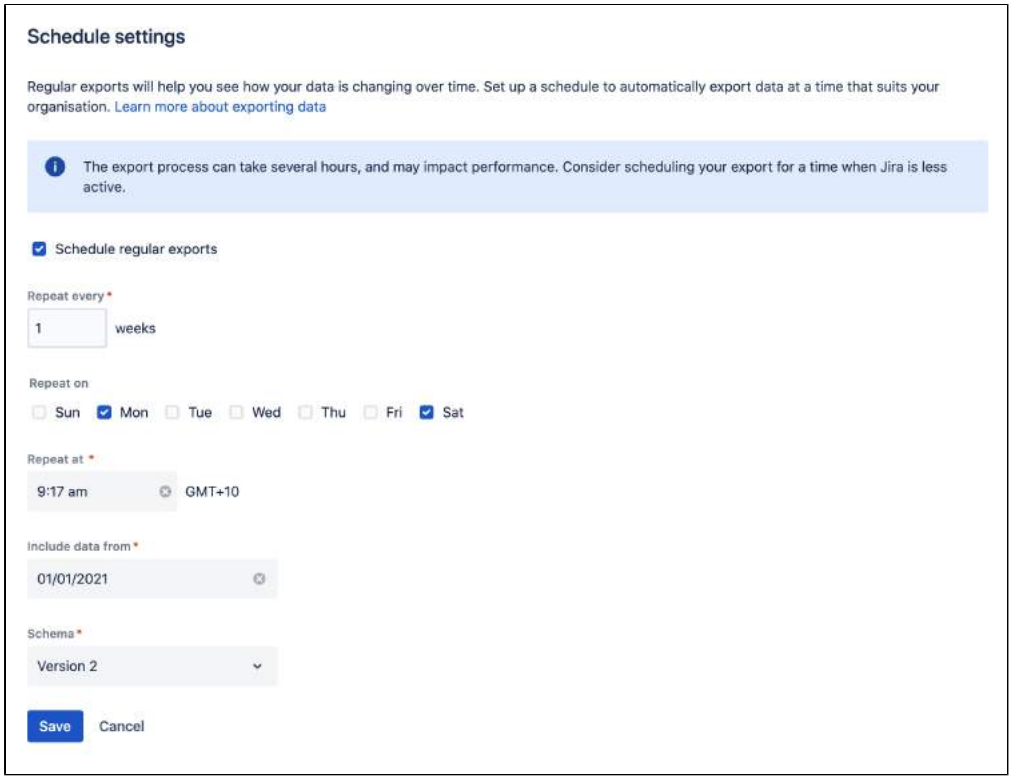

Timezones and recurring exports

We use your server timezone to schedule exports (or system timezone if youve overridden the server time in the application). The export schedule isnt updated if you change your timezone. If you do need to change the timezone, youll need to edit the schedule and re-enter the export time.

You can schedule exports to happen as often as you need. If you choose to export on multiple days, the first export will occur on the nearest day after you save the schedule. Using the example in the screenshot above, if you set up your schedule on Thursday, the first export would occur on Saturday, and the second export on Monday. We dont wait for the start of the week.

Export schema

The export schema defines the structure of the export. We version the schema so that you know your export will have the same structure as previous exports. This helps you avoid problems if youve built dashboards or reports based on this data.

We only introduce new schema versions for breaking changes, such as removing a field, or if the way the data is structured changes. New fields are simply added to the latest schema version.

Older schema versions will be marked as deprecated, and may be removed in future versions. You can still export using these versions, just be aware we wont update them with any new fields.

Check the status of an export

You can check the status of an export and view when your last export ran from the data pipeline screen.

The **Export details** table will show the most recent exports, and the current status.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

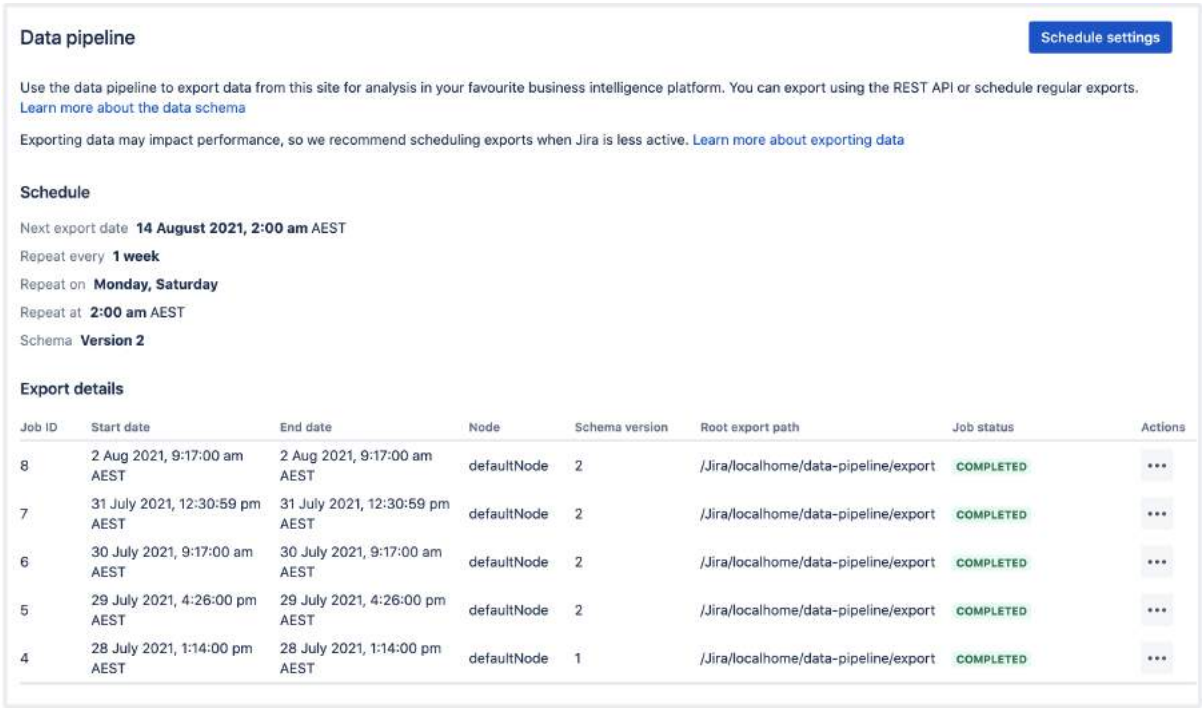

Select •••>**View details** to see the full details of the export in JSON format. Details include the export parameters, status, and any errors returned if the export failed.

For help resolving failed or cancelled exports, seeData pipeline troubleshooting.

Cancel an export

To cancel an export while it is in progress:

- Go to the **Data pipeline** screen.
- Select••• next to the export, and choose **Cancel** export.
- Confirm you want to cancel the export.

It can take a few minutes for the processes to be terminated. Any files already written will remain in the export directory. You can delete these files if you dont need them.

Exclude projects from the export

You can exclude spaces from the export by adding them to an opt-out list. This is useful if you dont need to report on that particular space, or if it contains sensitive content that youd prefer not to export.

To add spaces to the opt-out list, make a`POST`request to`<base-url>/rest/datapipeline/1.0/config/optout` and pass the space keys as follows.

```
{
 "type": "SPACE",
 "keys": ["HR","TEST"]
}
```

These spaces will be excluded from all future exports.

For full details, including how to remove spaces from the opt-out list, refer to theData pipeline REST API reference.

Automatic data export cancellations

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

If you shut down a node running a data export, the export will be cancelled. However, if the JVM is not notified after a crash or hardware-level failure,the export process may get locked. This means you'll need to manually mark the export as cancelled (through the UI, or via the REST API by making a`DELETE`request). This releases the process lock, allowing you to perform another data export.

Configuring the dataexport

You can configure the format of theexport data using the following system properties.

Default value	Description
plugin.data.pipeline.embedded.line.break.preserve	
`false`	Specifies whether embedded line breaks should be preserved in the output files. Line breaks can be problematic for some tools such as Hadoop. This property is set to`False`by default, which means that line breaks are escaped.
plugin.data.pipeline.embedded.line.break.escape.char	
`\\n`	Escaping character for embedded line breaks. By default, we'll print`\n`for every embedded line break.
plugin.data.pipeline.minimum.usable.disk.space.after.export	
5GB	To prevent you from running out of disk space, the data pipeline will check before and during an export that there is at least 5GB free disk space. Set this property, in gigabytes, to increase or decrease the limit. To disable this check, set this property to `-1`(not recommended).

Use the data pipeline REST API

You can use the data pipeline REST API to export data.

To start a data pipeline export, make a POST request to `<base-url>/rest/datapipeline/latest/export`.

Here is an example request, using cURL and a personal access token for authentication:

```
curl -H "Authorization:Bearer ABCD1234" -H "X-Atlassian-Token: no-check"
-X POST https://myexamplesite.com/rest/datapipeline/latest/
export?fromDate=2020-10-22T01:30:11Z
```

You can also use the API to check the status, change the export location, and schedule or cancel an export.

For full details, refer to theData pipeline REST API reference.

Output files

Each time you perform a data export, we assign a numerical job ID to the task (starting with1for your first ever data export). This job ID is used in the file name, and location of the files containing your exported data.

Location of exported files

Exported data is saved as separate CSV files.The files are saved to the following directory:

- `<shared-home>/data-pipeline/export/<job-id>` if you run Confluence in a cluster
- `<local-home>/data-pipeline/export/<job-id>` you are using non-clustered Confluence

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Within the <job-id> directory you will see the following files:

- users_job<job_id>_<schema_version>_<timestamp>.csv
- spaces_job<job_id>_<schema_version>_<timestamp>.csv
- pages_job<job_id>_<schema_version>_<timestamp>.csv
- comments_job<job_id>_<schema_version>_<timestamp>.csv
- analytics_events_job<job_id>_<schema_version>_<timestamp>.csv

To load and transform the data in these files, you'll need to understand the schema. SeeData pipeline export schema.

Set a custom export path

By default, the data pipeline exports the files to the home directory, but you can use the REST API to set a custom export path.

To change the root export path, make a PUT request to <base-url>/rest/datapipeline/1.0/config/export-path.

In the body of the request pass the absolute path to your preferred directory.

For full details, including how to revert back to the default path, refer to theData pipeline REST API reference.

Sample Spark and Hadoop import configurations

If you have an existing Spark or Hadoop instance, use the following references to configure how to import your data for further transformation.

```
%python
# File location
file_location = "/FileStore/**/export_2020_09_24T03_32_18Z.csv"

# Automatically set data type for columns
infer_schema = "true"
# Skip first row as it's a header
first_row_is_header = "true"
# Ignore multiline within double quotes
multiline_support = "true"

# The applied options are for CSV files. For other file types, these will be ignored. Note escape &
quote options for RFC-4801 compliant files
df = spark.read.format("csv") \
  .option("inferSchema", infer_schema) \
  .option("header", first_row_is_header) \
  .option("multiLine", multiline_support) \
  .option("quote", "\"") \
  .option("escape", "\"") \
  .option("encoding", "UTF-8").load(file_location)

display(df)
```


Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

```
CREATE EXTERNAL TABLE IF NOT EXISTS some_db.datapipeline_export (
  `page_id` string,
  `instance_url` string,
  `space_key` string,
  `page_url` string,
  `page_type` string,
  `page_title` string,
  `page_status` string,
  `page_content` string,
  `page_parent_id` string,
  `labels` string,
  `page_version` string,
  `creator_id` string,
  `last_modifier_id` string,
  `created_date` string,
  `updated_date` string,
  `last_update_description` string
)
ROW FORMAT SERDE 'org.apache.hadoop.hive.serde2.OpenCSVSerde'
WITH SERDEPROPERTIES (
  "escapeChar" = "\\",
  'quoteChar' = '"',
  'separatorChar' = ','
) LOCATION 's3://my-data-pipeline-bucket/test-exports/'
TBLPROPERTIES ('has_encrypted_data'='false');
```

Troubleshooting issues with data exports

Exports can fail for a number of reasons, for example if your search index isnt up to date. For guidance on common failures, and how to resolve them, seeData pipeline troubleshooting in our knowledge base.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Data pipeline export schema

This page describes the structure and data schema of the Confluence data export files.

To learn more about how the set up and configure your data pipeline, seeData pipeline.

On this page:

- Output file format and structure
- Users file
- Spaces file
- Pages file
- Comments file
- Analytics events file

Output file format and structure

The output files are written inCSVformat and areRFC4180compliant. They have the following characteristics:

- Each file has a header. This includes files from exports that resulted in no data.
- New lines are separated by CRLF characters `\r\n`.
- Fields containing line breaks (CRLF), double quotes, and commas are enclosed in double quote.
- If double-quotes are present inside fields, then a double-quote appearing inside a field are escaped by preceding it with another double quote. For example:`"aaa"`, `"b""bb"`, `"ccc"`.
- Fields with no data (null values) are represented in the CSV export by two consecutive delimiters(as in,`,,`).
- Embedded break lines are escaped by default and printed asn.

Fields are available in all schema versions, unless specifically noted below.

Users file

Field	Description
`instance_url`	**Type**: URL **Description**: Base URL of the current instance. **Example:** `https://yoursitename.com`
`user_id`	**Type**: String **Description**: ID of the user **Example**: `ff8080817572401e01757240b3520000`
`user_name`	**Type**: String **Description**: User name of the user. **Example**: `jsmith`
`user_fullname`	**Type**: String **Description**: Full name of the user. **Example**: `John Smith`
`user_email`	**Type**: Email **Description**: Email address of the user **Example**: `jsmith@example.com`

Spaces file

Field	Description
space_key	**Type**: String **Description**: Unique identifier that forms part of the URL for thatspace **Example:** `AMF`
instance_url	**Type**: String **Description**: Base URL of the current instance **Example:** `https://example.com`
space_url	**Type**: URL **Description**: The space URL **Example:** `https://example.com/display/SPACEKEY`
homepage_url	**Type**: URL **Description**: The space's home page URL **Example:** `https://example/display/SPACEKEY/Page+name`
space_name	**Type**: String **Description**: Title of the space **Example:** `Design Team Space`
space_type	**Type**: String **Description**: Whether the space is a global or personal space **Example:** `global`
space_status	**Type**: String **Description**: Whether the status of the space is current or archived **Example:** `CURRENT`
creator_id	**Type**: User **Description**: ID of the user who created the space **Example:** `ff8080817572401e01757240b3520000`
last_modifier_id	**Type**: User **Description**: ID of the user who last modified the space **Example:** `ff8080817572401e01757240b3520000`
created_date	**Type**: Time **Description**: Space creation timestamp **Example:** `2021-02-26T04:14:38Z`

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

updated_date	**Type**: Time **Description**: Last modification timestamp **Example:** `2021-02-26T04:14:38Z`

Pages file

Field	Description
page_id	**Type**: Number **Description**: Unique ID of the page
instance_url	**Type**: String **Description**: Base URL of the current instance **Example:** `https://example.com`
space_key	**Type**: String **Description**: Space key of the space the page exists in
page_url	**Type**: String **Description**: URL of the page **Example:** `https:/example/display/SPACEKEY/Page+name`
page_type	**Type**: String **Description**: Whether the entity is a page or a blog post **Example:** `page`
page_title	**Type**: String **Description**: Title of the page
page_status	**Type**: String **Description**: Status of the page (the only value is current, this does not indicate that a page is in a space that has been archived)
page_content	**Type**: String **Description**: Content of the page in Confluence storage format (limited to 10,000 characters) **Example**: `<ac:layout><ac:layout-section ac:type="two_equal"><ac:layout-cell> <p>This is sample content in a layout</p></ac:layout-cell><ac:layout-cell> <p>With two columns</p></ac:layout-cell></ac:layout-section></ac:layout>`
page_parent_id	**Type**: Number **Description**: ID of the current page's direct parent

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

labels	**Type**: String **Description**: Comma separated list of labels of the page **Example:** `["personal", "expense]`
page_version	**Type**: String **Description**: Version number of the latest version page **Example**: 3
creator_id	**Type**: Number **Description**: ID of the user who created the page **Example:** `ff8080817572401e01757240b3520000`
last_modifier_id	**Type**: User **Description**: ID of the user who last updated the page **Example:** `ff8080817572401e01757240b3520000`
created_date	**Type**: Time **Description**: Creation timestamp **Example:** `2021-02-26T04:14:38Z`
updated_date	**Type**: Time **Description**: Last modification timestamp **Example:** `2021-02-26T04:14:38Z`
last_update_description	**Type**: String **Description**: Version comment entered when the page was last updated (limited to 2,000 characters)

Comments file

Field	Description
comment_id	**Type**: Number **Description**: Unique ID of the comment
instance_url	**Type**: String **Description**: Base URL of the current instance **Example:** `https://example.com`
comment_url	**Type**: String **Description**: Full URL of the comment

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

page_id	**Type**: Number **Description**: Unique ID of the page which contains the comment
parent_comment_id	**Type**: Number **Description**: If the comment is a reply, this is the ID of the parent comment (empty for top level comments)
comment_content	**Type**: String **Description**: Content of the comment in Confluence storage format (limited to 2,000 characters) **Example**: `<p>Sample comment on a page</p>`
creator_id	**Type**: User **Description**: ID of the user who created the comment **Example:** `ff8080817572401e01757240b3520000`
last_modifier_id	**Type**: User **Description**: ID of the user who last modified the comment **Example:** `ff8080817572401e01757240b3520000`
created_date	**Type**: Time **Description**: Creation timestamp **Example:** `2021-02-26T04:14:38Z`
updated_date	**Type**: Time **Description**: Last modification timestamp **Example:** `2021-02-26T04:14:38Z`

Analytics events file

Field	Description
instance_url	**Type**: String **Description**: Base URL of the current instance **Example:** `https://example.com`
event_id	**Type**: Number **Description**: Unique ID of the analytics event

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

event_name	**Type**: String **Description**: Name of the analytics event. Events include page_viewed, page_created, page_updated, blog_viewed, blog_created, blog_updated, comment_created, attachment_viewed, attachment_created. **Example**: `page_created`
created_date	**Type**: Time **Description**: Creation timestamp **Example:** `2021-02-26T04:14:38Z`
event_author_id	**Type**: User **Description**: ID of the user who performed the action that triggered the event **Example:** `ff8080817572401e01757240b3520000`
event_space_key	**Type**: String **Description**: Space key of the space the event was triggered in or affects (affected object) **Example:** SPACEKEY
event_container_id	**Type**: Number **Description**: ID of the containing entity. For pages this is the page ID, for attachments and comments, its the page ID of the page the attachment or comment appears on.
event_content_id	**Type**: Number **Description**: ID of the entity. For pages this is the page ID, for attachments, it is the attachment ID, and for comments it's the comment ID.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Configuring Confluence

This section focuses on settings and configurations within the Confluence application.

For guidelines on external configuration, see Configuring a Confluence Environment.

- Viewing System Information
- Configuring the Server Base URL
- Configuring the Confluence Search and Index
- Configuring Mail
- Configuring Character Encoding
- Other Settings
- Configuring System Properties
- Working with Confluence Logs
- Scheduled Jobs
- Configuring the Allowlist
- Configuring the Time Interval at which Drafts are Saved

Related pages:

- Customizing your Confluence Site
- Confluence administrator's guide

Viewing System Information

The System Information screen provides information about Confluence's configuration, which plugins are in use, and the environment in which Confluence has been deployed.

To view your system information go to > **General Configuration** > **System Information**.

Notes:

- The handy **memory graph** helps you keep track of Confluence's memory usage.
- Your system configuration information is helpful to Atlassian Support when diagnosing errors you may face using Confluence. When logging a **support request** or **bug report**, please provide as much detail as possible about your installation and environment.

Related pages:

- Cache Statistics
- Live Monitoring Using the JMX Interface
- Tracking Customizations Made to your Confluence Installation

Tracking Customizations Made to your Confluence Installation

The '**Modification**' section of the Confluence '**System Information**' screen lists the files that have been changed since your Confluence application was installed. You will find this information particularly useful when upgrading Confluence to a new version, because you will need to re-apply all customizations after the upgrade.

To see the modifications made to files in your Confluence installation:

1. Choose the**cog icon** , then choose**General Configuration**
2. Select '**System Information**' in the 'Administration' section of the left-hand panel.
3. Scroll down to the section titled '**Modification**'.

Screenshot: Modifications tracker on the Confluence System Information screen

	Modification
Modified	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.

Viewing System Properties

After adding memory, setting a proxy, or changing other Java options, it can be difficult to diagnose whether the system has picked them up. This page tells you how to view the system properties that your Confluence site is using.

You can see the expanded system properties on the 'System Information' screen of the Confluence Administration Console. You do not need to restart Confluence before viewing the information.

To see the system properties recognized by your Confluence installation:

1. Choose the**cog icon** , then choose**General Configuration**
2. Choose **System Information** in the left-hand panel.
3. Scroll down to the section titled **System Properties**.

Configuring the Server Base URL

The**Server Base URL**is the URL via which users access Confluence. The base URL**must**be set to the same URL by which browsers will be viewing your Confluence site.

Confluence will automatically detect the base URL during setup, but you may need to set it manually if your site's URL changes or if you set up Confluence from a different URL to the one that will be used to access it publicly.

You need to haveSystem Administratorpermissions in order to perform this function.

To change the Server Base URL:

1. Go to > **General Configuration**.
2. Select**Edit**.
3. Enter the new URL in the**Server Base URL**field.
4. **Save** your changes.

Example

If Confluence is installed to run in a non-root context path (that is, it has a context path), then the server base URL should include this context path. For example, if Confluence is running at:

```
http://www.foobar.com/confluence
```

then the server base URL should be:

```
http://www.foobar.com/confluence
```

Notes

- **Using different URLs.** If you configure a different base URL or if visitors use some other URL to access Confluence, it is possible that you may encounter errors while viewing some pages.
- **Changing the context path.** If you change the context path of your base URL, you also need to make these changes:
    1. Stop Confluence.
    2. Go to the Confluence installation directory and edit `<installation-directory>\conf\server.xml`.
    3. Change the value of the `path` attribute in the `Context` element to reflect the context path. For example, if Confluence is running at `http://www.foobar.com/confluence`, then your`path` attribute should look like this:

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

        In this example we've used`/confluence`as the context path. Note that you can't use`/resources` as your context path, as this is used by Confluence, and will cause problems later on.
    4. Save the file.
    5. Go tothe Confluence home directory and edit`<confluence-home>/confluence.cfg.xml`
    6. Change`confluence.webapp.context.path`to reflect the new context path. For example:

        ```
        <property name="confluence.webapp.context.path">/confluence</property>
        ```

7. Restart Confluence and check you can access it at http://www.foobar.com/confluence.

   You may also want toclear the Confluence plugins cachebefore restarting.

- **Proxies**. If you are running behind a proxy, ensure that the proxy name matches the base URL. For example: `proxyName="foobar.com" proxyPort="443" scheme="https"`. This will make sure we are passing the information correctly. For more information on proxing Atlassian applications, seeProxying Atlassian Server applications.
- This information needs to be added in the`Connector`element at `{CONFLUENCE_INSTALLATION}\conf\server.xml`.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Configuring the Confluence Search and Index

Confluence administrators can adjust the behavior of the Confluence search, and manage the index used by the search.

- Configuring Indexing Language
- Configuring Search
- Content Index Administration
- Enabling OpenSearch
- Rebuilding the Ancestor Table
- Setting Up Confluence to Index External Sites
- Setting Up an External Search Tool to Index Confluence

Related pages:

- Search
- Confluence Administrator's Guide

Configuring Indexing Language

Changing the indexing language to be used in your Confluence site may improve the accuracy of Confluence search results, if the majority of the content in your site is in a language other than English.

Confluence supports content indexing in:

- Arabic
- Brazilian
- Chinese
- CJK
- Custom Japanese
- Czech
- Danish
- Dutch
- English (default)
- Finish
- French
- German
- Greek
- Hungarian
- Italian
- Norwegian
- Persian
- Romanian
- Russian
- Spanish
- Swedish

To configure the indexing language:

1. Go to > **General Configuration**then choose **Edit**.
2. Select the **Indexing Language** from the dropdown list in the **Formatting and International Settings** section.
3. Choose **Save**.

Related pages:

- Choosing a Default Language
- Installing a Language Pack
- Content Index Administration
- How to Rebuild the Content Indexes From Scratch on Confluence Server

Configuring Search

There are a few ways to search for content in Confluence:

- Using the search panel, which allows you to quickly search and filter results.
- Using the advanced search page.
- Using a search macro embedded on a Confluence page (for example, the Livesearch Macro or QuickNav Gadget).

Read more about the differentsearch options in Confluence.

Related pages:

- Search

By default, the search panel feature is enabled, with the maximum number of simultaneous requests set to 40. These options can be modified as described below.

Set the number of simultaneous search requests

Confluence admins can set the maximum number of simultaneous searches users can perform using the search panel. By default, the maximum is set to 40. This limit applies to a single Confluence node. If you're running Confluence Data Center with multiple nodes, this number will increase.

If your Confluence server serves a large number of individuals who use this feature regularly, some of whom are being denied access to it, you may wish to increase this value.

To change the maximum number of simultaneous search requests in Confluence:

1. Choose the**cog icon** , then choose**General Configuration**
2. Choose **Further Configuration** in the left-hand panel.
3. Choose **Edit**.
4. Enter the appropriate number in the field beside **Max Simultaneous Requests**.
5. Choose **Save**.

Disable quick search options

The search panel feature offers a quick way for users to search and filter content in Confluence. We recommend keeping this feature enabled, unless it's causing significant performance issues on your site.

If you disable the quick search option:

- The search panel will no longer appear when users click the search field. When you enter a search query, we'll take you to the advanced search page.
- The Confluence QuickNav Gadget will no longer drop down a list of search results. When you enter a search query, we'll take you to the advanced search page.

Other search macros, including the Livesearch Macro and the Page Tree Search Macro, won't be affected if you disable the quick search option.

To disable quick search options from your Confluence site:

1. Choose the**cog icon** , then choose**General Configuration**
2. Choose**Further Configuration**in the left-hand panel.
3. Choose**Edit**.
4. Deselect the **Quick Search** checkbox.
5. Choose**Save**.

Content Index Administration

The search index is used by search, the dashboard, some macros, and all the other places where we show information about the content in your Confluence site. The search index is made up of:

- a **content index** which contains content such as the text of pages, blog posts, and comments
- a **change index** which contains data about each change, such as when a page was last edited

On this page:

- View the index queues
- Rebuild the search index
- Location of search indexes
- Index recovery in a cluster
- Check the size of your index
- Troubleshooting

Related pages:

- Scheduled Jobs
- Search
- Configuring the Confluence Search and Index

These indexes are updated automatically as people get work done on your site. Changes, such as a new page, comment, or edit to an existing page, aren't updated in each index immediately. They're placed into queues and regularly processed in batches (as often as every 5 seconds) in the background as you work.

View the index queues

It can take a while for the queues to process if there are thousands of changes to your site within a short period.

To check the contents of the queue:

1. Go to > **General Configuration**>**Content Indexing**.
2. Select the **Content queue** or **Change queue** tab.

Here you can see the number of items in the queue, the last time the queue was processed, and how long it took to process. This information is useful for troubleshooting if your users report issues with search or dashboard activity feeds.

Rebuild the search index

There are situations where you may need to rebuild the search reindex; for example, when users report issues with search, dashboard activity feeds, or when directed to as part of an upgrade.

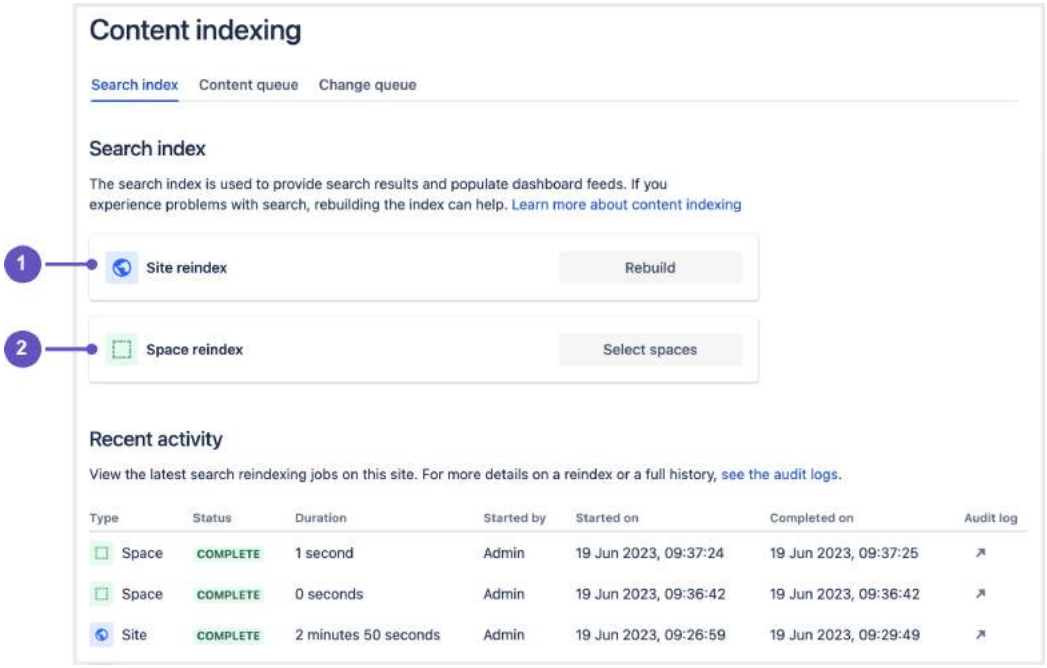

Screenshot: Search index UI screen in the admin console

You have the option to rebuild the search index for:

1. an entire site
2. a space or multiple spaces

You should run a **space reindex** when:

- you know the exact issue and the affected spaces
- you want to stagger or spread out a full site reindex
- the content index for a space is corrupted after importing it to your site
- a page is moved from one space to another, and the index for the page is corrupted in the process

You should run a **site reindex** when:

- a space reindex fails to resolve the issue
- a user can't be found or mentioned
- an admin can't find the target spaces when reindexing a space; this may mean the space directory index is broken

By rebuilding the search index for a site or space, you rebuild both the content index and change index. This can take some time for large sites. You should also consider this when deciding what type of reindex to run.

Reindexing a space

To reindex for one or more spaces:

1. Go to > **General Configuration**> **Content indexing**.
2. Next to **Spaces reindex**, select **Select spaces**.
3. Search for the spaces by entering the space name into the field, then select those you want to reindex.
4. Select **Rebuild** and follow the prompts to confirm you want to rebuild the index.

Reindexing a site

To reindex the entire site:

1. Go to > **General Configuration**> **Content indexing**.
2. Next to **Site reindex**, select**Rebuild** and follow the prompt to confirm you want to rebuild the index.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Track the progress of the reindex

You can get the status of a reindexing job in the **Recent activity**table. To learn more about a job (including any errors or issues that occurred), select the arrow at the end of the table row to see the **audit log** for that job.

Screenshot: Search index UI screen in the admin console

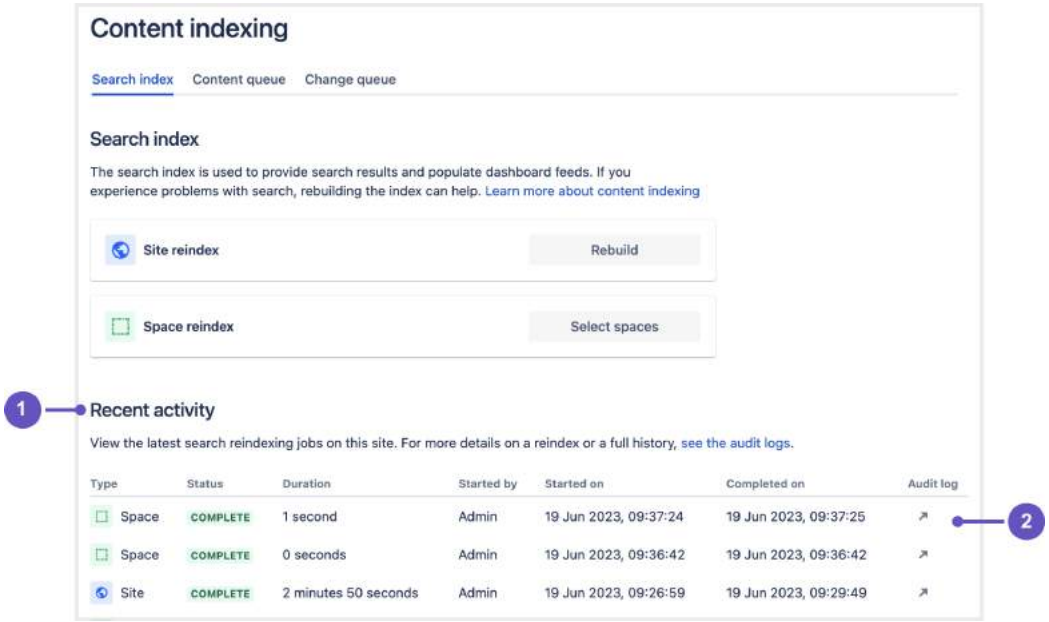

1. Recent activity table
2. See the audit log for a specific reindex job

For even more details, you can also check theConfluence indexing logs at `atlassian-confluence-index.log` (see Working with Confluence Logs for how to access these logs). Examples of the details available in the Confluence indexing logs are below:

Progress updates

Content reindexing happens concurrently in batches. The percentage of content that has been processed will be displayed regularly every time a batch of content is processed.

Example index log entry

```
2023-02-02 12:16:44,342 INFO [Indexer: 1] [confluence.internal.index.ConcurrentBatchIndexer] logProgress
Re-index progress: 38 of 61. 62% complete. Memory usage: 1 GB free, 2 GB total
```

However, not all content may be successfully indexed due to unhandled errors.

Unhandled errors

Unhandled errors that occur will impact content in the same batch. However, it won't impact the indexing of content in other batches. Unhandled errors will also be logged.

If you find an unhandled error, you should find out the root cause and resolve the issue before re-running the reindex.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Example index log entry

```
2023-02-01 12:24:50,043 ERROR [Indexer: 1] [confluence.internal.index.ConcurrentBatchIndexer]
lambda$null$2 An error occurred while re-indexing a batch. Only the particular batch which had an error
occur will not be re-indexed correctly.
-- referer: http://localhost:8080/confluence/plugins/servlet/rebuildindex | url: /confluence/rest
/prototype/latest/index/reindex | traceId: 0463502f0ab3faab | userName: admin
java.lang.RuntimeException: Some unhandled exception
....
```

Reindex complete

When the progress reaches 100%, reindexing is complete.

Example index log entry

```
2023-02-02 12:16:44,553 INFO [Indexer: 1] [confluence.internal.index.ConcurrentBatchIndexer] logProgress
Re-index progress: 100% complete. 61 items have been reindexed
```

Keep in mind that because of unhandled errors, it is possible that not all of your content has been successfully reindexed. The progress reflects how much content has been processed, rather than how much content has actually been successfully indexed.

Impact on end users

Users can continue to search and use Confluence but may experience some performance degradation, especially when running a site reindex. This is because rebuilding the index increases the load on your server.

Rebuilding an index can take several hours. The amount of time depends on the number, type, and size of pages and attachments on your site, the amount of memory allocated, and disk throughput.

If you have a very large site, there are some things you can do to reduce the impact on your users:

- If you're running Confluence on a single node,kick off the rebuild on a weekend, or during a scheduled maintenance window.
- If you're running Confluence in a cluster, remove the node rebuilding the index from your load balancer. Then, Confluence will then continue to use the existing index until the new index has been rebuilt successfully. Once propagation is complete, you can add the node back into the pool.

Propagate the search index to your cluster

For site reindexing, once the search index is rebuilt on the current node, we automatically propagate the index files to all other nodes in the cluster.

The index files will only be propagated to nodes that have joined the cluster. If Confluence isn't running on a node, we won't be able to propagate the index to that node.

If there's a problem, for example, if a node becomes unavailable, or there's insufficient disk space to copy the index, you will see an error status like `PROPAGATION FAILED`. Go to the audit log for the job to find details about the error.

For space reindexing, the search index is rebuilt across all nodes concurrently so node propagation is not required.

Disk space requirements

If you run Confluence in a cluster, before you do a site reindex ensure you have enough free space in your shared home directory to accommodate an additional reindex snapshot. This snapshot is required for node propagation.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Location of search indexes

You can find the Confluence index in the `<home-directory>/index`directory.

If you're running Confluence in a cluster, a full copy of the Confluence indexes are stored in the `<local-home>/index`directory on each Confluence node. A journal service keeps each index in sync.

If you need to see the contents of the search index for any reason, there is a tool you can use to browse the index directly. SeeHow to view the contents of the search index in Confluence Server and Data Center.

Index recovery in a cluster

If you run Confluence in a cluster, a snapshot of your site's search index is stored in the shared home directory. These snapshots are created by the Clean Journal Entries scheduled job which, by default, runs once per day.

When you start a Confluence node, it will check whether its index is current, and if not, it will request a recovery snapshot from the shared home directory. If a snapshot is not available, it will generate a snapshot from a running node (with a matching build number). Once the recovery snapshot is extracted into the index directory, Confluence will continue the startup process. The journal service will then make any further updates required to bring the index up to date.

If the snapshot can't be generated or is not received in time, existing index files will be removed and Confluence will perform a reindex on that node. If your index is very large or your file system is slow, you may need to increase the time Confluence waits for the snapshot to be generated using the`confluence.cluster.index.recovery.generation.timeout`system property.

Index recovery only happens on node startup, so if you suspect a problem with a particular cluster node's index, restart that node to trigger index recovery.

The index recovery snapshot isn't used when you manually rebuild your index from the UI. The rebuild process generates a brand new snapshot, before propagating it to other nodes in the cluster.

Check the size of your index

You can measure the index size in two ways, size on disk, or you can use the number of pages and blogs as a rough indication of the amount of content in the index.

To check the size on disk:

1. Go to `<local-home>/index`
2. Check the size of that directory. The way you do this will depend on your operating system.

To check the number of pages and blogs in the index:

1. Go to > **General Configuration** > **System information**
2. Scroll down to the **Confluence usage** section and check the **Content (Current Versions)** value.

Troubleshooting

If you have problems rebuilding the search index, the following may help.

Can't rebuild the index

If you're unable to rebuild the index from the Confluence UI, or if you still have problems with search after rebuilding the index, you may need to rebuild the index from scratch. The way you do this depends on whether Confluence is running in a cluster:

- How to Rebuild the Content Indexes From Scratch on Confluence Server
- How to Rebuild the Content Indexes From Scratch

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Can't access content indexing page

If the content indexing page does not load properly, and you see a "We can't check the status of your index, you may have lost your connection, refresh the page to try again" error, try updating your browser to the latest version.

Poor performance while rebuilding the index

If you experience stability problems while the index is being rebuilt,you canreduce the number of threads Confluence should use to rebuild the index. Set the `reindex.thread.count` system propertyto define the maximum number of threads that can be used.

If both`reindex.thread.count`and`index.queue.thread.count` are unset, the reindex thread count defaults to the number of CPUs on that Confluence server.

Out-of-memory errors while rebuilding the index

If you experience out of memory errors while rebuilding the index, increasing the heap memory available to Confluence may help. See How to fix out of memory errors by increasing available memory.

Rebuilt site index failed to propagate to other nodes in the cluster

This generally happens when there is not enough free disk space for the local home directory on each node to accommodate two copies of the index. SeeFailed to propagate index in Confluence Data Center 7.7 and laterto find out how to re-try the propagation.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Enabling OpenSearch

With OpenSearch autodiscovery, you can add Confluence search to your Firefox or Internet Explorer search box. By default, OpenSearch autodiscovery is enabled. This feature can be enabled or disabled as described below.

To enable or disable OpenSearch autodiscovery:

1. Choose the**cog icon** , then choose**General Configuration**
2. Choose **Further Configuration** in the left-hand panel.
3. Choose **Edit**.
4. Select the **Open Search** checkbox to enable this feature (deselect to disable).
5. Choose **Save**.

Information about OpenSearch

- Confluence supports the autodiscovery part of the OpenSearch standard, by supplying an OpenSearch description document. This is an XML file that describes the web interface provided by Confluence's search function.
- Any client applications that support OpenSearch will be able to add Confluence to their list of search engines.

Related pages:

- Search
- Confluence Administrator's Guide

Rebuilding the Ancestor Table

> The way you fix problems with the ancestor table changed significantly in **Confluence 6.14**.
>
> **If you're using Confluence 6.13 or earlier**, seeRebuilding the Ancestor Table in Confluence 6.13.5 or earlierto find out how to rebuild the ancestor table.

The `ancestor`table records the parent and descendant (child) relationship between pages. It is also used when determining whether a page will inherit view restrictions from a parent page.

Occasionally records in the ancestor table can become corrupted. The **Repair the Ancestors Table** scheduled job finds and automatically fixes problems in the ancestors table in all current spaces. The job runs daily.

The job ignores archived spaces, so if you suspect there is a problem with the ancestor table in a particular archived space, you will need to change the status of the space to 'current', then run the job manually.

Repair the ancestor table manually

If you suspect there is a problem, you can also run this job manually.

1. Go to> **General Configuration**> **Scheduled Jobs**.
2. Locate the **Repair the Ancestors Table** job and choose **Run**.

The job should complete quickly, and there won't be any impact on users.

Viewing the job results

If you want to see the result of the job each time it runs, you can change thelogging levelof`com.atlassian.confluence.pages.ancestors`to INFO.

You'll then see a message similar to the one below in the Confluence application log, each time the job runs:

```
Ancestors have been repaired. Found and fixed 3 broken pages.
It took 71 sec for 6407 spaces, average space processing time 0 sec.
```

Rebuilding the ancestor table in earlier Confluence versions

If you're using Confluence 6.13 or earlier the way you rebuild the ancestor table is different. SeeRebuilding the Ancestor Table in Confluence 6.13.5 or earlierfor more information.

Setting Up Confluence to Index External Sites

Confluence cannot easily index external sites, due to the way Lucenesearch works in Confluence,butth ere are two alternatives:

1. Embed External Pages Into Confluence
2. Replace Confluence Search

Related pages:

- Setting Up an External Search Tool to Index Confluence
- Configuring the Confluence Search and Index

Embedding external pages into Confluence

If you only have a small number of external sites to index, you may prefer to enable theHTML-include Macroa nd use it embed the external content inside normal Confluence pages.

The actual content of the external site won't be indexed.

Replacing the Confluence search

Use your own programmer resources to replace Confluence's internal search with a crawler that indexes both Confluence and external sites. This advanced option is easier than modifying the internal search engine. It requires removing Confluence internal search from all pages and replacing the internal results page with your own crawler front-end.

1. Setup a replacement federated search engine to index the Confluence site, as well as your other sites, and provide the results that way. You would need to host a web crawler, such as these open-source crawlers. Note that you can perform a search in Confluence via the Confluence API.
2. Replace references to the internal search by modifying the site layout so that it links to your search front-end
3. Host another site containing the search front-end. You may wish to insert it into a suitable context path in your application server so that it appears to be from a path under Confluence. Tomcat sets Confluence's paths from the Confluence install\confluence\WEBINF\web.xml file.

Setting Up an External Search Tool to Index Confluence

Any web crawler can be configured to index Confluence content. If a login is required to view content that will be indexed, you should create a Confluence user specifically for the search crawler to use. Grant this user view rights to all content you wish to index, but deny that user all delete and administration rights. This ensures that an aggressive crawler will not be able to perform actions that could modify the site.

External applications can also use the search function in theConfluence APIs.

Related pages:

- Setting Up Confluence to Index External Sites
- Configuring the Confluence Search and Index

Configuring Mail

- Configuring a Server for Outgoing Mail
- Configuring a Server for Incoming Mail
- Setting Up a Mail Session for the Confluence Distribution
- Configuring the Recommended Updates Email Notification
- The Mail Queue

- Customizing Email Templates

Configuring a Server for Outgoing Mail

Configuring your Confluence server to send email messages allows your Confluence users to:

- Receive emailed notifications and daily reports of updates.
- Send a page via email.

You can personalize email notifications by configuring the 'From' field to include the name and email address of the Confluence user who made the change.

You need System Administrator permissions in order to configure Confluence's email server settings.

On this page:

- Configuring Confluence to send email messages
- Testing the email settings

Related pages:

- The Mail Queue
- Setting Up a Mail Session for the Confluence Distribution

Configuring Confluence to send email messages

To configure Confluence to send outgoing mail:

1. Go to > **General Configuration** > **Mail Servers**. This will list all currently configured SMTP servers.
2. Click **Add New SMTP Server** (or **edit** an existing server).
3. Edit the following fields as required:
    - **Name**: By default, this is simply 'SMTP Server'.
    - **From Address**: Enter the email address that will be displayed in the 'from' field for email messages originating from this server.
      This field is mandatory. This must be an ordinary email address, you can'tenter variables in this field.
    - **From Name**: Enter the name that will be displayed in the 'from' field for email messages originating from this server. This is the text which appears before the user's registered email address (in square brackets).
      This field accepts the following variables, which reference specific details defined in the relevant Confluence user's profile:

Variable	Description
${fullname}	The user's full name.
${email}	The user's email address.
${email.hostname}	The domain/host name component of the user's email address.

      The default is '`${fullname} (Confluence)`'.
      Hence, if Joe Bloggs made a change to a page he was watching and the Confluence site's 'From Address' was set to `confluence-administrator@example-company.com`, then the 'From' field in his email notification would be: `Joe Bloggs (Confluence) <confluence-administrator@example-company.com>`.
    - **Subject Prefix**: Enter some text to appear at the beginning of the subject line.
4. Enter your **Hostname, Port,User name** and **Password** details.
   If your SMTP host uses the Transport Layer Security (TLS) protocol select **Use TLS**.
   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, seeSetting Up a Mail Session for the Confluence Distribution.

Testing the email settings

A Confluence administrator can test the email server as follows:

1. Set up a mail serveras described above.

2. Click **Send Test Email** to check that the server is working. Check that you get the test email in your inbox.
3. You can flush the email queue to send the email message immediately. Go to**Mail Queue**, and click **Flush Mail Queue**. See The Mail Queue.

A user can test that notifications are working as follows:

1. Go to your user profile (using the **Settings** link) and edit your email preferences. See Email Notifications.
2. Enable **Notify On My Actions**. (By default, Confluence does not send you notifications for your own changes.)
3. Go to a page you wish to get notifications about.
4. Choose**Watch** at the top-right of the page. See Watch Pages, Spaces and Blogs.
5. Edit the page, make a change, and save the page.
6. Check your email inbox. You may need to wait a while for the email message to arrive.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Configuring a Server for Incoming Mail

Configuring your Confluence server to receive emails from another POP or IMAP mail server allows your users to create pages and reply to page comments with email.

You need System Administrator permissions in order to configure Confluence's email server settings.

Configuring Confluence to receive email messages

To configure Confluence to receive incoming mail:

1. Go to > **General Configuration**> **Mail Servers**. This will list all currently configured servers.
2. Select **Add a new POP mail server** or **Add a new IMAP mail server**, or edit an existing server.
3. Complete the fields on this page, using the below table as a guide. All fields are required unless specified otherwise.

Name	The default is POP server or IMAP server. You may want to change it to another name if you are configuring more than one mail server of the same type.
To address	Specify a valid email address for Confluence to retrieve emails from.

Server hostname	Specify the hostname or IP address of your POP or IMAP mail server.
	
	Below are the mail server settings for Google and Microsoft:
	
	• Google (See official guide)
	  ◦ POP: pop.gmail.com
	  ◦ IMAP: imap.gmail.com
	• Microsoft (See official guide)
	  ◦ POP: outlook.office365.com
	  ◦ IMAP: outlook.office365.com

Protocol	Select whether your POP or IMAP mail server uses a standard (i.e. POP3 or IMAP) or secure (i.e. POP3Sor IMAPS) protocol. To use the OAuth 2.0 integration as your authorization method, select a secure protocol.
Server port	This is the port that will be used to retrieve mail from your POP OR IMAP account. When you select a protocol, the port will be changed to the default value. The defaults are: • POP: 110 • POP3S: 995 • IMAP: 143 • IMAPS: 993 You may specify your own custom port.

Authorization	Select a way to authenticate to the mail server. The default value is **Basic Authentication**.
	
	To use the OAuth 2.0 authorization method here, you will need to first configure Confluence as an OAuth 2.0 client in **Application links**. Learn how to configure an outgoing link with some notes below:
	
	• If Microsoft is the external provider, the following scopes must be entered in the Scopes field:
	  • `https://outlook.office.com/IMAP.AccessAsUser.All` (required for an IMAP mail server)
	  • https://outlook.office.com/POP.AccessAsUser.All (required for a POP3 mail server)
	  • offline_access (required for any mail server)
	• If Google is the external provider, https://mail.google.com is the only scope required for either POP3 or IMAP mail server

Username	This is the username used to authenticate your mail account.
Password	This field is only required if you are using basic authentication. This is the password for your mail account. Google and Microsoft have disabled using passwords as an authentication method. To connect to your Gmail or Microsoft Exchange Online account, youll need to use OAuth 2.0 as your authorization method. See the "Authorization" row in this table to learn how to do this.

4. Once completed, select **Authorize**. You will be redirected to your service provider's site to log in to your account and authorize the connection. After the connection is authorized successfully, you will be redirected back to your app.
5. Select **Test Connection**to check that Confluence can communicate with the mail server that you have just configured.
6. Select **Submit** to save the new mail server**.** If you forget to select Submit, the mail server configuration will not be saved even after its connection is authorized.

Troubleshooting

If you experience any errors while authorizing or testing the connection, check the application link has been configured correctly. See Linking to another application

If this keeps happening, inspect atlassian-confluence.log for specific error details. See Working with Confluence logs to learn how to do this. For help diagnosing any issues, reach out to Support.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Setting Up a Mail Session for the Confluence Distribution

The simplest way to set up a mail server is through the Confluence Administration console. SeeConfiguring a Server for Outgoing Mail.

If youwant to add different options or parameters you can also set up a mail session for the Confluence distribution. In the example below we'll set up Gmail.

To set up a mail session for the Confluence distribution:

1. Stop Confluence.
2. Move (don't copy) the following files from `<confluence-install>\confluence\WEB-INF\lib` to `<confluence-install>\lib:`

   `com.sun.activation_jakarta.activation-x.x.x.jar`
   `com.sun.mail_jakarta.mail-x.x.x.jar`

   (x.x.x. represents the version numbers on the jar files in your installation)

   Don't leave a renamed backup of the jar files in`\confluence\WEB-INF\lib`. Even with a different file name, the files will still be loaded as long as it remains in the directory.
3. Edit the `<confluence-install>\conf\server.xml` file and add the following at the end of the Confluence <context> tag, just before `</Context>`.
   Note: you're editing the <context> tag that contains the *Confluence* context path, not the one that contains the *Synchrony* context path.

   ```
   <Resource name="mail/GmailSMTPServer"
       auth="Container"
       type="javax.mail.Session"
       mail.smtp.host="smtp.gmail.com"
       mail.smtp.port="465"
       mail.smtp.auth="true"
       mail.smtp.user="yourEmailAddress@gmail.com"
       password="yourPassword"
       mail.smtp.starttls.enable="true"
       mail.transport.protocol="smtps"
       mail.smtp.socketFactory.class="javax.net.ssl.SSLSocketFactory"
   />
   ```

4. Restart Confluence.
5. Go to > **General Configuration** > **Mail Servers**.
6. Select either**Edit an existing configuration**, or**Add a new SMTP mail server**.
7. Edit the server settings as necessary, and set the**JNDI Location**as:

   ```
   java:comp/env/mail/GmailSMTPServer
   ```

   Note that the JNDI Location is case sensitive and must match the resource name specified in server.xml.
8. Save your changes and send a test email.

Configuring the Recommended Updates Email Notification

Confluence sends a regular email report to subscribers, containing the top content that is relevant to the person receiving the message, from spaces they have permission to view. This is called the 'Recommended Updates' notification.

If you have Confluence Administrator or System Administrator permissions, you can configure the default settings that determine how often the Recommended Updates notification is sent. When new users are added to Confluence, the default settings will be applied to their user profiles.

Confluence users can choose their personal settings, which will override the defaults. See Email Notifications.

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

Initial settings of the defaults

When you install Confluence, the initial values of the default settings are as follows:

- The default frequency is weekly.
- If your Confluence site has public signup enabled, the Recommended Updates notification is disabled by default. If public signup is not enabled, the notification is enabled by default.

You can change the above settings, specifying a different default value for the site.

Notes:

- The Recommended Updates notification is sent only to people who have a user profile in Confluence. If your Confluence site uses external user management, such as LDAP, then people will receive the report only after they have logged in for the first time. (The first login creates their user profile.)
- The daily email message is sent at 1 p.m. in the user's configured time zone.
- The weekly email message is sent at 1 p.m. on Thursdays in the user's configured time zone.

Configuring the Recommended Updates notification

You can set thedefault send option (send / do not send) and the default schedule (daily or weekly).

To configure the Recommended Updates email notification:

1. Choose the**cog icon** , then choose**General Configuration**
2. Click **Recommended Updates Email** in the left-hand panel.

Disabling the Recommended Updates notification for the entire site

You can also turn off the **recommended updates** notification for the entire site, by disabling the 'Confluence daily summary email' system app. SeeDisabling and enabling apps.

The Mail Queue

Email messages waiting to be sent are queued in a mail queue and periodically flushed from Confluence once a minute. A Confluence administrator can also manually flush messages from the mail queue.

If there is an error sending messages, the failed email messages are sent to an error queue from which you can either try to resend them or delete them.

To view the mail queue:

1. Choose the**cog icon** , then choose**General Configuration**
2. Choose **Mail Queue** in the left-hand panel. This will display the email messages currently in the queue.
3. Choose **Flush Mail Queue** to send all email messages immediately.
4. Choose **Error Queue** to view failed email messages. You can try to **Resend** the messages, which will flush the mails back to the mail queue, or you can **Delete** them from here.

Related pages:

- Configuring a Server for Outgoing Mail
- Setting Up a Mail Session for the Confluence Distribution

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

Configuring Character Encoding

Confluence and your database must be configured to use the same character encoding. To avoid problems with character encoding**always set all character encodings to UTF-8** (or the equivalent for your database, for example, UTF8MB4 for MySQL databases, or AL32UTF8 for Oracle databases).

On this page:

- Configuring Confluence character encoding
- Database character encoding
- Problems with character encodings

Related pages:

- Configuring Database Character Encoding

Configuring Confluence character encoding

By default, Confluence uses UTF-8 character encoding.

While it is possible to change the character encoding, it is **not recommended**. Changing the Confluence character encoding will change your HTTP request and response encoding and your filesystem encoding as used by exports and Velocity templates. You may also be prevented from restarting or upgrading Confluence, depending on your database.

To change the Confluence character encoding (not recommended):

1. Shut down Confluence and perform a database backup
2. Run:

```
UPDATE  BANDANA set BANDANAVALUE = REPLACE(BANDANAVALUE, 'UTF-8', 'UTF-16') where BANDANAKEY =
'atlassian.confluence.settings';
```

3. Verify:

```
SELECT BANDANAVALUE FROM BANDANA where BANDANAKEY = 'atlassian.confluence.settings';
```

4. Start Confluence

Database character encoding

Your database, and the JDBC connection to it, must be configured to use UTF-8 (or the equivalent for your database, for example, UTF8MB4 for MySQL databases, or AL32UTF8 for Oracle databases). There are a number of checks in place to warn you if your database character encoding is incorrect.

SeeConfiguring Database Character Encodingfor more information.

Problems with character encodings

SeeTroubleshooting Character Encodingsto find out how to test your character encoding.

Troubleshooting Character Encodings

If character encoding is not configured correctly in your Confluence site, you may experience problems like:

- Non-ASCII characters appearing as question marks (?)
- Page links with non-ASCII characters not working
- Single characters being displayed as two characters
- Garbled text appearing

To diagnose the problem, follow these steps.

1. Run the encoding test

Confluence includes an encoding test that can reveal problems with your configuration. You'll need to be a Confluence admin to do this.

1. Head to**<your-confluence-url>/admin/encodingtest.action**
2. Follow the prompts to paste a line of text and start the test. You can also paste text in a specific language, for example Japanese, if you're experiencing a particular problem with that language.

If the text displayed in the encoding test is different to what you entered, then there are problems with your character encoding settings. Here's what a successful test looks like.

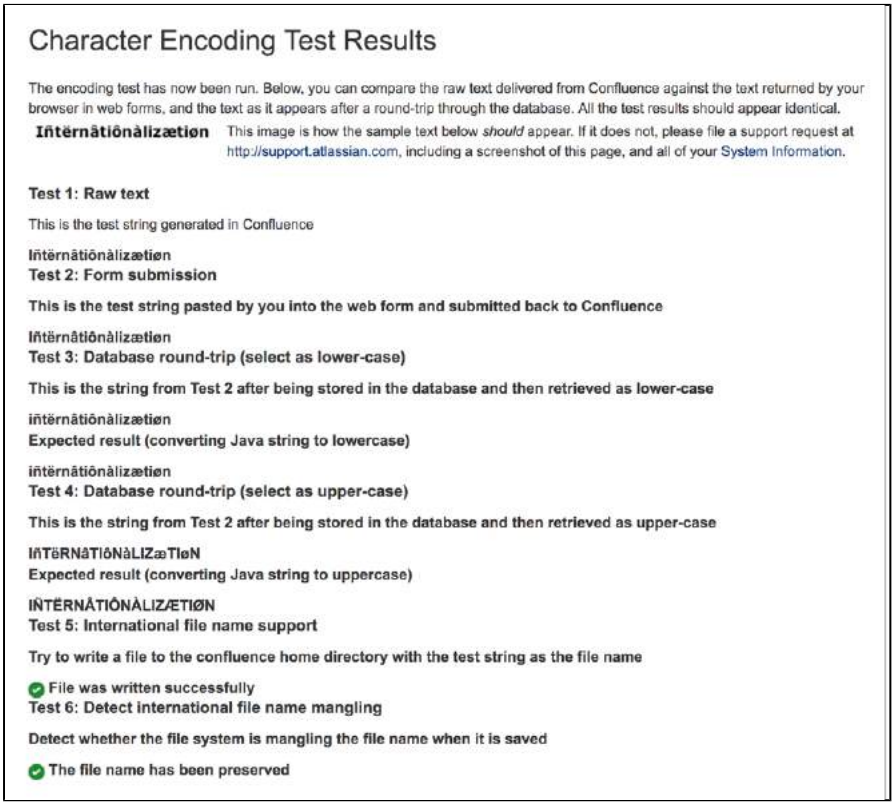

2. Use the same encoding for your database

Your database and Confluence must use the same character encoding. SeeConfiguring Database Character Encodingfor more information.

3. Get help

If you're still having problems with character encoding,create a support request, and our support team will help you solve the problem.

Include the following details to help us identify your problem:

- screenshots of the problem occurring

- results of the encoding test
- information about your database (including version)
- A copy of the information on yourSystem Information page.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

"" Euro character not displaying properly

The (euro) symbol is a three byte character, with byte values in file (UTF-8) of 0xE2, 0x82, 0xAC.

Sometimes, if the character encoding is not set consistently among all participating entities of the system, Confluence, server and the database, one may experience strange behavior.

> ...
> I write a page with a Euro sign in it (). All is well, the Euro sign shows up in the wiki markup 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 recognize the symbol.

There are two potential reasons for this behavior:

Database and Confluence is using utf-8 encoding. The connection is not.

When data transferred to it via the connection which does not use utf-8 encoding gets encoded incorrectly. Hence, updating the connection encoding may resolve this problem from now on, yet it probably would notaffect already existing data.

Database is not using utf-8. Confluence and your connection are.

If your Database encoding is not set to UTF-8, yet is using some other encoding such as *latin1*, it could be one of the potential reasons why you lose the "" characters at some stage. It could be occurring due to **caching**. When Confluence saves data to the database, it may also keep a local cached copy. If the database encoding is set incorrectly, the Euro character may not be correctly recorded in the database, but Confluence will continue to use its cached copy of that data (which is encoded correctly). The encoding error will only be noticed when the cache expires, and the incorrectly encoded data is fetched from the database.

> For instance the *latin1* encoding would store and display all 2-byte UTF8 characters correctly except for the euro character which is replaced by '?' before being stored. As Confluence's encoding was set to UTF-8,the 2-byte UTF-8 characters were stored in *latin1* database assuming that they were two *latin1* different characters, instead of one utf8 character. Nevertheless, this is not the case for 3-byte utf8 characters, such as the Euro symbol.

Please ensure that you set the character encoding to UTF-8 for all the entities of your system as advised in this guide.

MySQL 3.x Character Encoding Problems

MySQL 3.x is known to have some problems upper- and lower-casing certain (non-ASCII) characters.

Diagnosing the problem

1. Follow the instructions for Troubleshooting Character Encodings.
2. If the upper- and lower-cased strings displayed on the Encoding Test are different, then your database is probably affected.

An example (faulty) output of the Encoding Test is shown below:

Screenshot: Encoding Test Output (excerpt)

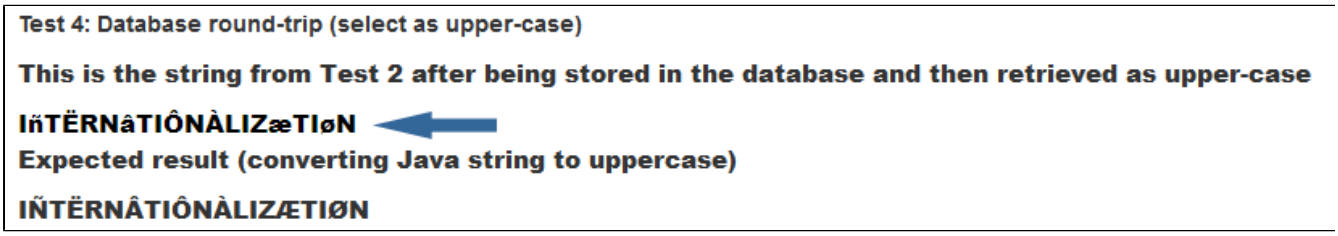

Solution

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.

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

Related pages:

- Global Permissions Overview
- Disabling and enabling apps
- Attachment Storage Configuration

Introduction to Confluence's WebDAV Client Integration

By default, all WebDAV clients have permission to write to Confluence. Write permissions include the ability for a WebDAV client to create, edit, move or delete content associated with spaces, pages and attachments in a Confluence installation.

On the '**WebDAV Configuration**' screen in the Confluence Administration Console, you can:

- Deny a WebDAV client write permissions to a Confluence installationusing a regular expression (regex)
- Disable or enable strict path checking
- Enable or disable access to specific virtual files/folders

Note:

- The 'WebDav Configuration' page is only available if the WebDAV plugin has been enabled. This plugin is bundled with Confluence, and can be enabled or disabled by theSystem Administrator.
- The settings on the 'WebDav Configuration' page**donotapply**toexternal 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

You can successfully connect but you can't see content when using HTTPS, so this technique won't work for Confluence Cloud. Use a third-party WebDAV client instead.

To use Finder to view and manage Confluence content:

1. In Finder choose**Go**>**Connect to Server**
2. Enter your Confluence URL in this format:

```
http://<confluence base URL>/plugins/servlet/confluence/default
```

For example if your Confluence URL ishttp://ourconfluence.sample.com/wikiyou would enter:

```
http://ourconfluence.sample.com/wiki/plugins/servlet/confluence/default
```

3. Enter your**Confluence**username and password and click**Connect**

> Use your username (jsmith), not your email address, unless your email address is your username.

Confluence will appear as a shared drive in Finder. You can use the same URL to connect using a third party WebDav client, like CyberDuck.

Accessing Confluence in Explorer in Microsoft Windows

This section covers the two methods for configuring a WebDAV client natively in Microsoft Windows:

- As a network drive
- As a web folder

If possible, use the network drive method as this will enable more comprehensive WebDAV client interaction with Confluence than that provided by a web folder. However, your Confluence instance must meet several environmental constraints if you use this method. If you cannot configure your instance to meet these requirements, then use the web folder method or third-party WebDAV client software.

If you're using SSL you may need to add @SSLto the end of your server URL, for example:

```
http://<confluence server url>@SSL/confluence/plugins/servlet/confluence/default
```

If you run into any problems with the procedures in this section, please refer to theWebDAV Troubleshooting page.

Windows Network Drive

To map a Confluence WebDAV client network drive, your Confluence instance must be configured so that*all*o f the following criteria is met:

- Has no context root
- There's an issuethat can prevent Network Drives from being mapped. Please use the Network Folders steps below as a workaround.

The reason for these restrictions results from limitations in Microsoft's Mini-Redirector component. For more information, please refer to Microsoft'sserver discovery issue.

To map a Confluence WebDAV client network drive in Microsoft Windows:

1. In Windows go to**Map Network Drive.**
   SeeMap a network drive in the Windows documentation to find out how to get to this in your version of Windows.
2. Specify the following input to map the WebDAV client as a network drive:
   - **Drive:**`<Any drive letter>`(for example,`Z:`)
   - **Folder:**`\\<hostname>\webdav`(for example,`\\localhost\webdav`)
3. Click**Finish**

When prompted for login credentials, specify your Confluence username and password.

Windows Web Folder

To map a Confluence WebDAV client web folder:

1. Go to**My Network Places**and choose**Add a network place**and follow the prompts.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

2. In the 'Internet or network address' field, enter the URL for the Confluence WebDAV location (for example,`http://<confluence server url>/confluence/plugins/servlet/confluence/defaultorhttp://<confluence server url>/plugins/servlet/confluence/default`) and click**Next**

3. Enter your Confluence username and password
4. Provide a meaningful name for your web folder andproceed with the wizard
5. Click**Finish**

Screenshot: A Confluence WebDAV Client Web Folder in Windows XP

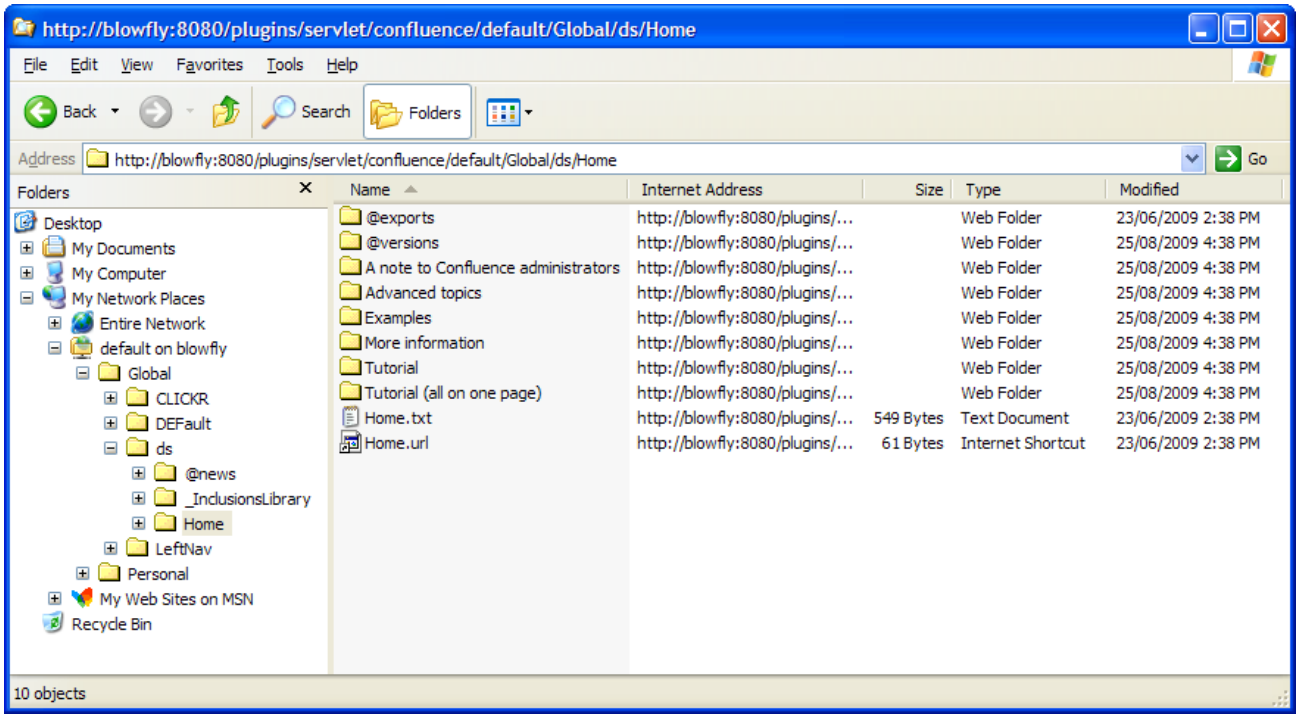

Setting up a WebDAV client in Linux

There are many tools and mechanisms available for configuring WebDAV clients in these operating systems. Therefore, we have chosen to demonstrate this using the file managerKonqueror, which is part of the LinuxK Desktop Environment.

To set up a Confluence WebDAV client in Konqueror:

1. Open Konqueror
2. In the 'Location' field, enter the URL for the Confluence WebDAV location using the 'protocol'`webdavs` (for example,`webdavs://<confluence server url>/confluence/plugins/servlet/confluence/defaultorwebdavs://<confluence server url>/plugins/servlet/confluence/default`) and press`Enter`.

3. Enter your Confluence username and password if prompted

You should be able to click to load many, but not all files. In practice, you would normally save a modified file locally, then drag it to the Konqueror window to upload it to Confluence.

Restricting WebDAV Client Write Access to Confluence

In earlier versions of the WebDAV plugin, separate options for restricting a WebDAV client's write permissions (that is, create/move, edit and delete actions), were available. However, in the current version of this plugin, they have been simplified and combined into a general write permission restriction that covers all of these actions.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

WebDAV clients are now denied write permission to your Confluence installation by setting a regex that matches specific content within the WebDAV client's user agent header. Upon setting a regex, it will be added to a list of restricted WebDAV clients. Any WebDAV clients whose user agent header matches a regex in this list will be denied write permission to your Confluence installation.

Example: A PROPFIND method header generated by a Microsoft Web Folder WebDAV client, showing the user agent header field:

```
PROPFIND /plugins/servlet/confluence/default HTTP/1.1
Content-Language: en-us
Accept-Language: en-us
Content-Type: text/xml
Translate: f
Depth: 1
Content-Length: 489
User-Agent: Microsoft Data Access Internet Publishing Provider DAV
Host: 127.0.0.1:8082
Connection: Keep-Alive
```

Unlike earlier versions of the WebDAV plugin, which could only restrict write permissions for**all**WebDAV clients, the current version of this plugin allows you to restrict write permissions to specific WebDAV clients.

To restrict a WebDAV client's write access permissions to your Confluence installation:

1. Choose the**cog icon** , then choose**General Configuration**
2. Choose '**WebDav Configuration**' in the left panel
3. Enter a regex that matches a specific component of the user agent header sent by the WebDAV client you want to restrict.
4. Click the '**Add new regex**' button
   Repeat steps 3 and 4 to add a regex for each additional WebDAV client you want to restrict.
5. Hit**Save**

To restore one or more restricted WebDAV client's write access permissions to your Confluence installation:

1. Choose the**cog icon** , then choose**General Configuration**
2. Click**WebDav Configuration**under 'Configuration' in the left panel
3. Select the regex(es) from the list that match(es) the user agent header sent by the restricted WebDAV client(s) you want to restore
4. Click the**Remove selected regexes**button
5. Hit**Save**

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Screenshot: WebDAV configuration

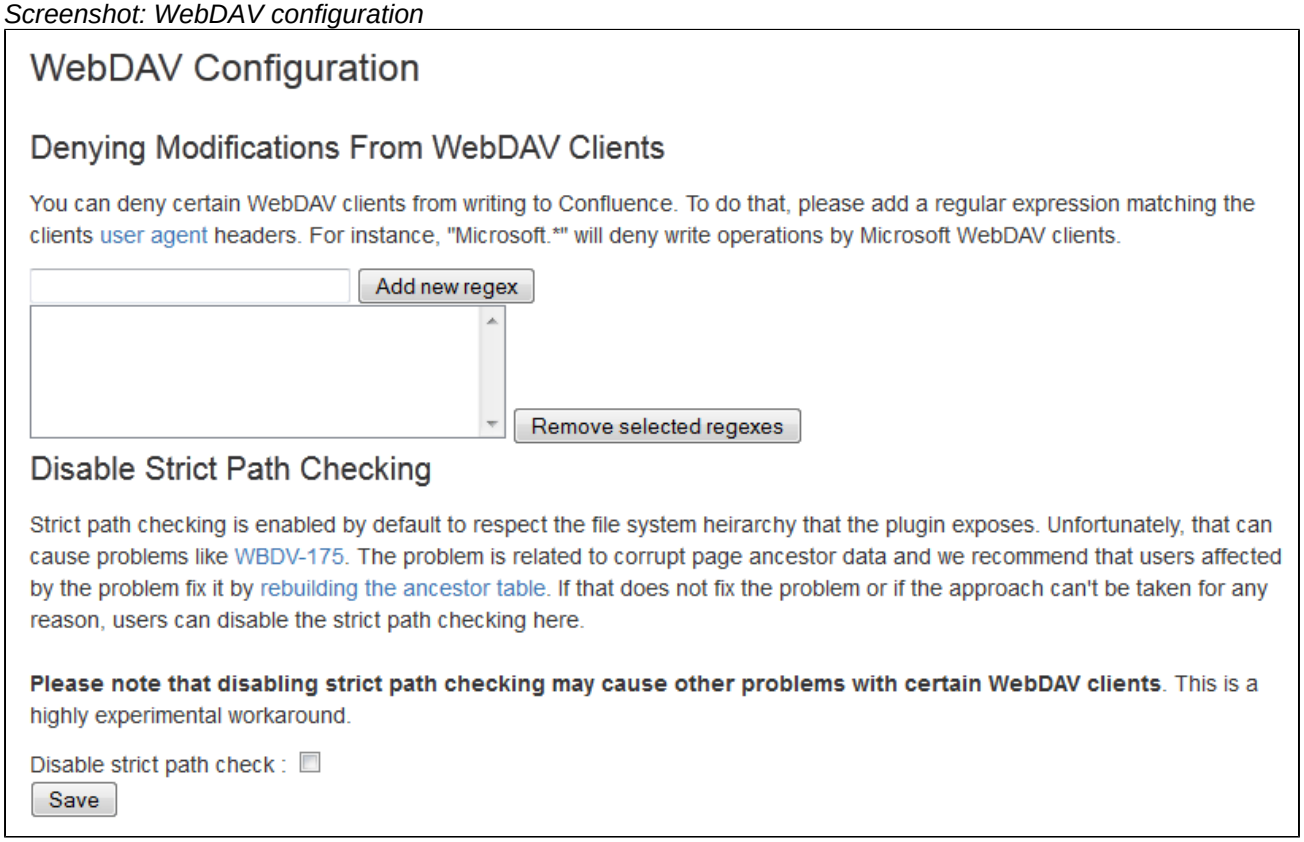

Disabling Strict Path Checking

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 minimize these problems.

To disable the WebDAV plugin's strict path checking option:

1. Choose the**cog icon** , then choose**General Configuration**
2. Click**WebDav Configuration**under 'Configuration' in the left panel
3. Clear the 'Disable strict path check' check box
4. Hit**Save**

Virtual Files and Folders

In the unlikely event that you have problems with the WebDAV client's performance or stability, you can enable access to automatically generated (that is, virtual) files and folders.

Note:

By default, these options are hidden on the 'WebDAV Configuration' page. To make them visible, append the parameter`?hiddenOptionsEnabled=true`to the end of your URL and reload the page. For example:

```
<Confluence base URL>/admin/plugins/webdav/config.action?hiddenOptionsEnabled=true
```


Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Screenshot: The Hidden Virtual Files and Folders Option

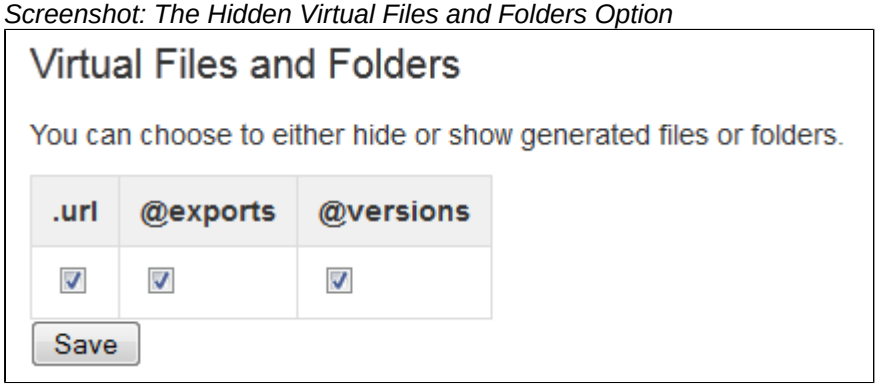

To enable or disable access to virtual files and folders:

1. Choose the**cog icon** , then choose**General Configuration**
2. Click**WebDav Configuration**under 'Configuration' in the left panel
3. Amend your URL as described in the note above and reload the 'WebDav Configuration' page
4. Select or clear the check box options in the 'Virtual Files and Folders' section as required
5. Hit**Save**

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Configuring HTTP Timeout Settings

When macros such as the RSS Macro make HTTP requests to servers which are down, a long timeout value is used. You can set this timeout value through a system parameter to avoid this.

To configure the HTTP Timeout Settings:

1. Choose the**cog icon** , then choose**General Configuration**
2. Select '**General Configuration**' under the 'Configuration' heading in the left-hand panel.
3. Find the '**Connection Timeouts**' section in the lower portion of the screen.
4. Click '**Edit**' to adjust the settings:
    - **Adjust External connections enabled**: This setting allows system administrators to disable external connections so macros like the RSS Macro won't be allowed to make connections to an external server. It provides protection against external servers providing insecure HTML, timing out or causing performance problems. The default setting is 'true'.
    - **Connection Timeout (milliseconds)**: Sets the maximum time for a connection to be established. A value of zero means the timeout is not used. The default setting is ten seconds (10000).
    - **Socket Timeout (milliseconds)**: Sets the default socket timeout (SO_TIMEOUT) in milliseconds, which is the maximum time Confluence will wait for data. A timeout value of zero is interpreted as an infinite timeout. The default setting is ten seconds (10000).

Configuring Number Formats

There are two number format settings in Confluence:

- Long number format. For example:###############
- 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 `[`**`searchterms`**`@google]` instead.

On this page:

- Creating shortcut links
- Using shortcut links
- Deleting shortcut links

Here is a screenshot showing the shortcuts currently defined onhttp://confluence.atlassian.com:

Key	Expanded Value	Default Alias	Operations
cache	http://www.google.com/search?q=cache:		Remove
imdb	http://us.imdb.com/Title?		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**
2. Choose **Shortcut Links** in the left-hand panel.
3. Enter a **Key** for your shortcut. This is the shortcut name a user will use to reference the URL.
4. Enter the **Expanded Value**. This is the URL for the link. You can use '%s' in the URL to specify where the user's input is inserted. If there is no '%s' in the URL, the user's input will be put at the end.
5. Enter a **Default Alias**. This is the text of the link which will be displayed on the page where the shortcut is used, with the user's text being substituted for '%s'.
6. Choose **Submit**.

Using shortcut links

Enter a shortcut link on the **Advanced** tab of the Insert Link dialog. See Links for details.

Specify in the link what should be appended to the end of the shortcut URL, followed by an at-sign (@) and the key of the shortcut. Shortcut names are case-insensitive. So, for example, using the keys shown in the above screenshot:

To link to...	Type this	Resulting URL	Demonstration
a issue	CONF-1000@JIRA	http://jira.atlassian.com/secure/QuickSearch.jspa?searchString=CONF-1000	CONF-1000
a Google search	Atlassian Confluence@Google	http://www.google.com/search?q=Atlassian+Confluence	Atlassian Confluence@Google

Deleting shortcut links

Shortcut links are listed on the **Shortcut Links** tab of the Administration Console. Click **Remove** to delete the shortcut.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Configuring Time and Date Formats

You can change how times and dates appear throughout your Confluence site to suit your organization's preferred date format.

On this page:

- Site date and time
- Relative dates
- Date lozenges

Related pages:

- Choosing a Default Language
- Installing a Language Pack

Site date and time

To change the time and date formats for your entire site:

1. Go to > **General Configuration**.
2. Select **Edit**.
3. Enter your preferred**Time Format**, **Date Time Format** and **Date Format**.
4. Select **Save**.

Confluence uses the Java SimpleDateFormat class. Head toJava SimpleDateFormat to see all allowed values, or use one of the common format examples below. Note: letters are case sensitive.

Format	Example
dd MMMM yyyy	05 June 2019
MMMM d, yyyy	June 5, 2019
d MMM yy	5 Jun 19
dd-MM-yy	05-06-19
d-M-yy	5-6-19
h:mm a	3:25 PM
HH:mm	15:25

Here's what the page history for the same page might look like with different date and time formats.

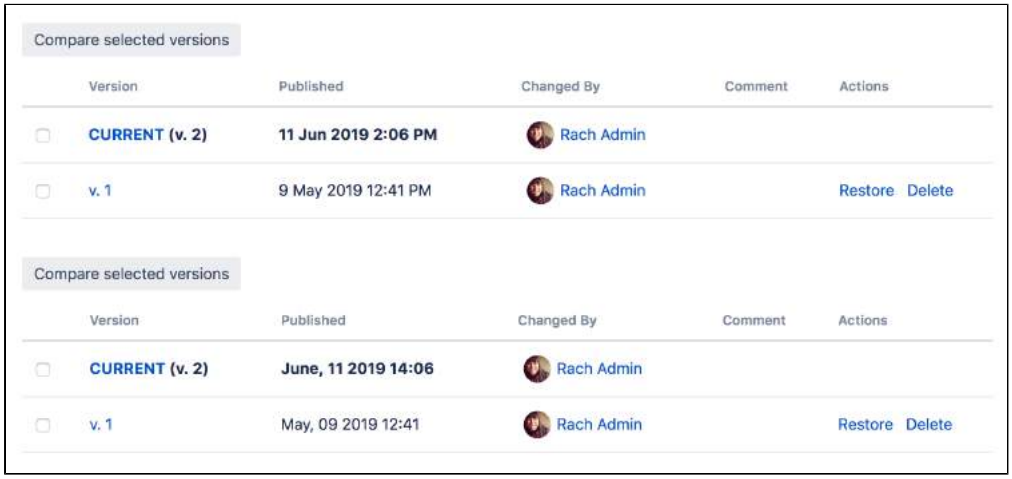

Relative dates

Some parts of Confluence use relative dates when the change happened recently. For example a comment might have been added "yesterday", or a page modified "about 2 hours ago".

It's not possible to customize this format. Full dates are displayed in your preferred format once the change is more than 1 day old.

Date lozenges

To insert a date lozenge, in the editor type ***//*** or select + > **Date** from the toolbar.

The date format displayed will depend on the language settings of the current user, and not the global default language settings. This means the lozenge will appear differently for different users.

The examples below show how the date lozenge will appear based on local language setting.

Tip: To change you local language setting, see Edit Your User Settings. By default, it is set to"Automatically recognize your browser settings".

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

User locale	Java 11 view mode	Java 17 view mode	Editor datepicker
cs-CZ	1. 9. 2022	1. 9. 2022	1. 9. 2022
da-DK	1. sep. 2022	1. sep. 2022	1. sep. 2022
de-DE	01.09.2022	01.09.2022	01.09.2022
et-EE	1. sept 2022	1. sept 2022	1. sept 2022
en-GB	1 Sep 2022	1 Sept 2022	1 Sept 2022
en-US	Sep 1, 2022	Sep 1, 2022	Sep 1, 2022
es-ES	1 sept. 2022	1 sept 2022	1 sept 2022
fr-FR	1 sept. 2022	1 sept. 2022	1 sept. 2022
is-IS	1. sep. 2022	1. sep. 2022	01/09/2022
it-IT	1 set 2022	1 set 2022	1 set 2022
hu-HU	2022. szept. 1.	2022. szept. 1.	2022. szept. 1.
nl-NL	1 sep. 2022	1 sep. 2022	1 sep. 2022
no-NO	1. sep. 2022	1. sep. 2022	1. sep. 2022
pl-PL	1 wrz 2022	1 wrz 2022	1 wrz 2022
pt-BR	1 de set de 2022	1 de set. de 2022	1 de set. de 2022
ro-RO	1 sept. 2022	1 sept. 2022	1 sept. 2022
sk-SK	1. 9. 2022	1. 9. 2022	1. 9. 2022
fi-FI	1.9.2022	1.9.2022	01.09.2022
sv-SE	1 sep. 2022	1 sep. 2022	1 sep. 2022
ru-RU	1 сент. 2022 г.	1 сент. 2022 г.	1 сент. 2022 г.
zh-CN	2022年9月1日	2022年9月1日	2022年9月1日
ja-JP	2022/09/01	2022/09/01	2022/09/01
ko-KR	2022. 9. 1	2022. 9. 1	2022. 09. 01

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Enabling the Remote API

XML-RPC and SOAP remote APIs were deprecated in Confluence 5.5. We recommend using the fully supported Confluence Server REST API wherever possible.

To use the XML-RPC and SOAP remote APIs youneed to enable the APIs from the **Administration Console**. You'll need System Administrator permissions to do this.

To enable the remote API:

1. Choose the**cog icon** , then choose**General Configuration**
2. Click **Further Configuration** in the left-hand panel.
3. Click **Edit**.
4. Click the check box next to **Remote API (XML-RPC & SOAP)**.
5. Click **Save**.

Enabling Threaded Comments

Comments on pages or blog posts are displayed in one of two views:

- **Threaded**: Shows the comments in a hierarchy of responses. Each reply to a comment is indented to indicate the relationships between the comments.
- **Flat**: Displays all the comments in one single list and does not indicate the relationships between comments.

By default, comments are displayed in **threaded** mode. A Confluence Administrator (see Global Permissions Overview) can enable or disable the threaded view for the entire Confluence site.

To enable or disable the threaded view:

1. Choose the**cog icon** , then choose**General Configuration**
2. Select **Further Configuration** in the left-hand panel
3. Choose**Edit**
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.

Confluence is available in these languages right out of the box:

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

The following languages are still bundled, but we no longer translate new features for these languages.

- Eesti (Eesti | Estonia)
- slenska (sland | Iceland)
- Slovenina (Slovensk republika | Slovak Republic)
- Romn (Romnia | Romania)

You can make additional languages available by installing language pack apps.You'll need to be a Confluence administrator to install a language pack.

Related pages:

- Choosing a Default Language
- Configuring Indexing Language
- Installing Marketplace apps

Installing additional language packs

We no longer provide community translations for Confluence. You can find a small number of third-party language packs on theAtlassian Marketplace.

Showing User Interface Key Names for Translation

This feature is useful if you are troubleshooting translations of the Confluence user interface. After opening the Confluence dashboard, you can add the following action to the end of your Confluence URL:

```
?i18ntranslate=on
```

For example`http://myconfluencesite.com?i18ntranslate=on`

This will cause each element of the user interface to display its special **key name**. This makes it easier to find the context for each key within the user interface.

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

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

```
?i18ntranslate=off
```


Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

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
- Confluence Data Center deployed in AWS
- Verifying your settings
- Recognized system properties

Related pages:

- Recognized System Properties
- How to fix out of memory errors by increasing available memory

Linux

To configure System Properties in Linux installations:

1. Edit the`<installation-directory>/bin/setenv.sh`file.
2. Find the section**CATALINA_OPTS=**
   (this is JAVA_OPTS= in Confluence 5.5 and earlier)
3. Refer to the list of parameters inRecognized System Properties.

Add all parameters in a space-separated list, inside the quotations. Make sure to keep the string ${CATALINA_OPTS}" in place.

Windows (starting from .bat file)

To Configure System Properties in Windows Installations When Starting from the .bat File:

1. Edit the`<installation-directory>/bin/setenv.bat`file.
2. Find the section**set CATALINA_OPTS=%CATALINA_OPTS%**
   (this is JAVA_OPTS=%JAVA_OPTS%in Confluence 5.5 and earlier)
3. Refer to the list of parameters inRecognized System Properties.

Add all parameters in a space-separated list. Make sure to keep the string %CATALINA_OPTS% in place.

Windows service

There are two ways to configure system properties when youStart Confluence Automatically on Windows as a Service, either viacommand lineorin the Windows Registry

Setting properties for Windows services via command line

To set properties for Windows services via a command line:

1. Identify the name of the service that Confluence is installed as in Windows (Go to**Control Panel**>**Administrative Tools**>**Services**):

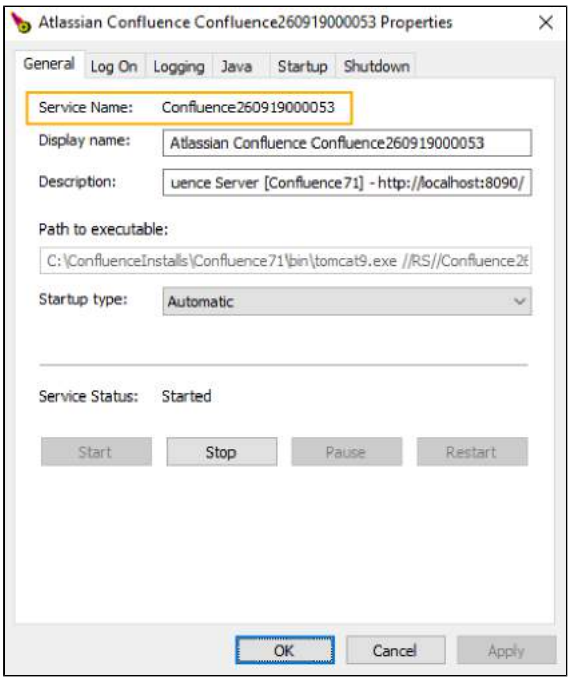

In the above example, the service nameis*Confluence260919000053*.

2. Open the command window (Choose**Start**>**cmd.exe**)
3. `cd`to the`bin`directory of your Confluence instance andrun the following command:

```
tomcat9w //ES//<SERVICENAME>
```

In the above example, it would be`tomcat9w //ES//Confluence`*`260919000053`*
The Tomcat version number may be different if you are using an earlier version of Confluence.

4. Click on the`Java`tab to see the list of current start-up options:
5. Append any new option on its own new line by adding to the end of the existing Java Options. Refer to the list of parametersinRecognized System Properties.

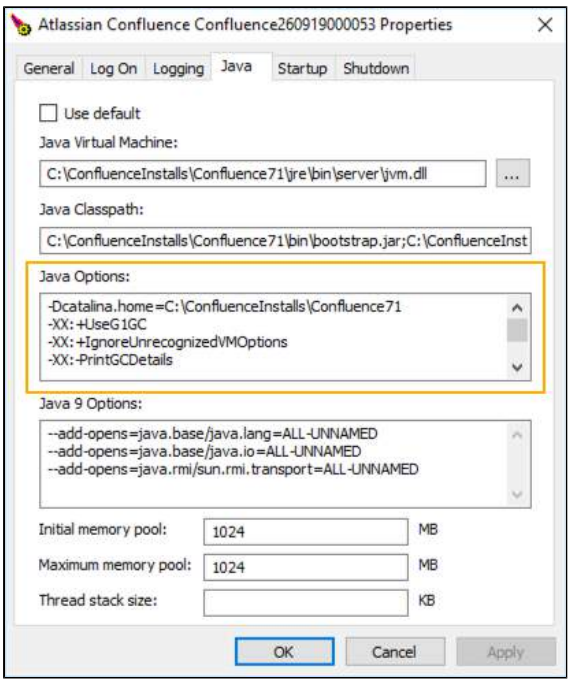

Setting properties for Windows services via the Windows registry

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

In some versions of Windows, there is no option to add Java variables to the service. In these cases, you must add the properties by viewing the option list in the registry.

1. Go to the Registry Editor (**Start**>**regedit.exe**).
2. Find the Services entry:
   64bit: `HKEY_LOCAL_MACHINE >> SOFTWARE >> WOW6432Node >> Apache Software Foundation >> Procrun 2.0 >> Confluence service name`
   32bit: `HKEY_LOCAL_MACHINE >> SOFTWARE >> Apache Software Foundation >> Procrun 2.0 >> Confluence service name`

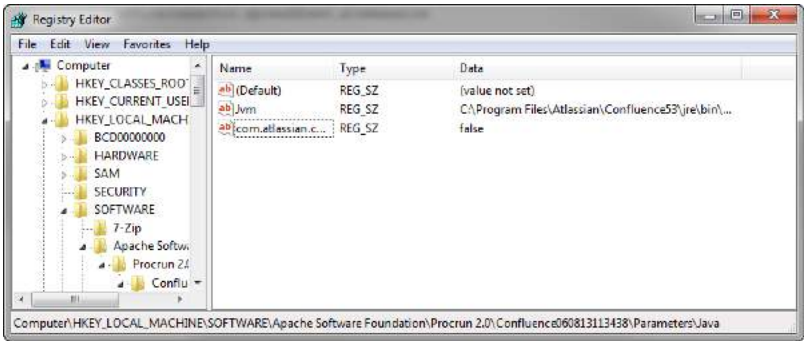

3. To change existing properties double-click the appropriate value.
4. To change additional properties, double-click options.
5. Refer to the list of parametersinRecognized System Properties. Enter each on a separate line.

Confluence Data Center deployed in AWS

If you've used the Quick Start or CloudFormation template to deploy Confluence Data Center in AWS, you will pass system properties via the Cloud Formation Template, and not using the methods described above.

1. In the AWS Console, choose**Update Stack**
2. Under**Advanced**, enter system properties in the Catalina Properties field as follows:

```
-Xms1024m -Xmx1024m -Dsystemproperty=value
```

3. Changes are applied when a new nodes are provisioned.

Verifying your settings

To see what Confluence is using, check Viewing System Properties.

Recognized system properties

SeeRecognized System Propertiesfor the full list of system properties available to your Confluence version.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Recognized System Properties

Confluence supports some configuration and debugging settings that can be enabled through Java system properties. System properties are usually set by passing the `-D`flag to the Java virtual machine in which Confluence is running. See thefull instructions:Configuring System Properties.

Since	Default Value	Effect
atlassian.forceSchemaUpdate		
1.0	`false`	By default, Confluence will only run its database schema update when it detects that it has been upgraded. This flag will force Confluence to perform the schema update on system startup.
confluence.home		
1.0	Any filesystem path	If this system property is set, Confluence will ignore the contents of the `confluence-init.properties` file, and use this property as the setting for the Confluence Home directory.
confluence.dev.mode		
1.0	`false`	Enables additional debugging options that may be of use to Confluence developers (additionally it changes spring bean creation to use lazyinitializationby default to decrease startup time). Do not enable this flag on a production system.
confluence.disable.mailpolling		
2.4	`false`	If set to "true", will prevent Confluence from retrieving mail for archiving within spaces. Manually triggering "check for new mail" via the web UI will still work. This property has no effect on outgoing mail
confluence.i18n.reloadbundles		
1.0	`true`	Setting this property will cause Confluence to reload its i18n resource bundles every time an internationalized string is looked up. This can be useful when testing translations, but will make Confluence run *insanely slowly*.
confluence.ignore.debug.logging		
1.0	`true`	Confluence will normally log a severe error message if it detects that DEBUG level logging is enabled (as DEBUG logging generally causes a significant degradation in system performance). Setting this property will suppress the error message.
confluence.jmx.disabled		
3.0	`false`	If set to "true", will disable Confluence's JMX monitoring. This has the same effect as setting the "enabled" property to false in `WEB-INF/classes/jmxContext.xml`
confluence.optimize.index.modulo		
2.2	`20`	Number of index queue flushes before the index is optimized. This property was removed in Confluence 5.2 when optimize index was removed.
confluence.plugins.bundled.disable		
2.9	`false`	Starts confluence without bundled plugins. May be useful in a development environment to make Confluence start quicker, but since bundled plugins are necessary for some of Confluence's core functionality, this property should not be set on a production system.

atlassian.indexing.contentbody.maxsize		
3.0	1048576	When a file is uploaded, its text is extracted and indexed. This allows people to search for the content of a file, not just the filename. If the amount of content extracted from the file exceeds the limit set by this property (default is 1MB, in bytes), the file's contents will still be indexed and searchable, but will not appear when the file is returned in search results. Increasing this limit may make displaying search results slower. See Configuring Attachment Size for more info.
atlassian.mail.fetchdisabled		
3.5	`false`	Disables mail fetching services for IMAP and POP
atlassian.mail.senddisabled		
3.5	`false`	Disables sending of mail
atlassian.disable.caches		
2.4	`true`	Setting this property will disable conditional get and expires: headers on some web resources. This will significantly slow down the user experience, but is useful in devlopment if you are frequently changing static resources and don't want to continually flush your browser cache.
confluence.html.encode.automatic		
2.9		Setting this property forces the antixss encoding on or off, overriding the behavior dictated by settings. The default behavior differs between Confluence versions.
org.osgi.framework.bootdelegation		
2.10	empty	Comma-separated list of package names to provide from application for OSGi plugins. Typically required when profiling Confluence. For example: "com.jprofiler.**, com.yourkit.**".
confluence.diff.pool.size		
3.1	`20`	Maximum number of concurrent diffs. When that number is exceeded, additional attempts by RSS feeds to create diffs are ignored and logged. (The RSS requests succeed, they are just missing diffs).
confluence.diff.timeout		
3.1	`1000`	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	`30000`	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`	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		

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

3.2	`false`	Setting this property changes Confluence to use the Experimental XML Importer. It is designed to be a more stable implementation but, at the time of the release of 3.2, the importer is largely untested and thus not supported.
atlassian.webresource.disable.minification		
3.3	`false`	Disables automatic minification of JavaScript and CSS resources served by Confluence.
index.queue.thread.count		
3.3	*See "Effect"*	Sets the number of threads to be used for the reindex job. The value has to be in the range of 1 to 50 (inclusive), i.e. at least one thread but no more than 50 threads will be used. There is no default value, i.e. • If you don't set `index.queue.thread.count`, the number of threads to be used are calculated based on the number of objects that need to be reindexed and the number of processors available (a maximum of 50 threads will be used). • If you set `index.queue.thread.count=2`, then two threads will be used to reindex the content (regardless of the number of objects to be reindexed or the number of processors available) • If you set `index.queue.thread.count=200`, then ten threads (the maximum allowed) will be used to reindex the content. Note: For Confluence versions from 3.3 to 5.6 the maximum threads is 10.
index.queue.batch.size		
3.3	1500	Size of batches used by the indexer. Reducing this value will reduce the load that the indexer puts on the system, but indexing takes longer. Increasing this value will cause indexing to be completed faster, but puts a higher load on the system. Normally this setting does not need tuning.
password.confirmation.disabled		
3.4	`false`	This property disables the password confirmation functionality that Confluence uses as an additional security measure. With this property set, Confluence will *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`	Setting this property to "false" disables the detection of browser language headers, effectively restoring Confluence behavior to that of earlier releases. Setting this property to "true" enables the detection of the language headers sent by the browser. Confluence will change the UI language based on the browser headers. See documentation on how users can choose a language preference.
upm.pac.disable		
Universal Plugin Manager 1.5	`false`	When this property is set to true, then UPM will not try to access the The Atlassian Marketplace. This is useful for application servers that do not have access to the Internet. See the UPM documentation.
confluence.reindex.documents.to.pop		
3.5.9	`20`	Indicates how many objects each indexing thread should process at a time during a full re-index. Please note that this number does not include attachments.
confluence.reindex.attachments.to .pop		

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

3.5.9	`10`	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`	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`	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`	Disables the session invalidation on log out. As of 4.0 the default behavior is to invalidate the JSession assigned to a client when they log out. If this is set to true the session is kept active (but the user logged out). This may be valuable when using external authentication systems, but should generally not be needed.
officeconnector.spreadsheet.xlsxmaxsize		
4.0.5	2097152	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	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	`false`	Allows Confluence to start up against the home directory of a newer version of Confluence. Note that running Confluence like that is unsupported. You should only turn this on if you know what you are doing. See After Downgrading, Confluence Will No Longer Run for details.
confluence.skip.reindex		
	false	Generally a full reindex is not required when upgrading Confluence, but there may be some occasions where an upgrade task will kick-off the reindex process. Set this property to true, to skip rebuilding the search index when Confluence is upgraded. This can be useful if you have a very large site and wish to delay rebuilding the index until a time after the upgrade is complete.
reindex.thread.count		
5.2		Sets the number of threads to be used for a one-off reindex job. The value has to be in the range of 1 to 50 (inclusive), i.e. at least one thread but no more than 50 threads will be used. This system property does not affect the incremental indexing that Confluence does. From Confluence 7.14 the default value of this property is calculated based on either the number of available CPUs or free memory, whichever is lower.
reindex.attachments.thread.count		

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

5.2	4	Sets the number of concurrent threads to be used when reindexing attachments, and allows you to reduce the concurrency for these more memory-intensive index items. From Confluence 7.14 the default value of this property is calculated based on either the number of available CPUs or free memory, whichever is lower. Minimum 4 threads.
atlassian.confluence.export.word.max.embedded.images		
5.2	1000	This property limits the number of images that are included when you export a Confluence page to Word. When you export a page with many large images to Word, all the images are loaded into memory, which can then cause out-of-memory errors affecting your entire Confluence site. You can temporarily increase this limit, using this system property, if you need to export a page with lots of images. The default value of this property was increased from 50 to 1000 in Confluence 7.20.2.
confluence.mbox.directory		
5.4.1		Setting this property defines the directory on your Confluence Server where mailboxes can be imported from (for use with the Confluence Mail Archiving system app). Mailboxes are not able to be imported from any other location. We recommend administrators create a directory in the Confluence Home directory specifically for this purpose. Mail cannot be imported from the server until this system property is set.
confluence.search.max.results		
5.5	1000	Setting this property changes the maximum number of items Confluence Search will return.
confluence.upgrade.recovery.file.enabled		
5.5	`true`	By default, Confluence creates an upgrade recovery file before and after an upgrade. The operation can take a long time on large databases and can be safely turned off if there is a process to back up the 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		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	When a file is uploaded, its text is extracted and indexed. This allows people to search for the content of a file, not just the filename. Confluence will only extract content from a Word document up to the limit set by this property (defaults to 16MB, in bytes), before proceeding to index it. This will mean that only part of the file's contents is searchable. The check uses the uncompressed file size, not the compressed size on disk in the case of .docx files. See Configuring Attachment Size for more info.
cluster.login.rememberme.enabled		

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

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		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	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		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		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		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		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	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		

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

5.8.7		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) .
confluence.clickjacking.protection.disable		
5.8.15	false	Security features prevent Confluence from being embedded in an `<iframe>` . To disable this protection, set this property to `true` which will allow Confluence to be embedded in an `<iframe>` .
confluence.cluster.index.recovery.query.timeout		
5.9.1	10	In Confluence Data Center, t he amount of time, in seconds, that a confluence node needing index recovery will wait for another node to offer an index snapshot, before it gives up and stops attempting to recover the index.
confluence.cluster.index.recovery.generation.timeout		
5.9.1	120	In Confluence Data Center, the amount of time, in seconds, that a confluence node needing index recovery will wait for an index snapshot to be created by another node, before it gives up and and stops attempting to recover the index.
confluence.cluster.index.recovery.num.attempts		
5.9.1	1	In Confluence Data Center, the number of times that a node needing index recovery will attempt to recover its index. Set this property to 0 to disable index recovery on that node (for example, when you want to force a node to automatically rebuild its own index on startup).
com.atlassian.confluence.officeconnector.canary.memory_value		
5.9.1	1024	Sets the memory (in megabytes) available to the Office Connector Canary process, which is a workaround for a known issue with the Import from Word option. See JVM crashes during Import from Word in Confluence for more information.
com.atlassian.confluence.officeconnector.canary.timeout		
5.9.1	120000	Sets the maximum timeout (in milliseconds) for the Office Connector Canary process, which is a workaround for a known issue with the Import from Word option. See JVM crashes during Import from Word in Confluence for more information.
atlassian.plugins.enable.wait		
5.9.5	300	Set this property to increase the default time to wait for plugins to start up. This is useful if you have problems with plugins not starting up in time, causing Confluence to fail to start.
confluence.cluster.hazelcast.max.no.heartbeat.seconds		
5.9.7	30	In Confluence Data Center, this sets how long (in seconds) a node can be out of communication with the cluster before it's removed from the cluster. See balancing uptime and data integrity for more info on when you may want to change this setting.
confluence.startup.remigration.disable		
5.10.8	False	Set this property to true if you repeatedly experience issues with Confluence creating a new page version as it attempts to migrate pages containing unmigrated wiki-markup macros each time a plugin is install or updated. See ~~CONFSERVER-37710~~ CLOSED for more details.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

cluster.safety.time.to.live.split.ms		
6.0.0	60000	In Confluence Data Center, this is the amount of time (in milliseconds) that the cluster safety job will wait to allow the nodes to rejoin after a split brain is detected. If the node still can't find the cluster safety number in the cache after this time, the node will panic.
confluence.cph.max.entries		
6.0.0	2000	This is the maximum number of pages that can be copied when you copy a page and its child pages. Increase this if you need to copy a page with more than the default number of child pages.
confluence.cph.batch.size		
6.0.0	10	This is the number of pages copied in each batch when you copy a page and its children. Increase or reduce this number if you experience problems copying a page with many child pages.
synchrony.port (formerly known as reza.port)		
6.0.0	8091	This is the port that Synchrony, the service that powers collaborative editing, runs on. You should only need to change this if port 8091 is not available. From 6.0.4, Confluence Server will accept either `reza.port` or `synchrony.port`.
synchrony.memory.max (formerly reza.memory.max)		
6.0.0	2g	This is the maximum heap size (Xmx) allocated to Synchrony, the service that powers collaborative editing. Change this value if you need to increase or decrease the heap size. From 6.0.4, Confluence Server will accept either `reza.memory.max` or `synchrony.memory.max`. The default value of this property was increased to 2 gigabytes in 7.10.0.
synchrony.stack.space		
6.0.0	2048k	This sets the stack size (Xss) of the Synchrony JVM. Increase if you experience stack overflow errors, or decrease if you experience out of memory errors from Synchrony. This property only applies when Synchrony is managed by Confluence.
synchrony.enable.xhr.fallback		
6.0.0	True	XML HTTP Request (XHR) fallback allows a user, who cannot connect to Synchrony via a WebSocket, to use the Confluence Editor. From Confluence 6.1 this is enabled by default, and only used when a WebSocket connection is not available. You should not need to disable this, unless directed by our support team.
synchrony.database.test.connection.on.checkin		
6.0.0	True	Verifies the connection to the database is valid at every connection checkin. Atlassian Support may suggest you set this property to False if you experience performance issues in sites that have very frequent page edits.
synchrony.proxy.enabled		

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

6.0.0	True	By default, Confluence uses an internal proxy to communicate between the Confluence JVM and Synchrony JVM. See Administering Collaborative Editing for more information. In Confluence 6.0, set this property to `true` to manually enable the internal proxy (useful if you have configured a reverse proxy and want to also use the internal Synchrony proxy). In Confluence 6.1 or later it should **not be necessary** to use this system property, as Confluence intelligently determines when to use the proxy. This property only applies when Synchrony is managed by Confluence. It has no effect on a Synchrony standalone cluster.
synchrony.bind (formerly known as reza.bind)		
6.0.0	0.0.0.0	This is the specific network interface that Synchrony listens on. It is unlikely that you will need to change this when Synchrony is managed by Confluence. In Confluence Data Center, when running a Synchrony standalone cluster this should be set to the same value as`synchrony.cluster.bind`. From 6.0.4, Confluence Server will accept either `reza.bind` or `synchrony.bind`.
synchrony.context.path		
6.0.0	/synchrony	This is the context path for Synchrony. There should be no need to change this in Confluence Server, or when Synchrony is managed by Confluence.
confluence.pdfexport.permits.size		
6.0.0	(number of cores)	This property sets the number of concurrent PDF exports that can be performed. It defaults to the number of cores on your server or cluster node.
confluence.pdfexport.timeout.seconds		
6.0.0	30	This property sets the amount of time (in seconds) a new PDF export request should wait before failing, if the maximum number of concurrent PDF exports (set inconfluence.pdfexport.permits.size) has already been reached.
confluence.flyingpdf.default.characters.per.line		
6.0.3	80	If the total characters in a table column heading exceeds the value of this property, Confluence will automatically adjust the width of the other table columns so that all columns will fit within one page when the page is exported to PDF.
synchrony.host		
6.0.4	127.0.0.1	This is the IP that Confluence uses to connect to Synchrony. It defaults to localhost. Change this if you need to allow Confluence to contact Synchrony via a custom hostname or IP address. This property only applies when Synchrony is managed by Confluence. It has no effect on a Synchrony standalone cluster.
synchrony.proxy.healthcheck.disabled		
6.1.0	false	The Synchrony proxy health check is used to check whether the Synchrony proxy is running and responding to requests. It requires a http connection. If a http connector is not present in your `server.xml` (for example you're using a https or AJP connector) the health check will fail even if the Synchrony proxy is operational. You can use this system property to disable the health check if necessary.
page.index.macro.max.pages		

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

6.1.4	1000	Sets the maximum number of pages that the Page Index macro can display. The Page Index macro can significantly slow down your Confluence instance and cause out of memory errors when used in a space with a large number of pages. If the number of pages in the space exceeds this limit, the Page Index macro will show a page count, and a message that there are too many pages to display. The default value of this property was reduced to 1000 in Confluence 7.11.
atlassian.indexing.attachment.maxsize		
6.2.2	104857600	When a file is uploaded, its text is extracted and indexed. This allows people to search for the content of a file, not just the filename. If the uploaded file is larger than the limit set by this property (default is 100MB, in bytes), text extraction and indexing will be skipped. See Configuring Attachment Size for more info.
officeconnector.excel.extractor.maxlength		
6.2.2	1048576	When a file is uploaded, its text is extracted and indexed. This allows people to search for the content of a file, not just the filename. Confluence will only extract content from an Excel spreadsheet up to the limit set by this property (default is 1MB, in bytes), before proceeding to index it. This will mean that only part of file's contents are searchable. See Configuring Attachment Size for more info.
atlassian.image_filter.transform.max_data_size		
6.2.2	16000000	Applying image effects to large images can cause out of memory errors. We prevent users from applying image effects to images with a data size greater than 16MB. Set this property, in bytes, to reduce the maximum data size. Note: this is the data size, not the size of the file on disk.
atlassian.image_filter.transform.max.pixel		
6.2.2	4000	Applying image effects to large images can cause out of memory errors. We prevent users from applying image effects to images larger than 4000x4000 pixels. Set this property, in pixels, to change the maximum image dimensions.
confluence.collab.edit.user.limit		
6.3.0	12	When collaborative editing is enabled, this sets the maximum number of users that can simultaneously edit a page. Reduce this number if you experience degraded performance when many people are editing.
jobs.limit.per.purge		
6.4.3	2000	The **Purge Old Job Run Details** scheduled job deletes details of old scheduled jobs from the database in batches. Set this property to change the number of records to remove in each batch.
all.jobs.ttl.hours		
6.4.3	2160	By default, the **Purge Old Job Run Details** scheduled job deletes details of successful scheduled jobs older than 90 days (or 2160 hours). Set this property, to change number of hours to keep details of successful jobs in the database.
unsuccessful.jobs.ttl.hours		

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

6.4.3	168	By default, the **Purge Old Job Run Details** scheduled job deletes details of unsuccessful (failed or aborted) scheduled jobs older than 7 days (or 168 hours). Set this property, to change number of hours to keep details of unsuccessful jobs in the database.
hide.system.error.details		
6.5.0	False	Set this property to `true` if you want to hide details on the error screen that appears in the browser when Confluence can't start up. This can be useful for public-facing sites, where you may not want to display details of the problem.
atlassian.recovery.password		
6.6.1		Allows an administrator to start Confluence in recovery mode and specify a temporary administrator password. This is useful if the administrator is locked out of the instance after a site import, or cannot reset their password by other methods. See Restore Passwords To Recover Admin User Rights for more information. This system property must be removed immediatley after the admin account has been recovered or a new admin account created.
confluence.extra.userlister.limit		
6.6.3, 6.7.1	10000	Set this property to change the maximum number of people the User List macro can display. This macro is known to cause out of memory errors when attempting to display a very large number of users.
conversion.sandbox.pool.size (formerly document.conversion.sandbox.pool.size)		
6.10.0	2	Use this property to increase the number of processes (sandboxes) in the external process pool. More processes means more tasks can be executed in parallel, but will consume more memory and CPU resources on each node. This property only applies to Data Center. This property was renamed in Confluence 6.12.
conversion.sandbox.startup.time.limit.secs (formerly document.conversion.sandbox.startup.time.limit.secs)		
6.10.0	30	When a document file is inserted into a page, thumbnails are generated of its contents, so it can be viewed inline and previewed. In Confluence Data Center this is handled by the external process pool . This property sets the amount of time (in seconds) that a process will wait for document conversion to start, before terminating the process. This property only applies to Data Center. This property was renamed in Confluence 6.12.
document.conversion.sandbox.request.time.limit.secs		
6.10.0	30	When a document file is inserted into a page, thumbnails are generated of its contents, so it can be viewed inline and previewed. In Confluence Data Center this is handled by the external process pool . This property sets the amount of time (in seconds) that a process will wait for document conversion to complete, before terminating the process, and marking thumbnail generation for that file as failed. This property only applies to Data Center. From Confluence 7.8.0 this property also applies to the HTML conversion that occurs when you insert a file using the Office Word or Office Excel macro.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

sandbox.termination.tolerance		
6.10.0	5	This property specifies how often a process in the external process pool will check if the request time limit (set in the request.time.limit.secs property above) has been exceeded. It's calculated by dividing the request time limit by the value of this property. For example, if the request time limit is 30 seconds, and the tolerance set in this property is 5, the process will check if the request time limit has been exceeded every 6 seconds. This property only applies to Data Center.
conversion.sandbox.memory.limit.megabytes (formerly document.conversion.sandbox.memory.limit.megabytes)		
6.10.0	512	When a document file is inserted into a page, thumbnails are generated of its contents, so it can be viewed inline and previewed. In Confluence Data Center this is handled by the external process pool . This property limits the amount of heap memory each process in the external process pool can consume. This property only applies to Data Center. This property was renamed in Confluence 6.12.
document.conversion.sandbox.log.level		
6.10.0	INFO	Use this property to change the logging level of document conversion in the external process pool to WARNING, INFO, or FINE.This is useful if you need to troubleshoot a problem with the sandbox. This property only applies to Data Center.
sandbox.error.delay.millis		
6.10.0	50	This property sets how often (in milliseconds) document conversion errors are captured for diagnostic purposes. This property only applies to Data Center.
document.conversion.sandbox.disable		
6.10.0	false	Set this property to true if you don't want to handle document conversion (thumbnail generation) in the external process pool . When disabled, document conversion will be handled in the Confluence JVM, as is the case in Confluence Server. This property only applies to Data Center.
conversion.sandbox.java.options		
6.10.0		When a document or image file is inserted into a page, thumbnails are generated of its contents, so it can be viewed inline and previewed. In Confluence Data Center this is handled by theexternal process pool. Use this property to override the default value of any of the following Confluence Server system properties (introduced in 7.0.1), and pass new values directly to the JVMs in the external process pool: • confluence.document.conversion.imaging.convert.timeout • confluence.document.conversion.imaging.convert.timeout • confluence.document.conversion.imaging.enabled.tif • confluence.document.conversion.imaging.enabled.psd

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

diagnostics.os.check.period.secs		
6.11.0	120	Set this property to change how often operating system diagnostics checks should be performed (in seconds). This property only applies to the Low free memory (OS-1001) and Low free disk space (OS-1002) alerts.
diagnostics.os.min.free.memory.megabytes		
6.11.0	256	This property applies to the free memory diagnostic alert (OS-1001). Set this property to change the threshold at which the amount of free memory (in megabytes) should trigger this alert.
diagnostics.os.min.free.disk.space.megabytes		
6.11.0	8192	This property applies to the free disk space diagnostic alert (OS-1002). Set this property to change the threshold at which the amount of free disk space (in megabytes) in the local or shared home directory should trigger this alert.
diagnostics.slow.http.request.secs		
6.11.0	60	This property applies to the HTTP request diagnostic alert (HTTP-1001). This alert is disabled by default. Set this property to change the threshold (in seconds) at which a slow HTTP request should trigger this alert.
diagnostics.slow.long.running.task.secs		
6.11.0	300	This property applies to the long running task diagnostic alert (JOB-1001). Set this property to change the threshold (in seconds) at which a long running task should trigger this alert.
diagnostics.slow.macro.rendering.secs		
6.11.0	30	This property applies to the macro rendering diagnostic alert (MACRO-1001). This alert is disabled by default. Set this property to change the threshold (in seconds) at which rendering a macro on a page should trigger this alert.
diagnostics.jvm.memory.check.period.secs		
6.11.0	10	Set this property to change how often JVM diagnostics checks should be performed (in seconds). This property only applies to the Thread memory allocation rate (JVM-1001) and Garbage collection (JVM-1002)alerts.
diagnostics.jvm.memory.allocation.rate.percent		
6.11.0	5	This property applies to the thread memory allocation rate diagnostic alert (JVM-1001). This alert is disabled by default. Set this property to change the threshold (as a percentage) at which the memory allocation to a particular thread, during the monitoring period (set in diagnostics.jvm.memory.allocation.monitoring.period.secs), should trigger this alert.
diagnostics.jvm.memory.allocation.monitoring.period.secs		

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

6.11.0	20	This property applies to the thread memory allocation rate diagnostic alert (JVM-1001). This alert is disabled by default. Set this property to change the monitoring period (in seconds) for this alert.
diagnostics.jvm.garbage.collector.percent		
6.11.0	10	This property applies to the garbage collection diagnostic alert (JVM-1002). This alert checks how much memory has been allocated to garbage collection during the monitoring period (set in diagnostics.jvm.garbage.collector.monitoring.period.secs). Set this property to change the threshold (as a percentage) at which the memory allocated to garbage collection should trigger this alert.
diagnostics.jvm.garbage.collector.monitoring.period.secs		
6.11.0	20	This property applies to the garbage collection diagnostic alert (JVM-1002). Set this property to change the monitoring period (in seconds) for this alert.
diagnostics.alert.retention.period.days		
6.11.0	30	Set this property to change how often diagnostic alert data should be retained in the database (in days).
diagnostics.alert.truncation.interval.minutes		
6.11.0	30	Set this property to change how often we check for, and remove diagnostic alert dat a that is older than 30 days (the limit set by diagnostics.alert.retention.period.days)
pdf.export.sandbox.disable		
6.12.0	false	Set this property to true if you don't want to handle PDF exports in the external process pool. When disabled, PDF exports will be handled in the Confluence JVM, as is the case in Confluence Server. This property only applies to Data Center.
pdf.export.sandbox.request.time.limit.secs		
6.12.0	180	When you export a space to PDF, Confluence exports the content of each page to HTML, converts that HTML to PDF, and then finally merges all the pages together into a single PDF file.In Confluence Data Center this is handled by the external process pool . This property sets the amount of time (in seconds) that a process should wait to complete, before being terminated. This time limit applies both to the time to convert the content from HTML to PDF, and the time to merge the final PDF file. This property only applies to Data Center.
pdf.export.sandbox.memory.requirement.megabytes		
6.12.0	64	In Confluence Data Center PDF exports are handled by the external process pool . This property sets the minimum memory (in megabytes) that a sandbox process must have available to start a PDF export. Ifconversion.sandbox.memory.limit. megabytesis set to less than the value of this property, PDF export will not start. We don't recommend setting this property to less than 64MB. This property only applies to Data Center.
synchrony.eviction.soft.job.threshold.hours		

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

7.0.1	72	This property changes the behavior of the Synchrony softeviction scheduled job. It sets the minimum time, in hours, since a Synchrony change log was last modified, to make it eligible to be cleaned up. By default, only Synchrony change logs last modified more than 3 days ago, for pages/blogs that do not have an active editing session, will be evicted.
synchrony.eviction.hard.job.threshold.hours		
7.0.1	360	This property changes the behavior of the Synchrony hardeviction scheduled job It sets the minimum age, in hours, of content eligible to be evicted by the hard eviction scheduled job. By default, all Synchrony data for any content that is 15 days or older will be evicted by this job, regardlessof whether it has been modified more recently.
confluence.document.conversion.imaging.convert.timeout		
7.0.1	30	When a complex imagefile (such as ICO, EMF, WMF, plus TIF and PSD if enabled) is inserted into a page, thumbnails are generated of its contents, so it can be viewed inline and previewed. This process is known to cause out of memory errors for large or complex files. This property sets the amount of time (in seconds) that Confluence will wait for conversion to complete for an image file, before terminating the process. This property applies to Confluence Server and Data Center.For Data Center also seeDocument conversion for Confluence Data Center
confluence.document.conversion.slides.convert.timeout		
7.0.1	30	When a presentationfile (such as PPT, PPTX, POT)is inserted into a page, thumbnails are generated of its contents, so it can be viewed inline and previewed. This process is known to cause out of memory errors for large or complex files. This property sets the amount of time (in seconds) that Confluence will wait for conversion to complete for a presentation file, before terminating the process. This property applies to Confluence Serve and Data Center.For Data Center seeDocument conversion for Confluence Data Center.
confluence.document.conversion.imaging.enabled.tif		
7.0.1	false	When a file is inserted into a page, thumbnails are generated of its contents, so it can be viewed inline and previewed. By default thumbnails are not generated for TIFF / TIF as they're known to cause out of memory errors. Set this property to 'true' to turn on thumbnail generation for TIFF files. If enabled, a timeout will apply to this type of file. This is set by theconfluence.document.conversion.imaging.convert.timeout system property. This property applies to Confluence Server and Data Center.For Data Center seeDocument conversion for Confluence Data Center
confluence.document.conversion.imaging.enabled.psd		

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

7.0.1	false	When a file is inserted into a page, thumbnails are generated of its contents, so it can be viewed inline and previewed. By default thumbnails are not generated for Photoshop PSD files as they're known to cause out of memory errors. Set this property to 'true' to turn on thumbnail generation for PSD files. If enabled, a timeout will apply to this type of file. This is set by theconfluence.document.conversion.imaging.convert.timeout system property. This property applies to Confluence Server and Data Center.For Data Center seeDocument conversion for Confluence Data Center.
officeconnector.powerpoint.extractor.maxlength		
7.0.1	1048576	When a file is uploaded, its text is extracted and indexed. This allows people to search for the content of a file, not just the filename. Confluence will only extract content from a PowerPoint presentation up to the limit set by this property (default is 1MB, in bytes), before proceeding to index it. This will mean that only part of file's contents are searchable. See Configuring Attachment Size for more info.
confluence.chart.macro.width.max		
7.2.0	3000	Maximum width, in pixels, that a chart macro can be set to display on a page. If a higher value is entered, the chart will automatically be reduced to this default.
confluence.chart.macro.height.max		
7.2.0	3000	Maximum height, in pixels, that a chart macro can be set to display on a page. If a higher value is entered, the chart will automatically be reduced to this default.
confluence.search.lucene.termFilterBitSetThreshold		
7.2.0	20	Set this property to change the behavior of the term filter in Confluence's Lucene based implementation of search. A bitset is only created when the number of matched documents times the threshold set in this property exceeds total number of documents. This will reduce memory usage and improve Confluence performance. You shouldn't need to change this threshold.
page-tree.partial-loading-batch-size		
7.3.0	200	Confluence limits the number of pages that initially display at each level of the page tree. This helps safeguard the performance of your site. Set this property to increase or decrease the number of pages to initially display at each level of the page hierarchy. At least one child page is always displayed, so if you set the value to 10 for example, you may see 11 or 12 pages.
page-tree.partial-loading.disable		
7.3.0	false	Confluence limits the number of pages that initially display at each level of the page tree. This helps safeguard the performance of your site. Set this property to`true` if you always want all pages to be displayed by default in the page tree.
confluence.word.import.maxsize		
7.3.3	20	Sets the maximum uncompressed size, in megabtyes, of a Microsoft Word document that can be imported into Confluence using the Import from Word feature. This is to prevent very large files from causing out of memory errors.
gatekeeper.request-timeout.seconds		

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

7.4.0	60	Some open-ended queries can take a long time to display when you Inspect Permissions in a Confluence Data Center site with a lot of spaces and users. To avoid this having an impact on your site, we have a timeout. Set this property to change the length of the timeout, in seconds. You can't set this lower than 30 seconds, or higher than 900 seconds (15 minutes).
plugin.audit.log.view.sysadmin.only (formerly audit.log.view.sysadmin.only)		
7.5.0	false	Set this property to `true` to restrict the Audit Log page in Confluence Administration to people with System Administrator global permissions. By default people with System Administrator or Confluence Administrator global permission can access this feature. This property applies to Confluence Server and Data Center. This property was renamed in 7.7.0.
plugin.audit.file.max.file.count		
7.5.0	100	In Confluence Data Center, audit events are written to an audit log file in the local home directory. These log files are rotated once they reach the file size set by plugin.audit.file.max.file.size. Set this property to change the total number of audit log files to keep in the file system.
plugin.audit.file.max.file.size		
7.5.0	100	In Confluence Data Center, audit events are written to an audit log file in the local home directory. Set this property to change the maximum file size, in megabytes, that a log file can reach before the log is rotated.
legacy.audit.migrator.num.threads		
7.5.0		Sets the number of threads to be used for migrating existing audit log events to the new audit log format when upgrading to Confluence 7.5 or later. There is no default value. When not set, Confluence will use 2 times the number of processors available. To increase throughput, you can set this property to use a specific number of threads. For example, if you set this property to 16, then 16 threads will be used to migrate the audit log events, regardless of the number of processors available. Note that more threads will result in higher database utilisation.
legacy.audit.migrator.batch.size		
7.5.0	1000	Sets the number of audit log events to migrate to the new audit log format, when upgrading to Confluence 7.5 or later. This migration can take a long time in large sites. When testing your upgrade, if you find the migration is causing performance issue on your Confluence and/or database server, you can set this property to decrease the number of records processed in each batch.
NumberItemPerPageOfPaginatedListAction		
7.5.0	30	Sets the number of items to be listed on the Undefined Pages and Restricted Pages tabs in Space Tools before Confluence paginates the results. Set this property to change the number of items to display per page.
CachingEnablingItemNumber		

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

7.5.0	1000	The Undefined Pages tab in Space Tools provides a list undefined links to pages that do not yet exist. As this list can be quite memory intensive to generate, the results are cached. Set this property to change the maximum number of undefined link references to cache.
confluence.child-macro.page-limit		
7.5.0	(For example, 4)	The Children Display macro may load slowly in spaces with a large number of pages, or complex permissions. Set this property to specify the maximum number of top-level child pages that can be displayed by the macro. Users can still configure the macro to set a lower value for the Number of Children parameter. This property should not be required from Confluence 7.17 if you are using the faster permissions service.
confluence.child-macro.max-depth		
7.5.0	(For example, 2)	The Children Display macro may load slowly in spaces with a large number of pages, or complex permissions. Set this property to specify the maximum depth of descendants that can be displayed by the macro. Users can still configure the macro to set a lower value for the Depth of Descendants parameter. This property should not be required from Confluence 7.17 if you are using the faster permissions service.
confluence.child-macro.disable-excerpt		
7.5.0	false	The Children Display macro may load slowly in spaces with a large number of pages, or complex permissions. Set this property to true tonever include excerpts. Users can still select a value for the Excerpt Display parameter, but excerpts will not be displayed. This property should not be required from Confluence 7.17 if you are using the faster permissions service.
CachingEnablingItemTimeout		
7.5.0	5	The Undefined Pages tab in Space Tools provides a list undefined links to pages that do not yet exist. As this list can be quite memory intensive to generate, the results are cached. Set this property to change the amount of time, in minutes, the results should be cached (time to live).
confluence.event.duration_checker.threshold_in_seconds		
7.6.2	60	When asynchronous events are generated faster than Confluence can process them, we write a message to the application log to advise that we'll process each task synchronously until the queue is cleared. Set this property to change how often this message appears, in seconds. This property is also available in the 7.4 Long Term Support release from 7.4.4.
confluence.mailserver.tls.hostname.verification.disabled		

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

7.6.2	false	Set this property to `true` if you need to disable TLS hostname verification for any SMTP mail servers you've configured in Confluence. This is not recommended as it can increase the risk of a Man In The Middle attack. This property is also available in Long Term Support releases from 6.13.17 and 7.4.5.
atlassian.image_filter.transform.max.pixel.drop_shadow		
7.7.0	2000	Applying a drop shadow image effect to large images can cause out of memory errors. We prevent users from applying a drop shadow to images larger than 2000x2000 pixels, or the value set byatlassian.image_filter.transform.max.pixel (whichever is smallest). Set this property, in pixels, to change the maximum image dimensions for drop shadow. This property is also available in the 7.4 Long Term Support release from 7.4.4.
pagePropertiesReportContentRetrieverMaxResult		
7.7.0	3000	The Page Properties Report macro displays a maximum of 3000 pages. Set this property to decrease the maximum number of pages the macro can display. We dont recommend you increase this limit, as it can have a performance impact on your site. This property is also available in the 7.4 Long Term Support release from 7.4.6.
confluence.webhooks.allow.all.hosts		
7.7.0	false	Set this property to true allow administrators to configure localhost URLs as a webhook endpoint. This is disabled by default for security reasons because all ports on the same network can be pinged by the UI.
plugin.audit.db.limit.rows		
7.7.0	10000000	The audit log has a hard limit of 10 million records. This is to avoid performance problems and your database from growing too large. Set this property to decrease or temporarily increase the hard limit. If you need to keep more events, the audit log can also write log events to a file. This property only applies to Data Center.
audit.log.search.disabled		
7.8.0	false	Set this property to `true` if you do not want to log the Search performed event in the audit log. This event is included by default when the End User Activity coverage area is set to Full, and records the search terms entered in search, advanced search, and in search macros (such as Livesearch and Page Tree Search). This property only applies to Data Center.
document.conversion.sandbox.memory.requirement.megabytes		
7.8.0	128	When a Word or Excel document file is inserted into a page using the Office Word or Office Excel macro, its contents are converted to a format that can be viewed inline and previewed. In Confluence Data Center this is handled by the external process pool. This property limits the amount of heap memory, in megabytes, this process in the external process pool can consume. This property only applies to Data Center.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

atlassian.pats.enabled		
7.9.0	true	Personal access tokens are used to authenticate REST API requests. By default, any user can create a personal access token. Set this property to `false` to prevent users from creating tokens. Any existing tokens can't be used to authenticate while this property is set to `false`.
atlassian.pats.eternal.tokens.enabled		
7.9.0	true	Personal access tokens are used to authenticate REST API requests.Set this property `false` to prevent users from creating tokens that never expire.
atlassian.pats.max.tokens.expiry.days		
7.9.0	365	Personal access tokens are used to authenticate REST API requests. Set this property, in days, to define the maximum days to until expiry.
atlassian.pats.max.tokens.per.user		
7.9.0	10	Personal access tokens are used to authenticate REST API requests. Set this property to limit the number of tokens an individual user can create.
atlassian.allow.insecure.url.parameter.login		
7.10.0	False	Set this property to `true` to use the `os_username+os_password` query parameter to log in to Confluence. This login method is blocked by default, and we strongly recommend you use a more secure method, such as personal access tokens.
com.atlassian.confluence.extra.calendar3.concurrent.task.max		
7.11.0	20	The max number of worker threads used by Team Calendars. This property only applies to Data Center.
com.atlassian.confluence.extra.calendar3.greenhopper.sprint.enabled		
7.11.0	true	If set to false integrating calendars with Jira sprints is disabled. This property only applies to Data Center.
com.atlassian.confluence.extra.calendar3.jira.timeout.socket		
7.11.0	10000	Timeout in milliseconds for calendars to connect to Jira. Increasing this may help if you experience timeouts on Jira calendars. This property only applies to Data Center.
com.atlassian.confluence.extra.calendar3.jira.issues.max		
7.11.0	1000	The maximum number of events that will be loaded from a Jira calendar. Increasing this may cause performance issues. This property only applies to Data Center.
com.atlassian.confluence.extra.calendar3.display.events.dashboard.maxperday		
7.11.0	10	The maximum number of calendar events that will be shown per day on the upcoming events view on the Confluence dashboard. This property only applies to Data Center.
com.atlassian.confluence.extra.calendar3.display.events.calendar.maxpercalendar		

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

7.11.0	200	The maximum number of events that can be displayed from a single calendar. This property only applies to Data Center.
com.atlassian.confluence.extra.calendar3.display.events.calendar.maxperdaysummary		
7.11.0	3	The maximum number of events that will be shown per day in the summary email.
com.atlassian.confluence.extra.calendar3.display.events.calendar.maxdailysummary		
7.11.0	4	The maximum number of calendar events that will be shown in total in the daily summary email. This property only applies to Data Center.
com.atlassian.confluence.extra.calendar3.display.events.calendar.maxweeklysummary		
7.11.0	8	The maximum number of calendar events that will be shown in total in the weekly summary email. This property only applies to Data Center.
com.atlassian.confluence.extra.calendar3.display.timeline.calendar.maxmonth		
7.11.0	6	The maximum number of months that will be shown in the timeline view of a calendar. This property only applies to Data Center.
plugin.data.pipeline.embedded.line.break.preserve		
7.12.0	false	Specifies whether embedded line breaks should be preserved in the output files. Line breaks can be problematic for some tools such as Hadoop. This property is set to`False`by default, which means that line breaks are escaped.
plugin.data.pipeline.embedded.line.break.escape.char		
7.12.0	\\n	Escaping character for embedded line breaks. By default, we'll print`\n`for every embedded line break.
reindex.attachments.freeMbPerThread		
7.14.0	500	Confluence calculates the number of threads to be used when reindexing attachments based on either the number of available CPUs or free memory, whichever is lower. See`reindex.attachments.thread.count` for more details. Set this property to change the amount of free memory per thread (in megabytes) used in this calculation. Reducing this value may result in more threads being used during reindexing, which allows for greater concurrency.
atlassian.oauth2.provider.enable.access.tokens		
7.17.0	true	This property changes theOAuth 2.0 behavior when linking to another application. Set this property to disable the ability to authenticate using access tokens on that node.
atlassian.oauth2.provider.skip.base.url.https.requirement		
7.17.0	false	This property changes theOAuth 2.0 behavior when linking to another application. Set this property to disable the HTTPS requirement for the base URL. When disabled, the OAuth 2.0 provider will be enabled even if the product is using HTTP.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

atlassian.oauth2.provider.skip.redirect.url.https.requirement		
7.17.0	false	This property changes theOAuth 2.0 behavior when linking to another application. Set this property to disable the HTTPS requirement for the Redirect URL. When disabled, the OAuth 2.0 provider will allow Redirect URLs using HTTP.
atlassian.oauth2.provider.max.lock.timeout.seconds		
7.17.0	10	This property changes theOAuth 2.0 behavior when linking to another application. Set this property to change the amount of time, in seconds, a request will await lock access before timing out.
atlassian.oauth2.provider.max.client.delay.seconds		
7.17.0	10	This property changes theOAuth 2.0 behavior when linking to another application. Set this property to change the maximum lifetime of authorization codes, in seconds. The limit is 600 seconds.
atlassian.oauth2.provider.prune.expired.authorizations.schedule		
7.17.0	* * * * * ?	This property changes theOAuth 2.0 behavior when linking to another application. Set this property to change the cron expression for the job that removes expired authorization codes. Default is 1 minute.
atlassian.oauth2.provider.access.token.expiration.seconds		
7.17.0	3600 (1 hour)	This property changes theOAuth 2.0 behavior when linking to another application. Set this property to change the maximum lifetime of access tokens, in seconds.
atlassian.oauth2.provider.prune.expired.tokens.schedule		
7.17.0	* * * * * ?	This property changes theOAuth 2.0 behavior when linking to another application. Set this property to change the cron expression for the job that removes expired access tokens. Default is 1 minute.
atlassian.oauth2.provider.access.token.expiration.seconds		
7.17.0	7776000	This property changes theOAuth 2.0 behaviour when linking to another application. Set this property to change the maximum lifetime of refresh tokens, in seconds. Default is 90 days (in seconds).
atlassian.oauth2.provider.invalidate.session.enabled		
7.17.0	true	This property changes theOAuth 2.0 behaviour when linking to another application. Set this property to change whether to invalidate a session after a successful authentication using an OAuth token.
atlassian.oauth2.provider.validate.client.secret		
7.17.0	true	This property changes theOAuth 2.0 behaviour when linking to another application. Set this property to change whether to validate the client ID and client secret when revoking and creating tokens.
atlassian.oauth2.provider.use.quotes.in.sql		

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

7.17.0	false	This property changes theOAuth 2.0 behavior when linking to another application. Set this property to change whether to add quotes to SQL statements. This is a sanity system property used for database requirements. PostgreSQL will always use quotes unless the do.not.use.quotes.in.sql property (below) is enabled.
atlassian.oauth2.provider.do.not.use.quotes.in.sql		
7.17.0	false	This property changes theOAuth 2.0 behavior when linking to another application. Set this property to change whether toadd quotes to SQL statements. This is a sanity system property used for database requirements.
atlassian.oauth2.provider.token.via.basic.authentication		
7.17.0	true	This property changes theOAuth 2.0 behavior when linking to another application. Set this property to change whether to extract tokens through the basic authentication password field for access token authentication.
confluence.plugins.collab.feedback.destination.folder		
7.20.0	<shared-home> /collab-data	Feedback reports are used to troubleshoot collaborative editing problems. Set this property to change where collaborative editing feedback reports should be saved. For non-clustered installations, the default location will be <local-home>/shared-home/collab-data.
confluence.plugins.collab.feedback.files.max		
7.20.0	200	Feedback reports are used to troubleshoot collaborative editing problems. Set this property to limit the number ofcollaborative editing feedback reports to store.
confluence.plugins.collab.feedback.operation.timeout.sec		
7.20.0	20	Feedback reports are used to troubleshoot collaborative editing problems. Set this property, in seconds, to change the operation timeout when creating a feedback report.
confluence.plugins.collab.feedback.concurrent.max		
7.20.0	5	Feedback reports are used to troubleshoot collaborative editing problems. Set this property to change the maximum number of reports that can be generated concurrently.
confluence.plugins.collab.feedback.operation.format		
7.20.0	csv	Feedback reports are used to troubleshoot collaborative editing problems. Set this property to change the format of the exported file. Can be csv or json.
confluence.plugins.collab.feedback.cleanup.threshold.hours		
7.20.0	120	Feedback reports are used to troubleshoot collaborative editing problems. Set this property, in hours, to change the length of time to retain reports. A scheduled job will delete any reports exceeding this limit, starting with the oldest.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

confluence.temp-files.ttl-in-hours		
7.20.0	24	This property changes the TTL(time-to-live in hours) of sub directories and files under the temp directory. Sub directories and files whose last modification time are earlier than this value will be cleaned during the scheduled temp directory job.
space.trash.content.pagination.size.max		
8.0.0	500	This is the maximum number of trashed items that can be viewed per page when you paginate the trash can for a space. Increase this if you need to increase trash page maximum size with more than the default 500 items per page.
synchrony.cluster.multicast.port		
8.0.0	54328	In a cluster, this property can be used to override the default port that synchrony multicast will listen to. It defaults to 54328.
confluence.filestore.attachments.s3.bucket.name		
8.1.0		This property is related to configuring S3 object storage. By setting this property and `confluence.filestore.attachments.s3.bucket.region`, Confluence will store attachment data in the specified bucket.
confluence.filestore.attachments.s3.bucket.region		
8.1.0		This property is related to configuring S3 object storage. By setting this property and `confluence.filestore.attachments.s3.bucket.name`, Confluence will store attachment data in the specified bucket.
synchrony.service.authtoken		
8.1.0	autogenerated string 32 characters long	This Confluence token is used to access Synchrony Handshaking REST API. This property is also available in Long Term Support releases from 7.13.14 and 7.19.6. It has also been backported to 8.0.1 and later.
http.header.security.disabled		
8.1.1	false	Setting this property will disable Tomcats `HttpHeaderSecurityFilter.` Header defaults for XSS, anti-click-jacking and HSTS are provided by this filter. Disabling this filter will disable HSTS.
http.header.security.hsts.max.age		
8.1.1	31536000	If HSTS is enabled, this will set the max-age in seconds for the header.
http.header.security.hsts.preload.enabled		
8.1.1	false	If HSTS is enabled, this will enable HSTS pre-loading.
http.header.security.hsts.include.subdomains		
8.1.1	false	If HSTS is enabled, this will enable HSTS for all subdomains.
macro.required.velocity.context.keys		

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

8.2.0	generalUtil, bootstrap	This property changes the velocity context items that are available for user macros. The property is a string of keys that can be found in Default Velocity context, split by a comma ",". If a user macro requires velocity context items on top of the default ones, admins must add the required items into this property before starting up Confluence. This property is also available in Long Term Support releases from 7.13.15 and 7.19.7.
com.ctc.wstx.inputBufferLength		
8.3.0	64000	This property is the size of input buffer (in chars) to use for reading XML content from the Woodstox input stream/reader. The Woodstox input reader is used as part of the process that marshals and unmarshals editor pages as XHTML. Previously, the limited buffer caused issues with macros containing large Javascript elements, which caused the contents to be broken into multiple CDATA sections.
confluence.clusterEvent.timeout		
8.3.0	10 (seconds)	This property defines the amount of time (in seconds) that a cluster node that publishes a cluster event will wait for the other nodes to acknowledge that theyve forwarded the event. This is the same timeout that the old mechanism uses, but there the timeout wasnt configurable like it is now.
confluence.show.setup.restore		
8.3.0	false	Setting this property to 'true' will display the option to restore from backup in the Confluence setup wizard. Since Confluence 8.3, we've changed the way we do backup and restore. We don't recommend this legacy option for restoring content to a new site because it contains issues resolved by the new system.
confluence.backuprestore.backup.ttl-in-hours		
8.3.0	72 (hours)	This property sets how long we keep XML backups in the restore directory before deleting them. Ensure this value is less than `confluence.backuprestore.jobs.ttl-in-days` to avoid accidental file deletion. For example, if this value is set to greater than 14 days (336 hours), the XML backup will be deleted when the scheduled job 'Backup Restore v2 trigger job clean up' runs at 14 days by default. This property relates to the scheduled job, 'Backup Restore v2 temporary backup zip cleaner'.
confluence.backuprestore.jobs.ttl-in-days		
8.4.0	14 (days)	This property sets the amount of time before we delete backup and restore jobs (metadata and files) from the database. We count the number of days since the job has finished, not when it was created. Ensure this value is greater than `confluence.backuprestore.jobs.ttl-in-hours` (72 hours by default) to avoid accidental file deletion.
confluence.pdfexport.allow.local.hosts		

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

8.4.0	false	This property affects space and site PDF layouts and PDF stylesheets. Setting the property as true means that Confluence will accept a localhost URL. For example, these URLs would be accepted if the property is set to true: • image src="http://localhost/path/to/img.png, or • image src="http://127.0.0.1/path/to/img.pngor • image src="http://[::1]/path/to/img.png. This property is also available in Long Term Support releases from 7.13.18 and 7.19.10.
emoticon.thumbnail.generator.permits.size		
8.4.0	the maximum number of processors available to the virtual machine; never smaller than one	This property sets the number of emoji resizes that are allowed to be processed concurrently. For example, if this value is 1 and 10 people upload emojis at the same time that need resizing, they will wait longer. If the value is set at 10, all requests will be processed and served immediately.
confluence.emoticons.max.file.size		
8.4.0	1 (MB)	This property controls the maximum file size for upload as an emoji. We only accept JPG, PNG, and GIF files.
com.atlassian.confluence.plugins.emoticons.max.allowed.uploads		
8.4.0	2000	This property controls the maximum number of emojis that are allowed to be uploaded to a Confluence site.
applinks.allow.all.hosts		
8.4.0	false	This property blocks local IPs and network link IPs (including AWS magic IP) during new app link creation. This helps prevent SSRF attacks on the local server. If you don't need to block local IPs and network link IPs, then set the property to`true` and restart the server. This property is also available in Long Term Support releases from 7.13.19 and 7.19.11.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Working with Confluence Logs

Confluence uses Apache's log4j logging service. This allows administrators to control the logging behavior and the log output file. There are six log4j logging levels.

If you request help from Atlassian Support, we will almost always ask for the Confluence application logs. The easiest way to get these logs is to go to > **General Configuration** > **Troubleshooting and support tools** and follow the prompts to create a **Support Zip**.

On this page:

- Application log files
- Change the log file configuration
- Change the destination of the log files
- Change the size and number of log files
- Change the logging levels
- Specific Confluence logging options
- Scanning log files for known problems
- Mark logs when troubleshooting issues
- Tomcat logs

Related pages:

- Enabling Detailed SQL Logging
- Enabling user access logging
- Generating a Thread Dump

Application log files

By default, the application log files can be found in the `<local-home>/logs` directory. This location is configurable, so you may need to check the log config for your location.

To make troubleshooting problems easier, the application log is split into several distinct log files:

- **atlassian-confluence.log**
  This is the main application log file, most entries will be written here. When you start Confluence, any log entries written to the console will also be repeated in this log.
- **atlassian-confluence-index.log**
  This file contains entries related to the search index.
- **atlassian-confluence-outgoing-mail.log**
  This file contains entries related to outgoing mail, such as notifications.
- **atlassian-confluence-security.log**
  This file contains entries related to your users and user directories.
- **atlassian-synchrony.log**
  This file contains entries related to Synchrony, which powers collaborative editing.
- **atlassian-diagnostics.log**
  This file contains entries for an experimental diagnostics feature which provides alerts for things like low disk space and memory.
- **atlassian-confluence-jmx.log**
  This file contains entries for Java Management Extensions API metrics which allow you to monitor the status of your instance in real time.
- **atlassian-confluence-ipd-monitoring.log**
  This file contains entries for in-product diagnostics including database connectivity and HTTP connection metrics which helps with identifying performance issues in your infrastructure.
- **atlassian-confluence-migrations.log**
  This file contains entries related to migration to Cloud.
- **atlassian-confluence-health-checks.log**
  This file contains details about issues and errors detected during the startup of your Confluence instance that need to be resolved.

In our documentation, when we refer to the "application log", we are referring to any of these files.

You can check the exact classes or packages that are logged to each file in the `log4j.properties` file under `LOGGING LOCATION AND APPENDER`.

Change the log file configuration

The logging behavior for Confluence and Synchrony is defined in the following properties file: `<CONFLUENCE-INSTALL>/confluence/WEB-INF/classes/log4j.properties`

This file is a standard log4j configuration file, as described in the Apache log4j documentation.

Change the destination of the log files

In log4j, an output destination is called an 'appender'. To change the destination of the log files, you need to stop Confluence and then change the settings in the '**Logging Location and Appender**' section of the `log4j.properties` file.

In the standard properties file, you will find entries for two appenders:

- `com.atlassian.confluence.logging.ConfluenceHomeLogAppender` This is a custom appender which controls the default logging destination. This appender allows the following settings:
  - MaxFileSize
  - MaxBackupIndex
- `org.apache.log4j.RollingFileAppender` If you want to log to a different location, uncomment the `RollingFileAppender` line and change the destination file in the line below it. Comment out the previous lines referring to the `ConfluenceHomeLogAppender`.

The Synchrony log destination can also be changed in the same way in file.

Confluence ships with the full suite of appenders offered by log4j. Read more about appenders in the log4j documentation.

For more detailed information seeConfiguring log4j in Confluence to send specific entries to a different log file in our Knowledge Base.

Note: If you change the location of your log files, they will no longer be included when you generate a support zip. This means you'll need to attach your logs to any support requests manually.

Change the size and number of log files

By default, Confluence keeps 5 application log files, which are overwritten as they reach 20 MB.

You can change the default log size and the number of log files to keep by editing the following values in `<CONFLUENCE-INSTALL>/confluence/WEB-INF/classes/log4j.properties`file.

```
log4j.appender.confluencelog.MaxFileSize=20480KB
log4j.appender.confluencelog.MaxBackupIndex=5
```

Change the logging levels

This can be done in the Confluence UI. See Configuring Logging for instructions on how to change the logging configuration of Confluence.

Specific Confluence logging options

Here's some specific log configurations you may need when troubleshooting.

Log the details of SQL requests made to the database

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

You may want to increase Confluence's logging so that it records individual SQL requests sent to the database. This is useful for troubleshooting specific problems. SeeEnabling Detailed SQL Logging.

Log the details of users accessing each Confluence page

Access logging using Tomcat Valve is enabled by default from Confluence 7.11. These logs are not part of the application log, and can be found at`<install directory>/logs/conf_access_log.<date>.log`.

You can however configure the application log to show which users are accessing which pages in Confluence. SeeHow to Enable User Access Loggingin our Knowledge Base.

Scanning log files for known problems

Atlassian Troubleshooting and support tools includes a log analyzer that can check for you Confluence logs for errors and match them against known problems in our knowledge base and issue tracker.

See Troubleshooting Problems and Requesting Technical Supportto find out how to set up a periodic scan of your log files.

Mark logs when troubleshooting issues

When troubleshooting it can be useful to mark your log files before and after you attempt to replicate the problem. This makes it easier for you to locate the specific part of the log to investigate.

To mark the application log files:

1. Go to > **General Configuration**>**Logging and Profiling**.
2. If you run Confluence in a cluster, select a cluster node.
3. Enter a message, for example "Reproduce directory sync issue".
4. Select**Rollover log files**if you want to start a new log file with your mark (this will delete the oldest log file).
5. Select**Mark**.

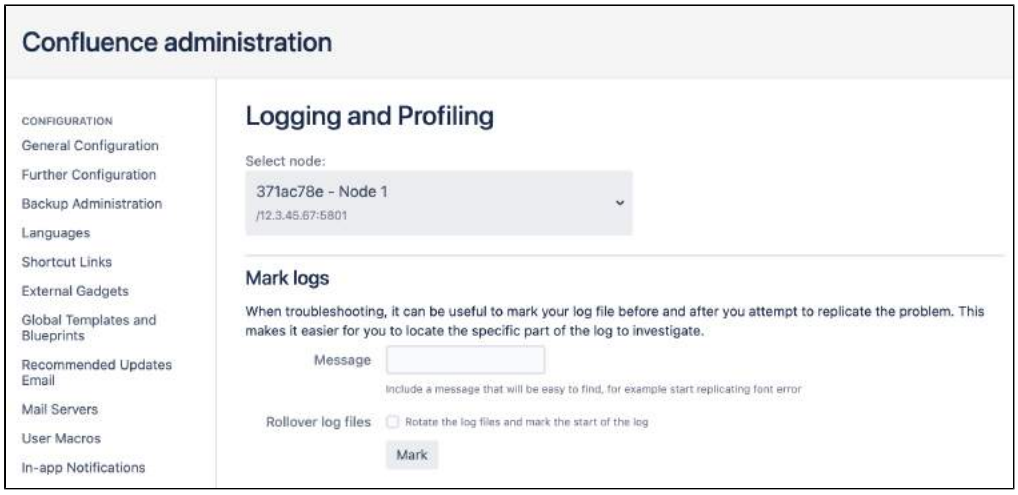

Screenshot: The logging and profiling screen in a cluster

Your message will be added to all the application log files (such as atlassian-confluence.log, and atlassian-confluence-security.log). You can mark your logs as often as you need to.

Here's an example:

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

```
...
2021-01-04 13:21:47,421 INFO [http-nio-8090-exec-5
[impl.admin.actions.MarkAllLogsAction] execute
********************************************
Reproduce directory sync issue
********************************************
2021-01-04 13:25:23,901 ERROR [test error] [atlassian.confluence.test] This
is a sample error java.lang.RuntimeException: Unable to find sample error
for testuser
...
```

Tomcat logs

There are some additional logs in your Confluence install directory that can be useful when troubleshooting issues with your Confluence site.

- `<install directory>/logs/catalina-<date>.log`
  This log records Tomcat operations, such as starting and stopping application server.
- `<install directory>/logs/conf_access_log.<date>.log`
  This is where you find Confluence's access logs. These logs are configured in the `server.xml`. SeeTomcat Access Log Valve documentation for further configuration options.
- `<install directory>/logs/gc.<date>.log`
  This is where you find the garbage collection logs. These logs provide useful information if you're experiencing long GC pauses.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Configuring Logging

There are many situations where you may want to change what is written to the Confluence application logs, particularly when troubleshooting a specific problem.

You can temporarily change the logging behaviourin the Logging and Profiling screen, while Confluence is running. Your changes will be discarded when you restart Confluence.

Alternatively, you can permanently change the logging behaviour in the log4j properties file.

On this page:

- Temporarily change logging behaviour while Confluence is running
- Permanently change logging level in the properties file
- Configuring Levels for java.util.logging in logging.properties

Terminology: In log4j, a 'logger' is a named entity. Logger names are case-sensitive and they follow a hierarchical naming standard. For example, the logger named `com.foo` is a parent of the logger named `com.foo.bar`.

Temporarily change logging behaviour while Confluence is running

When troubleshooting a problem, you can temporarily change the logging behaviour while Confluence is running. These changes aren't written to the`log4j.properties`file so are discarded when you next stop Confluence.

Change the logging level of an existing class or package

You need **System Administrator** global permissions to do this. Seelog4j Logging Levels for information on the specific levels.

To change the logging level of an existing class or package:

1. Go to> **General Configuration**> **Logging and Profiling**.
2. If you run Confluence in a cluster, select a cluster node.
3. Locate the relevant class or package, and select a value from the **New Level** dropdown.
4. **Save** your changes.

Alternatively, choose **Remove** if you want to stop logging a particular class or package.

Remember, your changes will not be written to the`log4j.properties`file so will be discarded when you next stop Confluence.

Add logging for an additional class or package

Sometimes our support team may ask you to enable some additional logging when troubleshooting a specific problem.You need**System Administrator**global permissions to do this.

To set the logging level for a new class or package:

1. Go to> **General Configuration**>**Logging and Profiling**.
2. If you run Confluence in a cluster, select a cluster node.
3. Enter the name in the **Class/Package name** field. It will be something like`com.atlassian.confluence.example`.
4. Select **Add entry** to add the new class or package to the list.
5. Locate the relevant class or package, and select a value from the**New Level**dropdown.
6. **Save**your changes.

Turn page profiling on or off

SeeTroubleshooting Slow Performance Using Page Request Profiling for more information on when and how to use page profiling.

Turn detailed SQL logging on or off

SeeEnabling Detailed SQL Loggingfor more information on when and how to use detailed SQL logging.

Preset logging configurations

Confluence provides two preset log configurations:

- **Production**
  This is the recommended default configuration, which aims to provide the most important information without flooding your logs.
- **Diagnostic**
  This changes the logging level of most packages to debug. It can be useful when troubleshooting,but results in slower performance and will fill up your log files more quickly.

If you want to reset your logging configuration to the default, select**Production**.

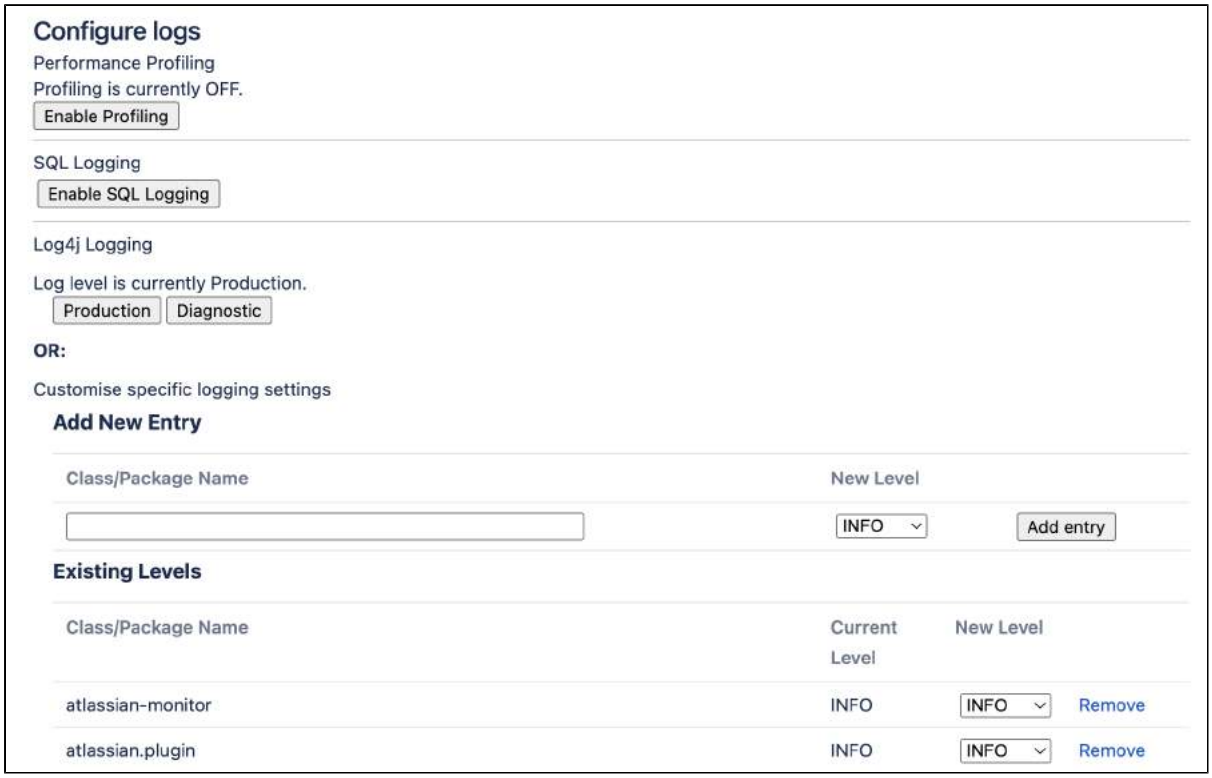

Screenshot: Changing Log Levels and Profiling

Permanently change logging level in the properties file

You need access to the installation directory to do this.

The`log4j.properties`file lists all the default classes and packages, plus some additional packages that you can choose to enable.

To permanently change the logging level of a class or package:

1. Stop Confluence.
2. Edit the`<CONFLUENCE-INSTALL>/confluence/WEB-INF/classes/log4j.properties`file.
3. Under **Logging Levels**, locate the package you want to change. Remove the # symbol to uncomment the package if necessary.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

4. Change the logging level to **DEBUG**, **INFO**, **WARNING**, **ERROR**, or **FATAL**. Seelog4j Logging Levels for accepted values. For example:

```
log4j.logger.com.atlassian.confluence.cache=DEBUG
```

5. Save the file, and restart Confluence for the changes to take effect.

If you're running Confluence in a cluster, you will need to repeat your change on every node. You can take each node down one by one, there's no need to stop the whole cluster.

If you want to log a class or package that's not already listed, simply add it to the file. Seethelog4j documentation for more information.

If you want to change the location, number, or size of the logs, seeWorking with Confluence Logs.

Configuring Levels for java.util.logging in logging.properties

A few libraries used by Confluence use java.util.logging rather than log4j or slf4j. These libraries include:

- com.sun.jersey
- org.apache.shindig

Confluence's `logging.properties` file is set to redirect java.util.logging at specific levels to log4j via slf4j.

To increase logging levels for these libraries you must first configure the `logging.properties` file in `<CONFLUENCE-INSTALL>/confluence/WEB-INF/classes/`. The logging levels are different from log4j. Seej ava.util.loggingin the Java 11 documentation.

For example, to increase logging for shindig change the following line in the `logging.properties` file:

```
org.apache.shindig.level = INFO
```

to

```
org.apache.shindig.level = FINE
```

And then use one of the methods above **as well** to configure the log4j level.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

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** (changes made here will not persist across restarts).
2. Manually modifying the `<Confluence-Install>\confluence\WEB-INF\classes\log4j.properties` file.

Default Log Level

The standard Confluence log level **WARN** is a way for Confluence to communicate with the server administrator. Logging at WARN level and higher should be reserved for situations that require some kind of attention from the server administrator, and for which corrective action is possible.

*See*log4j manualfor more information.

Troubleshooting SQL Exceptions

If you get an exception similar to those shown below, it is a good idea to increase the logging levels of your Confluence instance. If you request Atlassian support, this additional logging will help us work out the cause of the error.

Increased logging levels will enable us to diagnose errors like these:

```
org.springframework.dao.DataIntegrityViolationException: (HibernateTemplate): data integrity violated by
SQL ''; nested exception is java.sql.BatchUpdateException: Duplicate entry '1234' for key 1
 at org.springframework.jdbc.support.SQLStateSQLExceptionTranslator.translate
(SQLStateSQLExceptionTranslator.java:88)
caused by: java.sql.BatchUpdateException: Duplicate entry '1234' for key 1
 at com.mysql.jdbc.ServerPreparedStatement.executeBatch(ServerPreparedStatement.java:647)
```

or

```
(HibernateTemplate): data integrity violated by SQL ''; nested exception is java.sql.BatchUpdateException:
ORA-00001: unique constraint (CONFLUENCE.SYS_C0012345) violated
```

This document outlines the steps to take to increasing logging on your system.

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

Configure access logs

Access logs record every request made to your site which can be useful for auditing purposes and when troubleshooting a problem with an integration, app, or feature

View access logs

The log file is located in `<install-directory>/logs/conf_access_log<date>.log`.

Here's an example of the log output:

```
[18/Dec/2020:12:57:26 +1100] testuser http-nio-8090-exec-6 0:0:0:0:0:0:0:1 GET /index.action HTTP/1.1 200
1020ms 7881 http://localhost:8090/login.action?os_destination=%2Findex.action&permissionViolation=true
Mozilla/5.0 (Macintosh; Intel Mac OS X 10_14_6) AppleWebKit/537.36 (KHTML, like Gecko) Chrome/87.0.4280.88
Safari/537.36
[18/Dec/2020:12:57:26 +1100] testuser http-nio-8090-exec-1 0:0:0:0:0:0:0:1 GET /images/icons/profilepics
/default.svg HTTP/1.1 200 46ms 1206 http://localhost:8090/index.action Mozilla/5.0 (Macintosh; Intel Mac OS
X 10_14_6) AppleWebKit/537.36 (KHTML, like Gecko) Chrome/87.0.4280.88 Safari/537.36
[18/Dec/2020:13:02:20 +1100] testuser http-nio-8090-exec-5 0:0:0:0:0:0:0:1 GET /rest/quickreload/latest
/2129924?since=1608256850024&_=1608256758069 HTTP/1.1 200 12ms 152 http://localhost:8090/pages/viewpage.
action?pageId=2129924 Mozilla/5.0 (Macintosh; Intel Mac OS X 10_14_6) AppleWebKit/537.36 (KHTML, like
Gecko) Chrome/87.0.4280.88 Safari/537.36
```

Access log pattern

The Confluence access log uses the Apace Tomcat access log valve. The information recorded about each request is configurable.

The default pattern is:

```
%t %{X-AUSERNAME}o %I %h %r %s %Dms %b %{Referer}i %{User-Agent}i
```

This will log the date and time, username, remote logical username, remote host name (or IP), the first line of the request, the HTTP status code of the response, time taken to process the request (in milliseconds), bytes sent (excluding headers), the referer, and user agent.

See Access Log valve attributes in the Tomcat 9 documentation for more information on each of the attributes.

Change the log retention

The access log is configured to keep logs for a maximum of 30 days. You can choose to increase or decrease this limit, but be aware you'll need to allow enough disk space to accomodate the log files.

To change how long access logs are kept:

1. Stop Confluence.
2. Edit the `<install-directory>/conf/server.xml` file.
3. Locate the access log valve, and change the value of `maxDays`.

    ```
    <Valve className="org.apache.catalina.valves.AccessLogValve"
           directory="logs"
           maxDays="30"
           pattern="%t %{X-AUSERNAME}o %I %h %r %s %Dms %b %{Referer}i %{User-Agent}i"
           prefix="conf_access_log"
           requestAttributesEnabled="true"
           rotatable="true"
           suffix=".log"
           />
    ```

4. Save the file, and restart Confluence.

If you're running Confluence in a cluster, you will need to repeat this process on each node. You don't need to stop the whole cluster, you can update each node in turn.

Disable access logging

To disable access logging:

1. Stop Confluence.
2. Edit the `<install-directory>/conf/server.xml` file.
3. Remove the entire access log valve, shown here.

```
<Valve className="org.apache.catalina.valves.AccessLogValve"
       directory="logs"
       maxDays="30"
       pattern="%t %{X-AUSERNAME}o %I %h %r %s %Dms %b %{Referer}i %{User-Agent}i"
       prefix="conf_access_log"
       requestAttributesEnabled="true"
       rotatable="true"
       suffix=".log"
       />
```

4. Save the file, and restart Confluence.

If you're running Confluence in a cluster, you will need to repeat this process on each node. You don't need to stop the whole cluster, you can update each node in turn.

Other access log options

For a lightweight alternative access log solution, you can also choose to enable access logging in the application logs.SeeHow to Enable User Access Loggingto find out how to do this.

This option is best suited to smaller sites, as the additional log entries may cause the application log to fill up and rotate too quickly.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Scheduled Jobs

The administration console allows you to schedule various administrative jobs in Confluence, so that they are executed at regular time intervals. The types of jobs which can be scheduled cover:

- Confluence site backups
- Storage optimization jobs to clear Confluence's temporary files and caches
- Index optimization jobs to ensure Confluence's search index is up to date
- Mail queue optimization jobs to ensure Confluence's mail queue is maintained and notifications have been sent.

You'll needSystem Administrator permissions in order to edit and manually run jobs.

On this page:

- Accessing Confluence's scheduled jobs configuration
- Running a job manually
- Changing a job's schedule
- Viewing a job's execution history
- Cron expressions

Related pages:

- Trigger Module
- Scheduling a Backup

Accessing Confluence's scheduled jobs configuration

To access Confluence's Scheduled Jobs configuration page:

1. Go to > **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

Running a job manually

To run a job manually head to the Scheduled Jobs list and choose **Run** next to the job. It will runimmediately.

Not all jobs can be run manually.

Changing a job's schedule

To change a job's schedule:

1. Choose **Edit** next to the job you want to change.
2. Enter the new day or time to run the job as a cron expression - there's more info about cron expressions below. There are also some jobs that can only be scheduled to run at regular intervals in seconds, we call these simple jobs.
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 scheduled using cron expressions

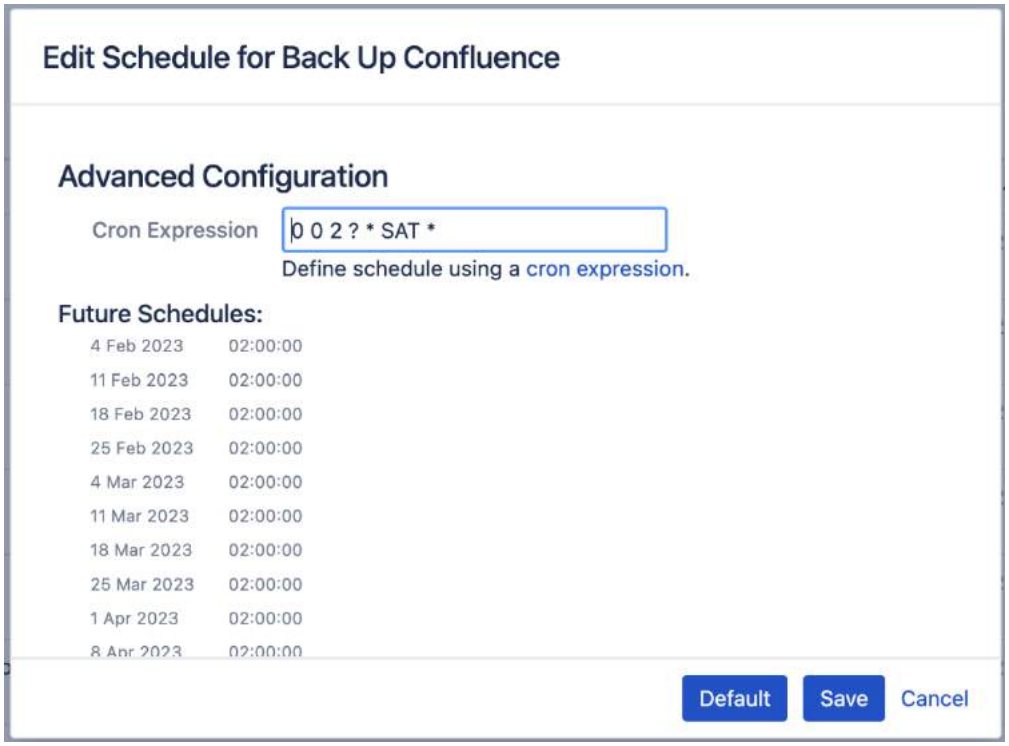

Screenshot: Configuring a simple job schedule using seconds

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

When you enable or edit a simple job, it won't run immediately after it has been enabled. Instead, it'll run at `interval seconds`, where `interval seconds` is the interval you've set.

For example, if you enable the simple job "Clean up re-index jobs" to run every 60 seconds, the first time it will run is 60 seconds after you save the schedule, and then every 60 seconds thereafter.

Disabling or re-enabling a job

By default, all jobs in Confluence are enabled.

Use the **Disable** / **Enable** links in the action column to disable and re-enable each job.

Not all jobs in Confluence can be disabled.

Viewing a job's execution history

To see when a job was last run, and how long the job took to run, click the **History** link beside the job.

If a job has not run at least once the History link won't appear.

Screenshot: Job Execution History

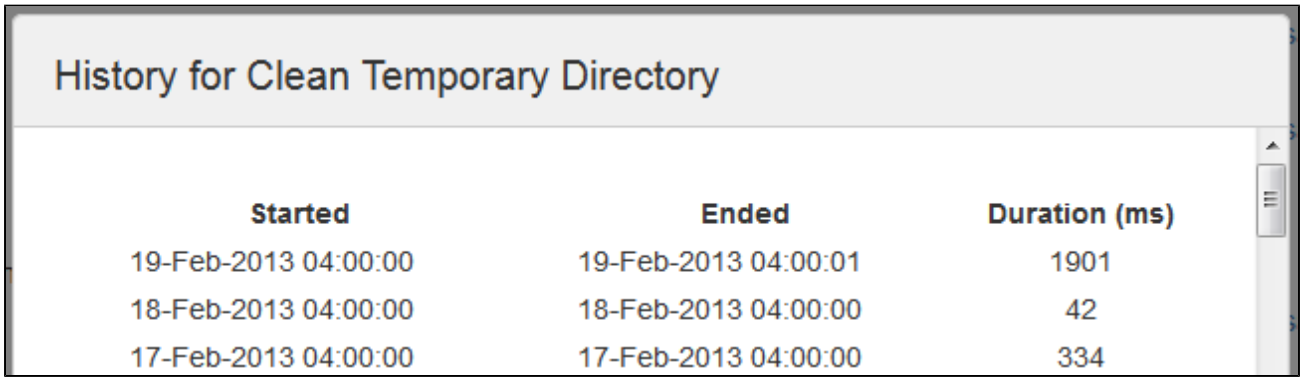

⚠Execution history is not available in Confluence Data Center.

Jobs overview

Here's a summary of some of the scheduled jobs that you may want to adjust.

Job Name	Description	Execution Behavior	Default Schedule
Back Up Confluence	Performs a backup of your entire Confluence site.	Per cluster	At 2am every day
Backup Restore v2 temporary backup zip cleaner	Cleans up any backup XML files saved to the restore directory for 72 hours or longer. If a user has chosen to save a backup permanently, it will not be deleted as part of this scheduled job.	Per cluster	Every hour

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Backup Restore v2 trigger job clean up	Cleans up any backup or restore jobs (including metadata and related files) done on the site 14 days ago or longer. If a user has chosen to save a backup permanently, it will not be deleted as part of this scheduled job.	Per cluster	At 1am 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 processedto ensure that its size does not grow indefinitely.	Per cluster	At 2am every day
Clean Temporary Directory	Cleans up temporary files generated in the `<confluence-home>/temp` directory. This temp directory is created by exports etc. 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 while keeping the job's frequency to 24 hours.	Per cluster	At 12am every day
Flush Change Index Queue	Flushes the Change Index Queue so Confluence's search results stay up to date.	Per node	Every minute
Flush Content Index Queue	Flushes the Content Index Queue so Confluence's search results stay up to date.	Per node	Every minute
Flush Edge Index Queue	Flushes the Edge Index Queue so Confluence's search results stay up to date.	Per node	Every 30 seconds
Flush Local Task Queue	Flushes the local task queue. (These are internal Confluence tasks that are typically flushed at a high frequency.)	Per node	Every minute
Flush Mail Queue	Sends notifications that have been queued up in the mail queue . This doesn't include batched notifications. Edit the **Send batched notifications** job if you also want to change how often notifications are sent for changes to a page or blog post.	Per node	Every minute

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Send batched notifications	Sends email notifications containing all changes to a page or blog post since the last time the job ran. Increase the time for fewer emails or reduce the time if more immediate notifications are important in your site.	Per cluster	Every 10 minutes
Flush Task Queue	Flushes the task queue. (These are internal Confluence tasks that are typically flushed at a high frequency.)	Per node	Every minute
Send 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
Purge Old Job Run Details	Confluence stores the details of each scheduled job that is run in the `scheduler_run_details` table in your database. In order to keep this table small for troubleshooting and debugging, the **Purge Old Job Run Details** job regularly removes the details of: • successful jobs run more than 90 days ago • unsuccessful jobs run more than 7 days ago You can override these settings using the following system properties; `jobs.limit.per.purge`, `all.jobs.ttl.hours` and `unsuccessful.jobs.ttl.hours`.	Per cluster	at 11pm every day
Property Entry Gardening	When a page is created from a blueprint, some data is left behind in the os_property table after the page is published. This job cleans up leftover data, which could contain personally identifiable information.	Per cluster	At 12am every day
Clean up unpublished Blueprint Page Entities	This job cleans up metadata stored about draft pages created from blueprints, which could contain personally identifiable information.	Per cluster	At 2:23am every day
Synchrony data eviction (soft)	Evicts Synchrony data for any content that has not been modified in the last 3 days, and does not have an active editor session.See How to remove Synchrony data for more information.	Per cluster	Every 10 minutes
Synchrony data eviction (hard)	Evicts Synchrony data for any content that is 15 days or older, regardlessof whether it has been modified more recently. See How to remove Synchrony data for more information.	Per cluster	Disabled by default
Versions Removal (Soft)	Deletes any historical page or attachment versions that don't meet the retention rules. Deletion happens in batches to minimise performance impact. After changing a retention rule, the job may need to run multiple times before all historical versions are removed. This job will only impact Confluence Data Center instances where retention rules are customizable.	Per cluster	Every 10 minutes
Versions Removal (Hard)	Deletes any historical page or attachment versions that don't meet the retention rules in one go, rather than in batches. This job can be run manually when required, but may have an impact on your site performance. This job will only impact Confluence Data Center instances where retention rules are customizable.	Per cluster	Disabled by default

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Trash Removal (Soft)	Purges any items from Space's trash that don't meet theretention rules. Deletion happens in batches to minimize performance impact. After changing a retention rule, the job may need to run multiple times before all trash is removed. This job will only impact Confluence Data Center instances where retention rules are customizable.	Per cluster	Every 10 minutes
Trash Removal (Hard)	Purges any items from Space's trash that don't meet theretention rulesin one go, rather than in batches. This job can be run manually when required, but may have an impact on your site performance. This job will only impact Confluence Data Center instances where retention rules are customizable.	Per cluster	Disabled by default

Type of jobs

There are some jobs you can schedule to repeat at intervals defined by a simple values like seconds. We call these simple jobs. When you edit/enable these simple jobs, they won't be executed instantly. Instead, they'll run at `interval seconds`, where interval seconds is the interval

For example, if the Check Cluster Safety job is set to repeat every 30 seconds, when it's enabled, it will be scheduled to run in interval seconds plus 30 seconds, which is 30 + 30 = 60 seconds.

Now you want to edit the Check Cluster Safety job to repeat every 40 seconds, when it's enabled, it will be scheduled to run in interval seconds plus 40 seconds, which is 40 + 30 =

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

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

7	Year	1970-2099	No

* Excluding special characters.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Configuring the Allowlist

Confluence administrators can choose to allow incoming and outgoing connections and content from specified sources for use in the:

- RSS Feed Macro
- HTML Include macro(disabled by default)
- gadgets
- Shared Links Blueprint
- Widget Connector Macro

by adding URLs to the allowlist.

Confluence will display an error if content has been added that is not from an allowed source, and prompt the user to add the URL to the allowlist.

Application linksare automatically added to the allowlist. You don't need to manually add them.

Add allowed URLs to theallowlist

To add a URL to the allowlist:

1. Go to > **General Configuration**> **Allowlist**.
2. Enter the URL or expression you want to allow.
3. Choose the **Type** of expression (see below for examples of the types available).
4. Choose **Allow Incoming** if you need to allow CORS requests (see below).
5. Choose **Allow anonymous users** if you need to allow unauthenticated users.
6. Choose **Add**.

Your URL or expression appears in the allowlist.

To test that your allowlistedURL is working as expected you can enter a URL in the**Test a URL**field. Icons will indicate whether incoming and / or outgoing traffic is allowed for that URL.

Expression types

When adding a URL to the allowlist, you can choose from a number of expression types.

When deciding the best expression type to use, aimfor a more restrictive URL, rather than less restrictive, to best protect your site.

Type	Description	Example
Domain name	Allows all URLs from the specified domain.	`https://www.example.com`
Exact match	Allows only the specified URL.	`https://www.example.com/thispage`
Wildcard Expression	Allows all matching URLs. Use the wildcard * character to replace one or more characters.	`https://*example.com`
Regular Expression	Allows all URLs matching the regular expression.	`http(s)?://www\.example\.com`

Allow Incoming

Allow Incoming enables CORS requests from the specified origin. The URL must match the 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`.

Allow anonymous users

You can use the **Allow anonymous users** option to allow outbound requests on behalf of unauthenticated users.

This isn't recommended for URLs that may contain private data, such as URLs from application links. If you do need to provide anonymous access, consider using an exact URL or wildcard based rule to limit access to just the required resources.

Change default settings for application links

When you create an application link, the URL is automatically added to the Confluence allowlist. By default, outbound requests from these URLs is only allowed for authenticated users.

To change the default behaviour for all application links, including new application links:

1. Go to > **General Configuration** > **Allowlist**.
2. Select **Configure Settings**.
3. Select either:
   - **Allow all users** to allow outbound requests for all users, including anonymous users
   - **Allow authenticated users** to deny outbound requests for anonymous users
   - **Restrict by default** to deny outbound requests for all users (the applink will not be added to the allowlist at all)
4. **Save** your changes.

All existing application links, and any new application links added to the allowlist, will use this setting.

Disable the allowlist

The allowlist is enabled by default. You can choose to disable the allowlist however this will allow all URLs, including malicious content.

We **strongly discourage you from disabling the allowlist**, as it will leave you vulnerable to Server-Side Request Forgery (SSRF) attacks, such as the one disclosed in

~~CONFSERVER-61399~~ PUBLISHED .

To disable the allowlist:

1. Go to > **General Configuration** > **Allowlist**.
2. Choose **Turn off allowlist**.
3. Choose **Confirm**.

All URLs will now be allowed. This is not recommended.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Configuring the Time Interval at which Drafts are Saved

Confluence saves a draft of your page once every thirty seconds by default. Confluence administratorscan configure how often drafts are saved.

To set the time interval at which drafts are saved:

1. Choose the**cog icon** , then choose**General Configuration**
2. Click**Further Configuration**in the left-hand panel.
3. Edit the setting for**Draft Save Interval**.

This setting applies regardless of whether collaborative editing is on or off. When collaborative editing is on, every keystroke is alsosaved by Synchrony in real time.

Related pages:

- Drafts

Configuring Confluence Security

This section gives guidelines on configuring the security of your Confluence site:

- Confluence Security Overview and Advisories
- Proxy and HTTPS setup for Confluence
- Configuring Secure Administrator Sessions
- Confluence Cookies
- Using Fail2Ban to limit login attempts
- Securing Confluence with Apache
- Best Practices for Configuring Confluence Security
- Hiding the People Directory
- Configuring Captcha for Spam Prevention
- Hiding External Links From Search Engines
- Configuring Captcha for Failed Logins
- Configuring XSRF Protection
- User Email Visibility
- Anonymous Access to Remote API
- Configuring RSS Feeds
- Preventing and Cleaning Up Spam

Related pages:

- Permissions and restrictions
- Configuring a Confluence Environment
- Confluence administrator's guide

Confluence Security Overview and Advisories

This document is for system administrators who want to evaluate the security of the Confluence web application. The page addresses overall application security and lists the security advisories issued for Confluence. As a public-facing web application, Confluence's application-level security is important. This document answers a number of questions that commonly arise when customers ask us about the security of our product.

Other topics that you may be looking for:

- For information about user management, groups and permissions, please refer to the internal security overview.
- For guidelines on configuring the security of your Confluence site, see the administrator's guide to configuring Confluence security.

Application Security Overview

Password Storage

When Confluence's internal user management is used, since version 3.5 of Confluence passwords are hashed through the saltedPKCS5S2 implementation provided byEmbedded Crowdbefore being stored in the database. There is no mechanism within Confluence to retrieve a user's password when password recovery is performed, a reset password linkis generated and mailed to the user's registered address.

When external user management is enabled, password storage is delegated to the external system.

On this page:

- Application Security Overview
- Finding and Reporting a Security Vulnerability
- Publication of Confluence Security Advisories
- Severity Levels
- Our Security Bugfix Policy
- Published Security Advisories

Buffer Overflows

Confluence is a 100% pure Java application with no native components. As such it is highly resistant to buffer overflow vulnerabilities possible buffer overruns are limited to those that are bugs in the Java Runtime Environment itself.

SQL Injection

Confluence interacts with the database through the Hibernate Object-Relational mapper. Database queries are generated using standard APIs for parameter replacement rather than string concatenation. As such, Confluence is highly resistant to SQL injection attacks.

Script Injection

Confluence is a self-contained Java application and does not launch external processes. As such, it is highly resistant to script injection attacks.

Cross-Site Scripting

As a content-management system that allows user-generated content to be posted on the web, precautions have been taken within the application to prevent cross-site scripting attacks:

- The wiki markup language in Confluence does not support dangerous HTML markup
- Macros allowing the insertion of raw HTML are disabled by default
- HTML uploaded as a file attachment is served with a content-type requesting the file be downloaded, rather than being displayed inline
- Only system administrators can make HTML-level customizations of the application

When cross-site scripting vulnerabilities are found in the Confluence web application, we endeavor to fix them as quickly as possible.

Transport Layer Security

Confluence does not directly support SSL/TLS. Administrators who are concerned about transport-layer security should set up SSL/TLS at the level of the Java web application server, or the HTTP proxy in front of the Confluence application.

For more information on configuring Confluence for SSL, see: Running Confluence Over SSL or HTTPS

Session Management

Confluence delegates session management to the Java application server in which it is deployed. We are not aware of any viable session-hijacking attacks against the Tomcat application server shipped with Confluence. If you are deploying Confluence in some other application server, you should ensure that it is not vulnerable to session hijacking.

Plugin Security

Administrators install third party plugins **at their own risk**. Plugins run in the same virtual machine as the Confluence server, and have access to the Java runtime environment, and the Confluence server API.

Administrators should always be aware of the source of the plugins they are installing, and whether they trust those plugins.

Administrator Trust Model

Confluence is written under the assumption that anyone given System Administrator privileges is trusted. System administrators are able, either directly or by installing plugins, to perform any operation that the Confluence application is capable of.

As with any application, you should not run Confluence as the root/Administrator user. If you want Confluence to listen on a privileged network port, you should set up port forwarding or proxying rather than run Confluence with additional privileges. The extra-careful may consider running Confluence inside a `chroot` jail.

Stack Traces

To help when debugging a problem, Confluence provides stack traces through the web interface when an error occurs. These stack traces include information about what Confluence was doing at the time, and some information about your deployment server.

This includes information such as operating system and version and Java version. With proper network security, this is not enough information to be considered dangerous. The username of the current user may be included.

Thread dumps include usernames and URLs by default. If you don't want to include this additional diagnostic information, you can disable Thread diagnostics.

Finding and Reporting a Security Vulnerability

Atlassian's approach to reporting security vulnerabilities is detailed in How to Report a Security Issue.

Publication of Confluence Security Advisories

Atlassian's approach to releasing security advisories is detailed in Security Advisory Publishing Policy.

Severity Levels

Atlassian's approach to ranking security issues is detailed in Severity Levels for Security Issues.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Our Security Bugfix Policy

Our approach to releasing patches for security issues is detailed in ourSecurity Bugfix Policy.

Published Security Advisories

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Proxy and HTTPS setup for Confluence

Many customers choose to run Confluence behind a reverse proxy, often with HTTPS enabled. Getting your proxy configuration right is essential, to avoid problems later when using Confluence.

Proxy and HTTPS access are both configured in Tomcat, Confluence's application server.

Sample connectors

To make setup and configuration as straightforward as possible, we've provided a number of sample connectors in the Tomcat`<install-directory>/conf/server.xml`file.

Sample connector	Description
DEFAULT - Direct connector with no proxy, for unproxied HTTP access to Confluence	This is the default option. Use this option when you don't have a reverse proxy and are *not* enabling HTTPS.
HTTP - Proxying Confluence via Apache or Nginx over HTTP	Choose this option if you have a reverse proxy, but are *not* enabling HTTPS.
HTTPS - Direct connector with no proxy, for unproxied HTTPS access to Confluence.	Choose this option if you want to use HTTPS without a reverse proxy.HTTPS will be terminated at Tomcat.
HTTPS - Proxying Confluence via Apache or Nginx over HTTPS	Use this option when you want to use a reverse proxy and enable HTTPS.This is the most common configuration.

We only provide HTTP/HTTPS connector examples. You can't use the AJP connector (for example, with Apache mod_jk), as Synchrony, which is required for collaborative editing, can't accept AJP connections.

If you plan to use collaborative editing, there are a number of proxy and SSL considerations you'll need to take into account when deciding the best way to configure your proxy.

Step-by-step guides

In addition to the sample connectors, we also provide a number of step-by-step guides to help you enable HTTPS and configure your proxy correctly.

HTTPS:

- Running Confluence Over SSL or HTTPS(terminating HTTPS at Tomcat)
- Running Confluence behind NGINX with SSL(terminating HTTPS at your proxy)
- Securing your Atlassian applications with Apache using SSL(terminating HTTPS at your proxy)

Reverse proxy:

- Using Apache with mod_proxy(Confluence)
- Running Confluence behind NGINX with SSL(Confluence)
- Proxying Atlassian server applications with Apache HTTP Server (mod_proxy_http)(any Atlassian product)
- Proxying Atlassian server applications with Microsoft Internet Information Services (IIS)(any Atlassian product)

Outbound proxy:

- Configuring Web Proxy Support for Confluence (Confluence)
- How to Configure Outbound HTTP and HTTPS Proxy for your Atlassian application(any Atlassian product)

Although we provide guides for some third-party solutions, and mention Apache and Nginx in the `server.xml` file, you can choose your own proxy solution.

Atlassian Supportcan't provide assistance with configuring third-party tools like NGINX, Apache, or IIS. If you have questions, check your proxy server's documentation, ask the Atlassian Community,or get help from a Solution Partner.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Configuring Web Proxy Support for Confluence

The content on this page relates to platforms which are not supported. Consequently, Atlassian Support **cannot guarantee providing any support for it**. Please be aware that this material is provided for your information only and using it is done so at your own risk.

Some of Confluence's macros, such as {rss} and {jiraissues} need to make web requests to remote servers in order to retrieve data. If Confluence is deployed within a data centre or DMZ, it may not be able to access the Internet directly to make these requests. If you find that the {rss} macro does not work, ask your network administrator if Confluence needs to access the Internet through a web proxy.

Configuring an outbound HTTP proxy in Confluence

Proxy support is configured by passing certain system properties to the Java Virtual Machine on startup.

- http.proxyHost
- http.proxyPort (default: 80)
- http.nonProxyHosts (default: <none>)
- https.proxyHost
- https.proxyPort

At a minimum, you need to define `http.proxyHost` to configure an HTTP proxy, and `https.proxyHost` to configure an HTTPS proxy. System property configuration is described in the Configuring System Properties.

Properties `http.proxyHost` and `http.proxyPort` indicate the proxy server and port that the http protocol handler will use, and `https.proxyHost` and `https.proxyPort` indicate the same for the https protocol handler.

```
-Dhttp.proxyHost=proxy.example.org -Dhttp.proxyPort=8080 -Dhttps.proxyHost=proxy.example.org -Dhttps.
proxyPort=8080
```

Property `http.nonProxyHosts` indicates the hosts which should be connected to directly and not through the proxy server. The value can be a list of hosts, each separated by a pipe character | . In addition, a wildcard character (asterisk) * can be used for matching. For example:

```
-Dhttp.nonProxyHosts=*.foo.com|localhost
```

If you're using Confluence 6.0 or later with Synchrony, you'll need to pass the following to ensure Confluence can connect directly to Synchrony. Replace `localhost|127.0.0.1` with your Synchrony IP if you have used the `synchrony.host` system property to change the IP Synchrony uses.

```
-Dhttp.nonProxyHosts=localhost|127.0.0.1
-Dhttps.nonProxyHosts=localhost|127.0.0.1
```

Note: You may need to escape the pipe character | in some command-line environments.

If 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 authenticatoravailable.

To configure NTLM authentication for your HTTP proxy, you need to define a domain system property, `http.auth.ntlm.domain`, in addition to the properties for host, port and username mentioned above:

```
-Dhttp.auth.ntlm.domain=MYDOMAIN
```

Configuring authentication order

Sometimes multiple authentication mechanisms are provided by an HTTP proxy. If you have proxy authentication failure messages, you should first check your username and password, then you can check for this problem by examining the HTTP headers in the proxy failure with a packet sniffer on the Confluence server. (Describing this is outside the scope of this document.)

To set the order for multiple authentication methods, you can set the system property `http.proxyAuth` to a comma-separated list of authentication methods. The available methods are: ntlm, digest and basic; this is also the default order for these methods.

For example, to attempt Basic authentication before NTLM authentication, and avoid Digest authentication entirely, you can set the `http.proxyAuth` property to this value:

```
-Dhttp.proxyAuth=basic,ntlm -Dhttps.proxyAuth=basic,ntlm
```

Troubleshooting

1. There's a diagnostic jsp file in CONF-9719 for assessing the connection parameters.
2. '**Status Code [407]**' errors are described in APR-160.
3. Autoproxies are not supported. See ~~CONFSERVER-16941~~ CLOSED .

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

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, seeConfiguring an SSL Connection to Active Directory.
- You want to set up your Jira application as atrusted applicationin 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, seeRunning Confluence Over SSL or HTTPS.

There's a Confluence SSL plugin that facilitates this process.

Importing SSL Certificates

For more information on these commands, see the Keytool documentation.

1. Add the root certificate to your default Java keystore with the following command. This is the certificate that was used to authorize the LDAP server's certificate. It will be either the one that was used for signing it, or will come from further up in the trust chain, possibly the root certificate. This is often a self-signed certificate, when both ends of the SSL connection are within the same network. Again, the exact alias is not important.

   ```
   keytool -importcert -alias serverCert -file RootCert.crt -keystore %JAVA_HOME%/jre/lib/security
   /cacerts (Windows)
   keytool -importcert -alias serverCert -file RootCert.crt -keystore $JAVA_HOME/jre/lib/security
   /cacerts (Linux/Unix/Mac)
   ```

2. Import your LDAP or Jira server's public certificate into the JVM Keystore. This is the certificate that the LDAP server will use to set up the SSL encryption. You can use any alias of your choosing in place of "JIRAorLDAPServer.crt".

   ```
   keytool -importcert -alias ldapCert -file JIRAorLDAPServer.crt -keystore %JAVA_HOME%/jre/lib
   /security/cacerts (Windows)
   keytool -importcert -alias ldapCert -file JIRAorLDAPServer.crt -keystore $JAVA_HOME/jre/lib
   /security/cacerts (Linux/Unix/Mac)
   ```

3. Verify that the certificate has been added successfully by entering the following command:

   ```
   keytool -list -keystore %JAVA_HOME%/jre/lib/security/cacerts (Windows)
   keytool -list -keystore $JAVA_HOME/jre/lib/security/cacerts (Unix/Linux)
   keytool -list -keystore /Library/Java/Home/lib/security/cacerts (Mac)
   ```

4. Ensure that you have updated CATALINA_OPTS to specify the path to the keystore, as specified in Connecting to SSL servicesbefore restarting Confluence.
   There is no need to specify an alias for Confluence to use. On connecting to the LDAP server, it will search through the keystore to find a certificate to match the key being presented by the server.

Troubleshooting

Check the following knowledge base articles:

- Unable to Connect to SSL Services due to PKIX Path Building Failed
- SSL troubleshooting articles

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Using Apache with mod_proxy

Atlassian applications allow the use of reverse-proxies, however Atlassian Support does not provide assistance for configuring them. Consequently, Atlassian**can not guarantee providing any support for them**.

If assistance with configuration is required, please raise a question onAtlassian Answers.

This page describes one possible way to use Apache HTTP Server 2.4 to proxy requests for Confluence running in a standard Tomcat container. You can find additional documentation that explains how to useNGINXfor the same purpose.

You might use this configuration when:

- You have an existing Apache website, and want to add Confluence (for example,**http://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**and want to make them both available on the regular HTTP port (80) (for example, at**http://www.example.com/confluence**and**http://www.example.com/jira)**. Each application can be restarted, managed and debugged separately.

Note: This page documents a configuration of Apache, rather than of Confluence itself. Atlassian will support Confluence with this configuration, but we cannot guarantee to help you debug problems with Apache. Please be aware that this material is provided for your information only, and that you use it at your own risk.

On this page:

- Base configuration
  - 1 Set the context path
  - 2 Set the URL for redirection
  - 3 Configure mod_proxy
  - 4 Restart Apache
  - 5 Disable HTTP Compression
  - 6 Change the Confluence Base URL
- Adding SSL
- More information

Base configuration

In these examples, we use the following:

http://www.example.com/confluence - your intended URL

http://example:8090 - the hostname and port Confluence is currently installed to

http://example:8091 - the hostname and port Synchrony, the service that powers collaborative editing, defaults to

/confluence - the intended context path for Confluence (the part after hostname and port)

/synchrony - the context path for Synchrony, the process that powers collaborative editing

You'll need to replace these URLs with your own URLs.

1 Set the context path

If you want to access Confluence without a context path, such as**www.example.com**,skip this step.

Set your Confluence application path (the part after hostname and port) in Tomcat. In this example the context path will be`/confluence`.

Edit `<installation-directory>conf/server.xml`, locate the "Context" definition:

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

and change it to:

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

In this example we've used `/confluence` as the context path. Note that you can't use `/resources` as your context path, as this is used by Confluence, and will cause problems later on.

Restart Confluence, and check you can access it at **http://example:8090/confluence**.

2 Set the URL for redirection

Next, set the URL for redirection. In the same `<installation-directory>conf/server.xml` file, use the example connectors as a starting point.

Comment out the default connector (for unproxied access).

In XML a comment starts with `<!--` and ends with `-->`, and is used to make sure only the relevant portions of the file are read by the application.

Add `<!--` and `-->` around the **default** connector. It should now look like this.

```
<!--
===========================================================
DEFAULT - Direct connector with no proxy, for unproxied HTTP access to Confluence.
===========================================================
-->
<!--
<Connector port="8090" connectionTimeout="20000" redirectPort="8443"
   maxThreads="48" minSpareThreads="10"
   enableLookups="false" acceptCount="10" debug="0" URIEncoding="UTF-8"
   protocol="org.apache.coyote.http11.Http11NioProtocol"/>
-->
```

Uncomment the connector listed under the **HTTP - Proxying Confluence via Apache or Nginx over HTTP** heading.

To uncomment a section, remove the `<!--` and `-->` surrounding the connector.

Here's an example showing the default connector commented out, and the HTTP connector uncommented. The headings remain commented out.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

```
<!--
==========================================================
DEFAULT - Direct connector with no proxy, for unproxied HTTP access to Confluence.
==========================================================
-->
<!--
<Connector port="8090" connectionTimeout="20000" redirectPort="8443"
   maxThreads="48" minSpareThreads="10"
   enableLookups="false" acceptCount="10" debug="0" URIEncoding="UTF-8"
   protocol="org.apache.coyote.http11.Http11NioProtocol"/>
-->
<!--
==========================================================
HTTP - Proxying Confluence via Apache or Nginx over HTTP
==========================================================
-->
<Connector port="8090" connectionTimeout="20000" redirectPort="8443"
   maxThreads="48" minSpareThreads="10"
   enableLookups="false" acceptCount="10" debug="0"URIEncoding="UTF-8"
   protocol="org.apache.coyote.http11.Http11NioProtocol"
   scheme="http" proxyName="<subdomain>.<domain>.com" proxyPort="80"/>
```

Insert your **proxyName**and**proxyPort**as shown in the last line below:

```
<Connector port="8090" connectionTimeout="20000" redirectPort="8443"
   maxThreads="48" minSpareThreads="10"
   enableLookups="false" acceptCount="10" debug="0" URIEncoding="UTF-8"
   protocol="org.apache.coyote.http11.Http11NioProtocol"
   scheme="http" proxyName="www.example.com" proxyPort="80"/>
```

If you plan to enable HTTPS, use the connector under**HTTPS - Proxying Confluence via Apache or Nginx over HTTPS**.

3 Configure mod_proxy

Use one of the examples below to edit your Apache `http.conf`file to proxy requests to the application server.

⚠You will need to enable the following required Apache modules if they are not already enabled:

- mod_proxy
- mod_proxy_http
- proxy_wstunnel
- mod_rewrite

(proxy_wstunnel and mod_rewrite are new requirements in Confluence 6.0)

The format of the `http.conf`file, and location of the modules may differ on your operating system. We recommend Windows usersspecify the absolute path to the module files.

Example 1: Configuration with context path

Use this example if you set a context path in step 1, and will access Confluence with a context path like this**http://www.example.com/confluence**.

In this example, users will connect to Synchrony, which is required for collaborative editing, directly via WebSockets.

⚠The order of directives in the config is important.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Apache HTTP server 2.4

```
# Put this after the other LoadModule directives
LoadModule proxy_module /usr/lib/apache2/modules/mod_proxy.so
LoadModule proxy_http_module /usr/lib/apache2/modules/mod_proxy_http.so
LoadModule proxy_wstunnel_module /usr/lib/apache2/modules/mod_proxy_wstunnel.so
LoadModule rewrite_module /usr/lib/apache2/modules/mod_rewrite.so

# Put this in the main section of your configuration (or virtual host, if using Apache virtual hosts)
ProxyRequests Off
ProxyPreserveHost On

<Proxy *>
    Require all granted
</Proxy>

ProxyPass /synchrony http://<domain>:8091/synchrony
<Location /synchrony>
    Require all granted
    RewriteEngine on
    RewriteCond %{HTTP:UPGRADE} ^WebSocket$ [NC]
    RewriteCond %{HTTP:CONNECTION} Upgrade$ [NC]
    RewriteRule .* ws://<domain>:8091%{REQUEST_URI} [P]
</Location>

ProxyPass /confluence http://<domain>:8090/confluence
ProxyPassReverse /confluence http://<domain>:8090/confluence

<Location /confluence>
    Require all granted
</Location>
```

Note: It's not possible to use Apache HTTP Server 2.2 with Confluence 6.0 or later. If you plan to use SSL, you will need version 2.4.10 or later.

Example 2: Configuration without context path

Use this example if you skipped step 1, and will access Confluence without a context path like this**http://www.example.com**.

As in the previous example, users will connect to Synchrony, which is required for collaborative editing, directly via WebSockets.

⚠️The order of directives in the config is important.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Apache HTTP server 2.4

```
# Put this after the other LoadModule directives
LoadModule proxy_module /usr/lib/apache2/modules/mod_proxy.so
LoadModule proxy_http_module /usr/lib/apache2/modules/mod_proxy_http.so
LoadModule proxy_wstunnel_module /usr/lib/apache2/modules/mod_proxy_wstunnel.so
LoadModule rewrite_module /usr/lib/apache2/modules/mod_rewrite.so

# Put this in the main section of your configuration (or virtual host, if using Apache virtual hosts)

  ProxyRequests Off
  ProxyPreserveHost On

  RewriteEngine On
  RewriteCond %{REQUEST_URI} !^/synchrony
  RewriteRule ^/(.*) http://<domain>:8090/$1 [P]

  <Proxy *>
      Require all granted
  </Proxy>

  ProxyPass /synchrony http://<domain>:8091/synchrony

  <Location /synchrony>
      Require all granted
      RewriteEngine on
      RewriteCond %{HTTP:UPGRADE} ^WebSocket$ [NC]
      RewriteCond %{HTTP:CONNECTION} Upgrade$ [NC]
      RewriteRule .* ws://<domain>:8091%{REQUEST_URI} [P]
  </Location>

  ProxyPass / http://<domain>:8090
  ProxyPassReverse / http://<domain>:8090

  <Location />
      Require all granted
  </Location>
```

Note: It's not possible to use Apache HTTP Server 2.2 with Confluence 6.0 or later. If you plan to use SSL, you will need version 2.4.10 or later.

4 Restart Apache

This is needed to pick up on the new configuration.To restart Apache, run the following command:

```
sudo apachectl graceful
```

5 Disable HTTP Compression

Having compression run on both the proxy and Tomcat can cause problems integrating with other Atlassian applications, such as Jira. Please disable HTTP compression as per ourCompressing an HTTP Response within Confluencedocs.

6 Change the Confluence Base URL

The last stage is to set theBase URLto the address you're using within the proxy, for example**http://www.example.com/confluence**.

Adding SSL

If you plan to enable HTTPS, seeSecuring your Atlassian applications with Apache using SSL, andmake sure you choose the HTTPS sample connector.

More information

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

- 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.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Running Confluence behind NGINX with SSL

This page describes how to set up NGINXas areverse proxyfor Confluence.

The configuration described on this page results in a scenario where:

- External clientconnections with NGINX are secured using SSL. Connections between NGINXand Confluence Server are unsecured.
- Confluence Server and NGINXrun on the same machine.

On this page

- Step 1: Set the context path
- Step 2: Configure the Tomcat connector
- Step 3: Configure NGINX
- Step 4: Restart Confluence and NGINX

We assume that you already have a running instance of NGINX. If not, refer to theNGINX documentationfor instructions on downloading and installing NGINX.SSL certificates must be installed on the server machine. You'll an NGINX version that supports WebSockets (1.3 or later).

If your team plans to use the Confluence Server mobile app, you'll need a certificate issued by a trusted Certificate Authority. You can't use the app with a self-signed certificate, or one from an untrusted or private CA.

⚠ Atlassian Supportcan't provide assistance with configuring third-party tools like NGINX. If you have questions, check the NGINX documentation, ask the Atlassian Community,or get help from a Solution Partner.

Step 1: Set the context path

⚠If you want to access Confluence without a context path (www.example.com),or via a sub-domain (confluence.example.com)skip this step.

Set your Confluence application path (the part after hostname and port) in Tomcat. Edit `<installation-directory>/conf/server.xml`, locate the "Context" definition:

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

and change it to:

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

In this example we've used `/confluence`as the context path. Note that you can't use `/resources`as your context path, as this is used by Confluence, and will cause problems later on.

Restart Confluence, and check you can access it at**http://example:8090/confluence**

Step 2: Configure the Tomcat connector

In the same`<installation-directory>conf/server.xml`file, use the example connectorsas a starting point.

Comment out the default connector (for unproxied access).

> In XML a comment starts with`<!--` and ends with`-->`, and is used to make sure only the relevant portions of the file are read by the application.
>
> Add `<!--` and`-->` around the **default** connector. It should now look like this.

```
<!--
=========================================================
DEFAULT - Direct connector with no proxy, for unproxied HTTP access to Confluence.
=========================================================
-->
<!--
<Connector port="8090" connectionTimeout="20000" redirectPort="8443"
   maxThreads="48" minSpareThreads="10"
   enableLookups="false" acceptCount="10" debug="0" URIEncoding="UTF-8"
   protocol="org.apache.coyote.http11.Http11NioProtocol"/>
-->
```

Uncomment the connector listed under the**HTTPS - Proxying Confluence via Apache or Nginx over HTTPS**heading.

To uncomment a section, remove the`<!--` and`-->` surrounding the connector.

Here's an example showing the default connector commented out, and the HTTPS connector uncommented. The headings remain commented out.

```
<!--
=========================================================
DEFAULT - Direct connector with no proxy, for unproxied HTTP access to Confluence.
=========================================================
-->
<!--
<Connector port="8090" connectionTimeout="20000" redirectPort="8443"
   maxThreads="48" minSpareThreads="10"
   enableLookups="false" acceptCount="10" debug="0" URIEncoding="UTF-8"
   protocol="org.apache.coyote.http11.Http11NioProtocol"/>
-->
...
<!--
=========================================================
HTTPS - Proxying Confluence via Apache or Nginx over HTTPS
=========================================================
-->
<Connector port="8090" connectionTimeout="20000" redirectPort="8443"
   maxThreads="48" minSpareThreads="10"
   enableLookups="false" acceptCount="10" debug="0" URIEncoding="UTF-8"
   protocol="org.apache.coyote.http11.Http11NioProtocol"
   scheme="https" secure="true" proxyName="<subdomain>.<domain>.com" proxyPort="443"/>
```

Insert your**proxyName**and**proxyPort**as shown in the last line below:

```
<Connector port="8090" connectionTimeout="20000" redirectPort="8443"
   maxThreads="48" minSpareThreads="10"
   enableLookups="false" acceptCount="10" debug="0" URIEncoding="UTF-8"
   protocol="org.apache.coyote.http11.Http11NioProtocol"
   scheme="https" secure="true" proxyName="www.example.com" proxyPort="443"/>
```

Make sure you've included correct values for`protocol` and`proxyName`.

Step 3: Configure NGINX

You will need to specify a listening server in NGINX, as in the example below. Add the following to your NGINX configuration.

Replace your server name and the location of your SSL certificate and key.

In this example, users will connect to Synchrony, which is required for collaborative editing, directly.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

```
server {
    listen www.example.com:80;
    server_name www.example.com;

    listen 443 default ssl;
    ssl_certificate     /usr/local/etc/nginx/ssl/nginx.crt;
    ssl_certificate_key /usr/local/etc/nginx/ssl/nginx.key;

    ssl_session_timeout  5m;

    ssl_protocols TLSv1 TLSv1.1 TLSv1.2;
    ssl_ciphers 'ECDHE-ECDSA-CHACHA20-POLY1305:ECDHE-RSA-CHACHA20-
POLY1305:ECDHE-ECDSA-AES128-GCM-SHA256:ECDHE-RSA-AES128-GCM-SHA256:
ECDHE-ECDSA-AES256-GCM-SHA384:ECDHE-RSA-AES256-GCM-SHA384:ECDHE-ECDSA-
AES128-SHA256:ECDHE-RSA-AES128-SHA256:ECDHE-ECDSA-AES128-SHA:ECDHE-RSA-
AES256-SHA384:ECDHE-RSA-AES128-SHA:ECDHE-ECDSA-AES256-SHA384:ECDHE-
ECDSA-AES256-SHA:ECDHE-RSA-AES256-SHA:ECDHE-ECDSA-DES-CBC3-SHA:ECDHE-
RSA-DES-CBC3-SHA:EDH-RSA-DES-CBC3-SHA:AES128-GCM-SHA256:AES256-GCM-
SHA384:AES128-SHA256:AES256-SHA256:AES128-SHA:AES256-SHA:DES-CBC3-
SHA:!DSS';
    ssl_prefer_server_ciphers   on;

location /synchrony {
        proxy_set_header X-Forwarded-Host $host;
        proxy_set_header X-Forwarded-Server $host;
        proxy_set_header X-Forwarded-For $proxy_add_x_forwarded_for;
        proxy_pass http://localhost:8091/synchrony;
        proxy_http_version 1.1;
        proxy_set_header Upgrade $http_upgrade;
        proxy_set_header Connection "Upgrade";
    }
  location /confluence {
        client_max_body_size 100m;
        proxy_set_header X-Forwarded-Host $host;
        proxy_set_header X-Forwarded-Server $host;
        proxy_set_header X-Forwarded-For $proxy_add_x_forwarded_for;
        proxy_pass http://localhost:8090/confluence;
    }
}
```

See https://nginx.org/en/docs/http/ngx_http_proxy_module.htmlfor more information.

Note: do not include`ssl on;` if you are configuring SSL and Confluence on the same server, as in this example.

If you're not sure what to include for`ssl_ciphers`,https://mozilla.github.io/server-side-tls/ssl-config-generator/is a useful resource.

If you experience 413Request Entity Too Large errors,make sure that the client_max_body_sizein the `/confluence`location block matches Confluence's maximum attachment size. You may also need to increase theclient_max_body_sizein the `/synchrony`location block if you experienceerrors when editing large pages.

If you plan to allow users to use the Confluence mobile app with your site, and you have configured a context path, as in the example above, you may also need to add the following line to your nginx configuration.

```
location /server-info.action {
     proxy_pass   http://localhost:8090/confluence/server-info.action;
}
```

If you're accessing Confluence via a sub-domain, your config will look like this:

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

```
server {
    listen confluence.example.com:80;
    server_name confluence.example.com;

    listen 443 default ssl;
    ssl_certificate     /usr/local/etc/nginx/ssl/nginx.crt;
    ssl_certificate_key /usr/local/etc/nginx/ssl/nginx.key;

    ssl_session_timeout  5m;

     ssl_protocols TLSv1 TLSv1.1 TLSv1.2;
    ssl_ciphers 'ECDHE-ECDSA-CHACHA20-POLY1305:ECDHE-RSA-CHACHA20-
POLY1305:ECDHE-ECDSA-AES128-GCM-SHA256:ECDHE-RSA-AES128-GCM-SHA256:
ECDHE-ECDSA-AES256-GCM-SHA384:ECDHE-RSA-AES256-GCM-SHA384:ECDHE-ECDSA-
AES128-SHA256:ECDHE-RSA-AES128-SHA256:ECDHE-ECDSA-AES128-SHA:ECDHE-RSA-
AES256-SHA384:ECDHE-RSA-AES128-SHA:ECDHE-ECDSA-AES256-SHA384:ECDHE-
ECDSA-AES256-SHA:ECDHE-RSA-AES256-SHA:ECDHE-ECDSA-DES-CBC3-SHA:ECDHE-
RSA-DES-CBC3-SHA:EDH-RSA-DES-CBC3-SHA:AES128-GCM-SHA256:AES256-GCM-
SHA384:AES128-SHA256:AES256-SHA256:AES128-SHA:AES256-SHA:DES-CBC3-
SHA:!DSS';
    ssl_prefer_server_ciphers   on;

    location / {
        client_max_body_size 100m;
        proxy_set_header X-Forwarded-Host $host;
        proxy_set_header X-Forwarded-Server $host;
        proxy_set_header X-Forwarded-For $proxy_add_x_forwarded_for;
        proxy_pass http://localhost:8090;
    }
    location /synchrony {
        proxy_set_header X-Forwarded-Host $host;
        proxy_set_header X-Forwarded-Server $host;
        proxy_set_header X-Forwarded-For $proxy_add_x_forwarded_for;
        proxy_pass http://localhost:8091/synchrony;
        proxy_http_version 1.1;
        proxy_set_header Upgrade $http_upgrade;
        proxy_set_header Connection "Upgrade";
    }
}
```

Step 4: RestartConfluence and NGINX

1. Restart Confluence and NGINX for all the changes to take affect.
2. Update Confluence's base URL to include the context path you set earlier - seeConfiguring the Server Base URL.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Running Confluence Over SSL or HTTPS

Atlassian applications can be accessed via HTTPS, however Atlassian Support does not provide assistance for configuring it. Consequently, Atlassian **cannot guarantee providing any support for 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 Community.

This page provides a basic outline of how to configure Confluence to enable access via HTTPS (HTTP Secure), so that your Confluence logins and data are encrypted during transport to and from Confluence. This isa good way to safeguard your Confluence data and user logins from being intercepted and read by outsiders.

In this article we use 'SSL' as a general term to refer to the protocol used to encrypt traffic. In most cases the protocol will be TLS.

On this page:

- Step 1. Create or request an SSL certificate
- Step 2. Modify your Confluence server.xml file
- Step 3. Specify the location of your certificate
- Step 4. Change your confluence base URL to HTTPS
- Step 5. Add a security constraint to redirect all URLs to HTTPS
- Notes
- Troubleshooting

These instructions cover terminating SSL at Tomcat, the application server shipped with Confluence.

If you want to terminate SSL at your web server or proxy, seeApache with mod_proxyorRunning Confluence behind NGINX withSSLfor examples of how to terminate SSL at an external web server.

You'll need the JDK for some of the steps in this guide. The JRE is not enough.

Running Confluence without HTTPS enabled may leave your site exposed to vulnerabilities, such as man-in-the-middle or DNS rebinding attacks. We recommend you enable HTTPS on your site.

Step 1. Create or request an SSL certificate

You'll need a valid certificate before you can enable HTTPS. If you already have a certificate, skip to step 2.

You can create your own self-signed certificate, or acquire one from a trusted Certificate Authority.

If your team plans to use the Confluence Server mobile app, you'll need a certificate issued by a trusted Certificate Authority. You can't use the app with a self-signed certificate, or one from an untrusted or private CA.

Option 1: Create a self-signed certificate

Self-signed certificates are useful if you require encryption but don't need to verify the identity of the requesting website. In general, you might use a self-signed certificate on a test environment and on internal corporate networks (intranets).

Because the certificate is not signed by a certificate authority (CA), users may receive a message that the site is not trusted and may have to perform several steps to accept the certificate before they can access the site. This usually will only occur the first time they access the site. Users won't be able to log in to your site at all via the Confluence Server mobile app if you use a self-signed certificate.

In this example, we'll use Java's `keytool` utility, which is included with the JDK.If you're not comfortable using command line utilitiesKeyStore Exploreris a usefulalternative to the command line.

To generate a self-signed certificate using `keytool`:

1. From the command line, run the appropriate command for your operating system:

Windows

```
"%JAVA_HOME%\bin\keytool" -genkeypair -keysize 2048 -alias tomcat -keyalg RSA -sigalg
SHA256withRSA
```

Linux (and MacOS)

```
$JAVA_HOME/bin/keytool -genkeypair -keysize 2048 -alias tomcat -keyalg RSA -sigalg SHA256withRSA
```

2. When prompted, createa **password**for the certificate (private key).
    - Only use alphanumeric characters. Tomcat has a known issue with special characters.
    - Make a note of the password, you'll need it in the next step.
    - The default password is '`changeit`'.
3. Follow the prompts to specify the certificate details.This info is used to construct the X.500 Distinguished Name (DN) of the entity.
    - **First and last name**: this is not your name, it is the Common Name (CN), for example 'confluence.example.com'. The CN must match the fully qualified hostname of the server running Confluence, or Tomcat won't be able to use the certificate for SSL.
    - **Organizational unit**: this is the team or department requesting the certificate, for example 'marketing'.
    - **Organization**: this is your company name, for example 'SeeSpaceEZ'.
    - **City, State / province, country code**: this is where you're located, for example Sydney, NSW, AU.
4. The output will look something like the example below. Hit 'y' to confirm the details.

```
CN=confluence.example.com, OU=Marketing, O=SeeSpaceEZ, L=Sydney, ST=NSW, C=AU
```

5. When asked for the **password** for 'tomcat', enter the password you created in step 2 (or hit return to use the same .
    - 'tomcat' is the alias we entered in the keytool command above, it refers to your.
    - Your keystore entry must have the same password as your private key. This is a Tomcat requirement.
6. You certificate is now ready. Go to step 2 below.

Option 2: Use a certificate issued by a Certificate Authority (recommended)

Production environments will need a certificate issued by a Certificate Authority (CA). These instructions are adapted from the Tomcat documentation.

First you will generate a local certificate and create a 'certificate signing request' (CSR) based on that certificate. You will submit the CSR to your chosen certificate authority. The CA will use that CSR to generate a certificate for you.

1. Use Java's `keytool` utility to generate a local certificate (follow the steps in option 1, above).
2. From the command line, run the following command to generate a certificate signing request.

```
keytool -certreq -keyalg RSA -alias tomcat -file certreq.csr -keystore <MY_KEYSTORE_FILENAME>
```


Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Replace<MY_KEYSTORE_FILENAME>with the path to and file name of the.keystorefile generated for your local certificate.

3. Submit the generated file called `certreq.csr` to your chosen certificate authority.
   Check your CA's documentation to find out how to do this.
4. The CA will send you a certificate.
5. Import the new certificate into your local keystore:

```
keytool -importcert -alias tomcat -keystore <MY_KEYSTORE_FILENAME> -file <MY_CERTIFICATE_FILENAME>
```

Some CAs require you to install an intermediate certificate before importing your certificate. You should follow your CA documentation to successfully install your certificate.

If you receive an error, and you use Verisign or GoDaddy, you mayneed to export the certificate to PKCS12 format along with the private key.

1. First, remove thecertificate added above from the keystore:

```
keytool -delete -alias tomcat -keystore <MY_KEYSTORE_FILENAME>
```

2. Then export to PKCS12 format:

```
openssl pkcs12 -export -in <MY_CERTIFICATE_NAME> -inkey <MY_PRIVATEKEY_NAME> -out
<MY_PKC12_KEYSTORE_NAME> -name tomcat -CAfile <MY_ROOTCERTIFICATE_NAME-
alsoCalledBundleCertificateInGoDaddy> -caname root
```

3. Then import from PKCS12to jks:

```
keytool -importkeystore -deststorepass <MY_DESTINATIONSTORE_PASSWORD> -destkeypass
<MY_DESTINATIONKEY_PASSWORD> -destkeystore <MY_KEYSTORE_FILENAME> -srckeystore
<MY_PKC12_KEYSTORE_NAME> -srcstoretype PKCS12 -srcstorepass <MY_PKC12_KEYSTORE_PASSWORD> -
alias tomcat
```

Step 2. Modify your Confluence server.xml file

The next step is to configure Confluence to use HTTPS.

1. Edit `<install-directory>/conf/server.xml`.
2. Uncomment the following lines:

```
<Connector port="8443" maxHttpHeaderSize="8192"
   maxThreads="150" minSpareThreads="25"
   protocol="org.apache.coyote.http11.Http11Nio2Protocol"
   enableLookups="false" disableUploadTimeout="true"
   acceptCount="100" scheme="https" secure="true"
   clientAuth="false" sslProtocol="TLSv1.2"
   sslEnabledProtocols="TLSv1.2" SSLEnabled="true"
   URIEncoding="UTF-8" keystorePass="<MY_CERTIFICATE_PASSWORD>"/>
```

3. Replace <MY_CERTIFICATE_PASSWORD> with the password you specified for your certificate.
4. Make sure that the attribute-value pair `SSLEnabled="true"` is part of the `Connector` element, as shown above. If this attribute is not present, attempts to access Confluence will time out.
5. Change the value of`maxThreads`to be at least 10 threads (or 25%) less than the size of your database connection pool. 48 is usually about right. SeeHTTP MaxThreads configurationfor more information about this.
6. Save the server configuration file.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

⚠Don't remove or comment out the `http` connector, as the Synchrony proxy health check, still requires HTTP.If you don't want to include the `http`connector, you can use the`synchrony.proxy.healthcheck.disabled`system property to disable the health check.

You should also **not** disable the internal Synchrony proxy (by setting the`synchrony.proxy.enabled`system property to `false`) as this is known to cause problems when you're terminating SSL at Tomcat.

ℹThe default connector port for Confluence is 8090.

The Confluence mobile app requires minimum TLS 1.2.

Step 3. Specify the location of your certificate

By default, Tomcat expects the keystore file to be named `.keystore` and to be located in the user home directory under which Tomcat is running (which may or may not be the same as your own home directory). This means that, by default, Tomcat will look for your SSL certificates in the following location:

- On Windows: `C:\users\#CURRENT_USER#\.keystore`
- On OS X and UNIX-based systems: `~/.keystore`

⚠ Don't store your keystore file in your Confluence installation directory as the contents of that directory are removed when you upgrade Confluence.

You may decide to move the certificate to a custom location. If your certificate is not in the default location, you'll need to update your server configuration file as outlined below, so that Tomcat can find the certificate.

1. Edit`<confluence-install-directory>/conf/server.xml`
2. Add the attribute `keystoreFile="<MY_CERTIFICATE_LOCATION>"` to the `Connector`element, so that the element looks like this:

```
<Connector port="8443" maxHttpHeaderSize="8192"
   maxThreads="150" minSpareThreads="25"
   protocol="org.apache.coyote.http11.Http11Nio2Protocol"
   enableLookups="false" disableUploadTimeout="true"
   acceptCount="100" scheme="https" secure="true"
   clientAuth="false" sslProtocol="TLSv1.2"
   sslEnabledProtocols="TLSv1.2" SSLEnabled="true"
   URIEncoding="UTF-8" keystorePass="<MY_CERTIFICATE_PASSWORD>"
   keystoreFile="<MY_CERTIFICATE_LOCATION>"/>
```

3. Replace the text `<MY_CERTIFICATE_LOCATION>` with the path to your certificate, including the path and the name of the `.keystore` file.
4. Save the configuration file.

Step 4. Change your confluence base URL to HTTPS

1. In your browser, go to > **General Configuration**.
2. Click **Edit**.
3. Change the Server Base URL to HTTPS. See the documentation on configuring the server base URL.
4. Restart Confluenceand access Confluence on `https://<MY_BASE_URL>:8443/`.

Step 5. Add a security constraint to redirect all URLs to HTTPS

Although HTTPS is now activated and available, the old HTTP URLs (http://localhost:8090) are still available. Now you need to redirect the URLs to their HTTPS equivalent. You will do this by adding a security constraint in `web.xml`. This will cause Tomcat to redirect requests that come in on a non-SSL port.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

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.

Confluence has two web.xml files. The other one is at `<CONFLUENCE_INSTALLATION>/conf/web.xml`. Please only add the security constraints to `<CONFLUENCE_INSTALLATION>/confluence/WEB-INF/web.xml`, as described above.

Notes

- **Background information on generating a certificate:** The '`keytool -genkeypair`' command generates a key pair consisting of a public key and the associated private key, and stores them in a keystore. The command packages the public key into an X.509 v3 self-signed certificate, which is stored as a single-element certificate chain. This certificate chain and the private key are stored in a new keystore entry, identified by the `alias` that you specify in the command. The Java 11 documentation has a good overview of the utility.

- **Custom SSL port:** If you have changed the port that the SSL connector is running on from the default value of 8443, you must update the `redirectPort` attribute of the standard HTTP connector to reflect the new SSL port. Tomcat needs this information to know which port to redirect to when an incoming request needs to be secure.
- **Multiple instances on the same host:** When running more than one instance on the same host, it is important to specify the *address* attribute in the `<CONFLUENCE_INSTALLATION>/conf/server.xml` file because by default the connector will listen on **all available network interfaces**, so specifying the address will prevent conflicts with connectors running on the same default port. See the Tomcat Connector documentation for more about setting the address attribute:

```
<Connector port="8443" address="your.confluence.url.com"
   maxHttpHeaderSize="8192" maxThreads="150" minSpareThreads="25"
   protocol="org.apache.coyote.http11.Http11Nio2Protocol"
   enableLookups="false" disableUploadTimeout="true"
   acceptCount="100" scheme="https" secure="true"
   clientAuth="false" sslProtocol="TLSv1.2"
   sslEnabledProtocols="TLSv1.2" SSLEnabled="true"
   URIEncoding="UTF-8" keystorePass="<MY_CERTIFICATE_PASSWORD>"
   keystoreFile="<MY_CERTIFICATE_LOCATION>"/>
```

- **HTTPS must be configured for your whole site.** It can't be enabled for individual pages or spaces.
- **Before you upgrade Confluence**, make a note of the changes you have made to your `server.xml`and `web.xml`files. It is always best to re-apply these changes manually after upgrading, rather than copying over your existing files.
- **TLS 1.2 or 1.3 recommended**. The Confluence Server mobile app requires TLS 1.2. If you use Jira and Confluence together, we recommend configuring both applications to use the same TLS version.

Troubleshooting

- Check the Confluence knowledge base articles on troubleshooting SSL
- SSL Configuration HOW-TOin the Apache Tomcat 9.0 documentation

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

- keytool - Key and Certificate Management Toolin the Java 11 documentation

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Using Apache to limit access to the Confluence administration interface

As well as limiting access to the Confluence administration console to users who really need it, and using strong passwords, you can consider limiting access to certain machines on the network or internet. If you are usingApache web server, this can be done with Apache's **Location** functionality.

To limit access to admin screens to specific IP addresses in Apache:

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. In your Apache Virtual Host, add the following lines to restrict the administration actions to the Systems Administrator.

```
Define segmentregex (?:;[^/]*)?(?:/)?(?:(?:;[^/]*)?(?:/)?)*
<LocationMatch (?i)/confluence${segmentregex}/admin>
  Include sysadmin_ips_only.conf
</LocationMatch>
<LocationMatch (?i)/confluence${segmentregex}/plugins${segmentregex}/servlet${segmentregex}
/oauth${segmentregex}/consumers${segmentregex}/list>
   Include sysadmin_ips_only.conf
</LocationMatch>
<LocationMatch (?i)/confluence${segmentregex}/plugins${segmentregex}/servlet${segmentregex}
/oauth${segmentregex}/view-consumer-info>
   Include sysadmin_ips_only.conf
</LocationMatch>
<LocationMatch (?i)/confluence${segmentregex}/plugins${segmentregex}/servlet${segmentregex}
/oauth${segmentregex}/service-providers${segmentregex}/list>
   Include sysadmin_ips_only.conf
</LocationMatch>
<LocationMatch (?i)/confluence${segmentregex}/plugins${segmentregex}/servlet${segmentregex}
/oauth${segmentregex}/service-providers${segmentregex}/add>
   Include sysadmin_ips_only.conf
</LocationMatch>
<LocationMatch (?i)/confluence${segmentregex}/plugins${segmentregex}/servlet${segmentregex}
/oauth${segmentregex}/consumers${segmentregex}/add>
   Include sysadmin_ips_only.conf
</LocationMatch>
<LocationMatch (?i)/confluence${segmentregex}/plugins${segmentregex}/servlet${segmentregex}
/oauth${segmentregex}/consumers${segmentregex}/add-manually>
   Include sysadmin_ips_only.conf
</LocationMatch>
<LocationMatch (?i)/confluence${segmentregex}/plugins${segmentregex}/servlet${segmentregex}
/oauth${segmentregex}/update-consumer-info>
   Include sysadmin_ips_only.conf
</LocationMatch>
<LocationMatch (?i)/confluence${segmentregex}/pages${segmentregex}/templates${segmentregex}
/listpagetemplates.action>
   Include sysadmin_ips_only.conf
</LocationMatch>
<LocationMatch (?i)/confluence${segmentregex}/pages${segmentregex}/templates${segmentregex}
/createpagetemplate.action>
   Include sysadmin_ips_only.conf
</LocationMatch>
<LocationMatch (?i)/confluence${segmentregex}/spaces${segmentregex}/spacepermissions.action>
   Include sysadmin_ips_only.conf
</LocationMatch>
<LocationMatch (?i)/confluence${segmentregex}/pages${segmentregex}/listpermissionpages.action>
   Include sysadmin_ips_only.conf
</LocationMatch>
<LocationMatch (?i)/confluence${segmentregex}/spaces${segmentregex}/removespace.action>
        Include sysadmin_ips_only.conf
```

```
</LocationMatch>
<LocationMatch (?i)/confluence${segmentregex}/spaces${segmentregex}/importmbox.action>
        Include sysadmin_ips_only.conf
</LocationMatch>
<LocationMatch (?i)/confluence${segmentregex}/spaces${segmentregex}/viewmailaccounts.action>
   Include sysadmin_ips_only.conf
</LocationMatch>
<LocationMatch (?i)/confluence${segmentregex}/spaces${segmentregex}/addmailaccount.action?>
   Include sysadmin_ips_only.conf
</LocationMatch>
<LocationMatch (?i)/confluence${segmentregex}/spaces${segmentregex}/importpages.action>
   Include sysadmin_ips_only.conf
</LocationMatch>
<LocationMatch (?i)/confluence${segmentregex}/spaces${segmentregex}/flyingpdf${segmentregex}
/flyingpdf.action>
   Include sysadmin_ips_only.conf
</LocationMatch>
<LocationMatch (?i)/confluence${segmentregex}/spaces${segmentregex}/exportspacehtml.action>
   Include sysadmin_ips_only.conf
</LocationMatch>
<LocationMatch (?i)/confluence${segmentregex}/spaces${segmentregex}/exportspacexml.action>
   Include sysadmin_ips_only.conf
</LocationMatch>
<LocationMatch (?i)/confluence${segmentregex}/plugins${segmentregex}/servlet${segmentregex}/embedded-
crowd>
   Include sysadmin_ips_only.conf
</LocationMatch>
<LocationMatch (?i)/confluence${segmentregex}/plugins${segmentregex}/servlet${segmentregex}/upm>
   Include sysadmin_ips_only.conf
</LocationMatch>
```

⚠ This configuration assumes you're running Confluence with the context path '/confluence'. If you are running with a different context path, or no context path, adjust the sample above accordingly.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Using Apache with mod_jk

It's not possible to use only mod_jk to proxy Confluence 6.0 or later. This is because Synchrony, which is required for collaborative editing, cannot accept AJP connections. The preferred configuration isUsing Apache with mod_proxy.

If you are unable to switch to mod_proxy, see [ARCHIVED] How to configure Apache mod_jk to proxy Confluence 6.x or later for a workaround.

Using mod_rewrite to Modify Confluence URLs

Note: This page documents a configuration of Apache, rather than of Confluence itself. Atlassian will support Confluence with this configuration, but we cannot guarantee to help you debug problems with Apache. Please be aware that this material is provided for your information only, and that you use it at your own risk.

Confluence requires URL rewriting for proper functionality, if Confluence is accessible via different domain names. If Confluence is configured for multiple domains *without* URL rewriting, you will experience an array of problems. See Various Issues Caused when Server Base URL Does Not Match the URL Used to Access Confluence.

An example of why you may want to access Confluence from different domains:

- From an internal network:
  http://wiki
- The externally visible domain:
  http://wiki.domain.com

Using URL rewriting to access Confluence over multiple domains

To configure Confluence over multiple domains:

1. Add a DNS entry mapping http://wiki to the externally visible IP address of the Confluence server.
2. Set Confluence's server base URL to http://wiki.domain.com.
3. Add Apache HTTP proxy, using the instructions from Running Confluence behind Apache.
4. Add the mod_rewrite module to change the URL.

Further information

You may be interested in the UrlRewriteFilter that is Java web filter that works in a similar way of the **Apache's mod_rewrite**.

Configuring Secure Administrator Sessions

Secure administrator sessions allows you to require administrators to re-enter their password before they can access administrative functions. This feature is sometimes known as "websudo" and is turned on by default.

Start a secure administrator session

When an administrator attempts to access an admin function (including some space admin functions like delete space), they will be prompted to re-enter their password. This starts the secure administrator session.

Administrators can click **Drop access** in the banner to manually end the session. This won't log them out of Confluence, it will just end the secure administrator session.

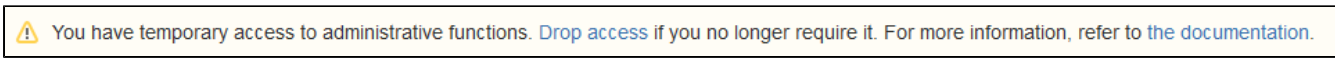

Change the secure administrator session timeout

The secure administrator session has a rolling timeout which defaults to 10 minutes. If there's no activity for a period of time, the administrator will be logged out of the session. They'll remain logged in to Confluence.

To change the timeout value:

1. Go to > **General Configuration**> **Security Configuration**.
2. Select **Edit**.
3. Under **Secure administrator sessions**, enter the**Minutes before automatic invalidation**.
4. **Save** your changes.

Turn off secure administrator sessions

If you're using single sign-on, or have other security measures in place, you may want to disable secure administrator sessions. We don't recommend doing this unless you need to.

To turn off secure administrator sessions:

1. Go to > **General Configuration**> **Security Configuration**.
2. Select **Edit**.
3. Under **Secure administrator sessions**, deselect the **Enable** checkbox.
4. **Save** your changes.

Troubleshooting

Known issues with single sign-on and just-in-time user provisioning

You may need to disable secure administrator sessions if your users are not stored in Confluence's internal user directory. See

CONFSERVER-60263 - Ability to have the Websudo functionality working with SAML / SSO
GATHERING INTEREST

for more information and some suggested workarounds.

Known issues for app developers

Secure administrator sessions can cause exceptions when developing against Confluence or deploying a plugin. SeeHow do I develop against Confluence with Secure Administrator Sessions?

Note that REST and XML-RPC APIs are not affected by secure administration sessions.

Confluence Cookies

This page lists cookies stored in Confluence users' browsers which are generated by Confluence itself. This page does not list cookies that may originate from 3rd-partyConfluence plugins.

On this page:

- Authentication cookies
- The 'remember me' cookie
- Other Confluence cookies

Authentication cookies

Confluence uses Seraph, an open source framework, for HTTP cookie authentication. Confluence uses two types of cookies for user authentication:

- The JSESSIONID cookie is created by the application server and used for session tracking purposes. This cookie contains a random string and the cookie expires at the end of every session or when the browser is closed. This cookie is ignored when Confluence is clustered.
- The 'remember me' cookie, `seraph.confluence`, is generated by Confluence when the user selects the **Remember me** check box on the login page. Remember me is enforced by default when Confluence is clustered.

You can read about cookies on the Wikipedia page about HTTP cookies.

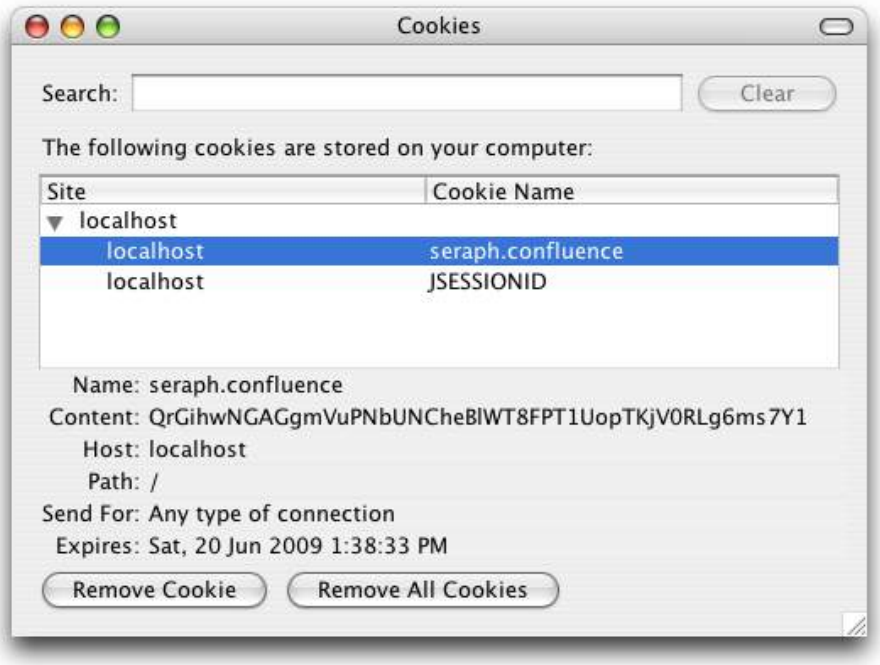

The 'remember me' cookie

The 'remember me' cookie, seraph.confluence, is a long-lived HTTP cookie. This cookie can be used to authenticate an unauthenticated session. Confluence generates this cookie when the user selects the Remember me check box on the login page.

The default time to live of this cookie is two weeks.

When running Confluence in a cluster,Remember me is enabled by default to allow users to move seamlessly between nodes.See How to configure the 'Remember Me' feature in Confluence if you need to configure this.

Cookie key and contents

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 theConfluence database.

Use of cookie for authentication

When a user requests a web page, if the request is not already authenticated via session-based authentication or otherwise, Confluence will match the 'remember me' cookie (if present) against the token (also if present), which is stored for the user in the Confluence database.

If the token in the cookie matches the token stored in the database and the cookie has not expired, the user is authenticated.

Life of 'remember me' cookies

You can configure the maximum age of the cookie. To do that you will need to modify the `CONFLUENCE-INSTALLATION/confluence/WEB-INF/classes/seraph-config.xml` file and insert the following lines below the other `init-param` elements:

```
<init-param>
    <param-name>autologin.cookie.age</param-name>
    <param-value>259200</param-value><!-- 3 days in seconds -->
</init-param>
```

Automatic cleanup of 'remember me' tokens

Every cookie issued by Confluence has a corresponding record in the database. A scheduled job runs on the 20th of every month to clean up expired tokens. The name of the trigger is `clearExpiredRememberMeTokensTrigger`.

Note: The only purpose of this job is to prevent the database table from growing too big. For authentication purposes, Confluence will ignore expired tokens even if they still exist in the database.

Is it possible to disable the 'remember me' feature?

Confluence does not offer an option for disabling the 'Remember Me' feature. See the workaround.

Other Confluence cookies

There are several cookies that Confluence uses to store basic 'product presentation' states.Confluence users' authentication details are not stored by these cookies.

Cookie Key	Purpose	Cookie Contents	Expiry
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.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

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.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Using Fail2Ban to limit login attempts

What is Fail2Ban?

We need a means of defending sites against brute-force login attempts. Fail2Ban is a Python application which trails logfiles, looks for regular expressions and works with Shorewall (or directly with iptables) to apply temporary blacklists against addresses that match a pattern too often. This can be used to limit the rate at which a given machine hits login URLs for Confluence.

Prerequisites

- Requires Python 2.4 or higher to be installed
- Requires Apache Reverse Proxy to be installed
- Needs a specific file to follow, which means your Apache instance needs to log your Confluence access to a known logfile. You **should adjust the configuration below** appropriately.

How to set it up

This list is a skeletal version of the instructions

- There's an RPM available for RHEL on the download page, but you can also download the source and set it up manually
- Its configuration files go into `/etc/fail2ban`
- The generic, default configuration goes into `.conf` files (`fail2ban.conf` and `jail.conf`). Don't change these, as it makes upgrading difficult.
- Overrides to the generic configuration go into `.local` files corresponding to the `.conf` files. These only need to contain the specific settings you want overridden, which helps maintainability.
- Filters go into `filter.d` this is where you define regexps, each going into its own file
- Actions go into `action.d` you probably won't need to add one, but it's handy to know what's available
- "jails" are a configuration unit that specify one regexp to check, and one or more actions to trigger when the threshold is reached, plus the threshold settings (e.g. more than 3 matches in 60 seconds causes that address to be blocked for 600 seconds)
- Jails are defined in `jail.conf` and `jail.local`. Don't forget the `enabled` setting for each one it can be as bad to have the wrong ones enabled as to have the right ones disabled.

Running Fail2Ban

- Use `/etc/init.d/fail2ban {start|stop|status}` for the obvious operations
- Use `fail2ban-client -d` to get it to dump its current configuration to STDOUT. Very useful for troubleshooting.
- Mind the CPU usage; it can soak up resources pretty quickly on a busy site, even with simple regexp
- It can log either to syslog or a file, whichever suits your needs better

Common Configuration

jail.local

```
# The DEFAULT allows a global definition of the options. They can be override
# in each jail afterwards.

[DEFAULT]

# "ignoreip" can be an IP address, a CIDR mask or a DNS host. Fail2ban will not
# ban a host which matches an address in this list. Several addresses can be
# defined using space separator.
# ignoreip = <space-separated list of IPs>

# "bantime" is the number of seconds that a host is banned.
bantime  = 600

# A host is banned if it has generated "maxretry" during the last "findtime"
# seconds.
findtime  = 60

# "maxretry" is the number of failures before a host get banned.
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 =
```


Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Securing Confluence with Apache

When opened in a viewport, the user will be redirected to: Proxy and HTTPS setup for Confluence.

Best Practices for Configuring Confluence Security

This page outlines a number of approaches you can use to make your Confluence site as secure as possible.There are many things to consider, such as the configuration of your underlying operating systems, application servers, database servers, network, firewall, routers, etc. It would be impossible to outline all of them here.

On this page:

Best practices

Not everything in this guide will be applicable to your environment, but the principles described can be adapted to most environments.

It's worth noting that none of these practices can provide 100% security. They are measures to reduce impact and to slow down an intruder in case your system does become compromised.

Subscribe to advisory alerts

Subscribe to advisory alerts and keep technical contact details up to date to make sure you receive security advisory alerts and other important technical updates.

Atlassian email and privacy preferences

Secure your installation and data directories

It's important to make sure your Confluence installation directory, home directory, and any storage locations you may define for attachments, space exports, or data pipeline exports are secure.

We strongly recommend you:

- run Confluence with a dedicated non-root user account.
- limit the user accounts who can access any Confluence directories.

To find out how to do this, see Creating a Dedicated User Account on the Operating System to Run Confluence.

You should also monitor your binaries. If an attacker compromises an account on your system, they will usually try to gain access to more accounts. This is sometimes done by adding malicious code, such as by modifying files on the system. Consider how you might regularly verify that no malicious changes have been made.

Secure your database

Make sure the Confluence database user (and all datasource database users) only have the amount of database privilege they really need.

Limit database access to just the Confluence host (using iptables or built in database security tools).Refer to your database documentation to find out how to do this.

Limit access to administrator functions

As a general rule, you shouldkeep the number of Confluence administrators as low as possible, and review these user accounts every so often to make sure the access level is still appropriate.

- Avoid shared administrator or user accounts, and easily guessed usernames like 'admin'.
- Provide administrators with two separate accounts, to allow them to separate day-to-day work such as creating pages, from administration tasks.
- Limit the number of people in the confluence-administrators group.Members of this 'super group' can access all admin functions and access all content, including restricted pages. Consider limiting the

members of this group and instead create a new group with system administrator global permissions.
Learn about the confluence-administrators super group

- Use secure administrator sessions to require admins to re-enter their password to access admin functions, and set a short timeout for the administrator session.
Learn how to turn on secure administrator sessions

- Use Apache to lock down the administration interface to specific IP addresses. This can be used as a template for your reverse proxy of choice.
Using Apache to limit access to the Confluence administration interface.

Limit incoming and outgoing connections

There are a number of ways you can limit incoming and outgoing connections, including using firewalls and proxy servers.

- Use the allowlist tolimit incoming and outgoing connections to avoid Server-Side Request Forgery (SSRF) attacks. Confluence relies on the allowlist for things like macros, that may be displaying content from external sites.
Learn how to turn on the allowlist

Manage user accounts

Good user management practices can help prevent user accounts being compromised.

- Consider integrating Confluence with an identity provider for single sign-on and two-factor authentication.
Learn about the various SSO options available

- Use personal access tokens for integrations. This provides your users a more secure way to authenticate API requests than basic authentication (username and password)
Learn how to manage personal access tokens

- Disable basic authentication. If you've configured single sign-on, you can disable basic authentication for login and REST requests. Basic authentication is less secure than single sign-on and personal access tokens.Learn how to disable basic authentication

- Disable user accounts when people leave your organisation. If required, you can also delete user accounts which will replace their name with an anonymized ID.
Delete and disable users

- Restrict the number of users with powerful roles or group memberships. If only one department should have access to particularly sensitive data, then restrict access to the data to only those users. Do not let convenience over-rule security. Do not give all staff access to sensitive data when there is no need.

Limit login attempts and monitor access and activity

There are several things you can do to reduce the risk of brute force or denial of service attacks.

- Fail2Ban can help you reduce the risk of brute force attacks.
Using Fail2Ban to limit login attempts

- Use rate limiting to block REST API requests from anonymous users if you don't have a reason to allow them, or limit the number of requests to reduce the risk of DoS attacks
Learn how to use rate limiting to block requests

- Review your audit log settings, to make sure you're logging important administrator and end user actions.
Learn which events you can write to the audit log

- Access logs can help you identify unusual activity. Logs are written to the install directory, and you may want to monitor these logs using your preferred monitoring tool.
Learn about access logging

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Perform regular security audits

Regular security audits can help you identify potential threats, and also provide an opportunity to review your security policies and procedures.

- Know who can help if a security breach occurs. What is the process if a potential threat is identified?
- Perform 'what if' planning exercises. Consider questions like 'What is the worst thing that could happen if a privileged user's password were stolen while on vacation? What can we do to minimize damage?'.
- Document your security measures, and regularly monitor that all measures are still in place, and are adequate. For example after upgrading or migrating, someone may forget to apply the rule to the new system or version.
- Perform a security check-up when preparing for a major upgrade. It's a good time to check your current configuration against our current recommendations.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Hiding the People Directory

The People Directory provides a list of all users in your Confluence system.

If you need to disable the People Directoryset the following system propertieson your application server command line:

- **To disable the People Directory foranonymous users:**

```
-Dconfluence.disable.peopledirectory.anonymous=true
```

- **To disable the People Directory entirely:**

```
-Dconfluence.disable.peopledirectory.all=true
```

This workaround will prevent the People directory from appearing on the dashboard, but if you navigate to the profile of a user, and then click on the "People" in the breadcrumb link (Dashboard >> **People** >> FullName >> Profile) or you go to the URL directly `<CONFLUENCE_INSTALL>/browsepeople.action`, you will be able to access the people directory.

To workaround this, set up your Apache webserver in front of Confluence and redirect requests to this URL.

Configuring Captcha for Spam Prevention

If your Confluence site is open to the public (you allow anonymous users to add comments, create pages etc) you may find that automated spam is being added, in the form of comments or new pages.

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**
2. Choose**Spam Prevention**in the left-hand panel
3. Choose**ON**to turn on Captcha
4. If you want to disable Captcha for certain groups:
    - Select**No one**if you want everyone to see Captchas.
    - Select**Signed in users**if you want only anonymous users to see Captchas.
    - If you want everyone to see Captchas except members of specific groups, select**Members of the following groups**and enter the group names in the text box.
      You can click the magnifying-glass icon to search for groups. Search for all or part of a group name and click the**Select Groups**button to add one or more groups to the list.
    - To remove a group from the list, delete the group name
5. Choose**Save**

Hiding External Links From Search Engines

Hiding external links from search engines helps to discourage spammers from posting links on your site. If you turn this option on, any URLs inserted in pages and comments will be given the 'nofollow' attribute, which prevents search engines from following them.

Shortcut links (e.g. CONF-2622@JIRA) and internal links to other pages within Confluence are not tagged.

To hide external links from search engines:

1. Choose the**cog icon** , then choose**General Configuration**
2. Click '**Security Configuration**' in the left panel.
3. This will display the '**Security Configuration**' screen. Click '**Edit**'.
4. Check the '**Hide External Links From Search Engines**' checkbox.
5. Click the '**Save**' button.

Background to the nofollow attribute

As part of the effort to combat the spamming of wikis and blogs (Confluence being both), Google came up with some markup which instructs search engines not to follow links. By removing the main benefit of wiki-spamming it's hoped that the practice will stop being cost-effective and eventually die out.

Configuring Captcha for Failed Logins

If you have confluence administrator permissions, you can configure Confluence to impose a maximum number of repeated login attempts. After a given number of failed login attempts (the default is three) Confluence will display a Captcha form asking the user to enter a given word when attempting to log in again. This will prevent brute force attacks on the Confluence login screen.

Similarly, after three failed login attempts via the XML-RPC or SOAP API, an error message will be returned instructing the user to log in via the web interface. Captcha will automatically be activated when they attempt this login.

On this page:

- Enabling, Disabling and Configuring Captcha for Failed Logins
- Notes

'Captcha' is a test that can distinguish a human being from an automated agent such as a web spider or robot.When Captcha is activated, users will need to recognize a distorted picture of a word, and must type the word into a text field. This is easy for humans to do, but very difficult for computers.

Screenshot: example of a Captcha test

Enabling, Disabling and Configuring Captcha for Failed Logins

By default, Captcha for failed logins is enabled and the number of failed login attempts is set to three.

To enable, disable and configure Captcha for failed logins:

1. Choose the**cog icon** , then choose**General Configuration**
2. Choose '**Security Configuration**' from the left menu.
3. Choose '**Edit**'.
4. To enable Captcha:
    - Select the '**Enable**' checkbox next to '**CAPTCHA on login**'.
    - Set the maximum number of failed logins next to '**Maximum Authentication Attempts Allowed**'. You must enter a number greater than zero.
5. To **disable** Captcha, deselect the '**Enable**' checkbox.
6. Choose '**Save**'.

Screenshot: Configuring Captcha for failed logins

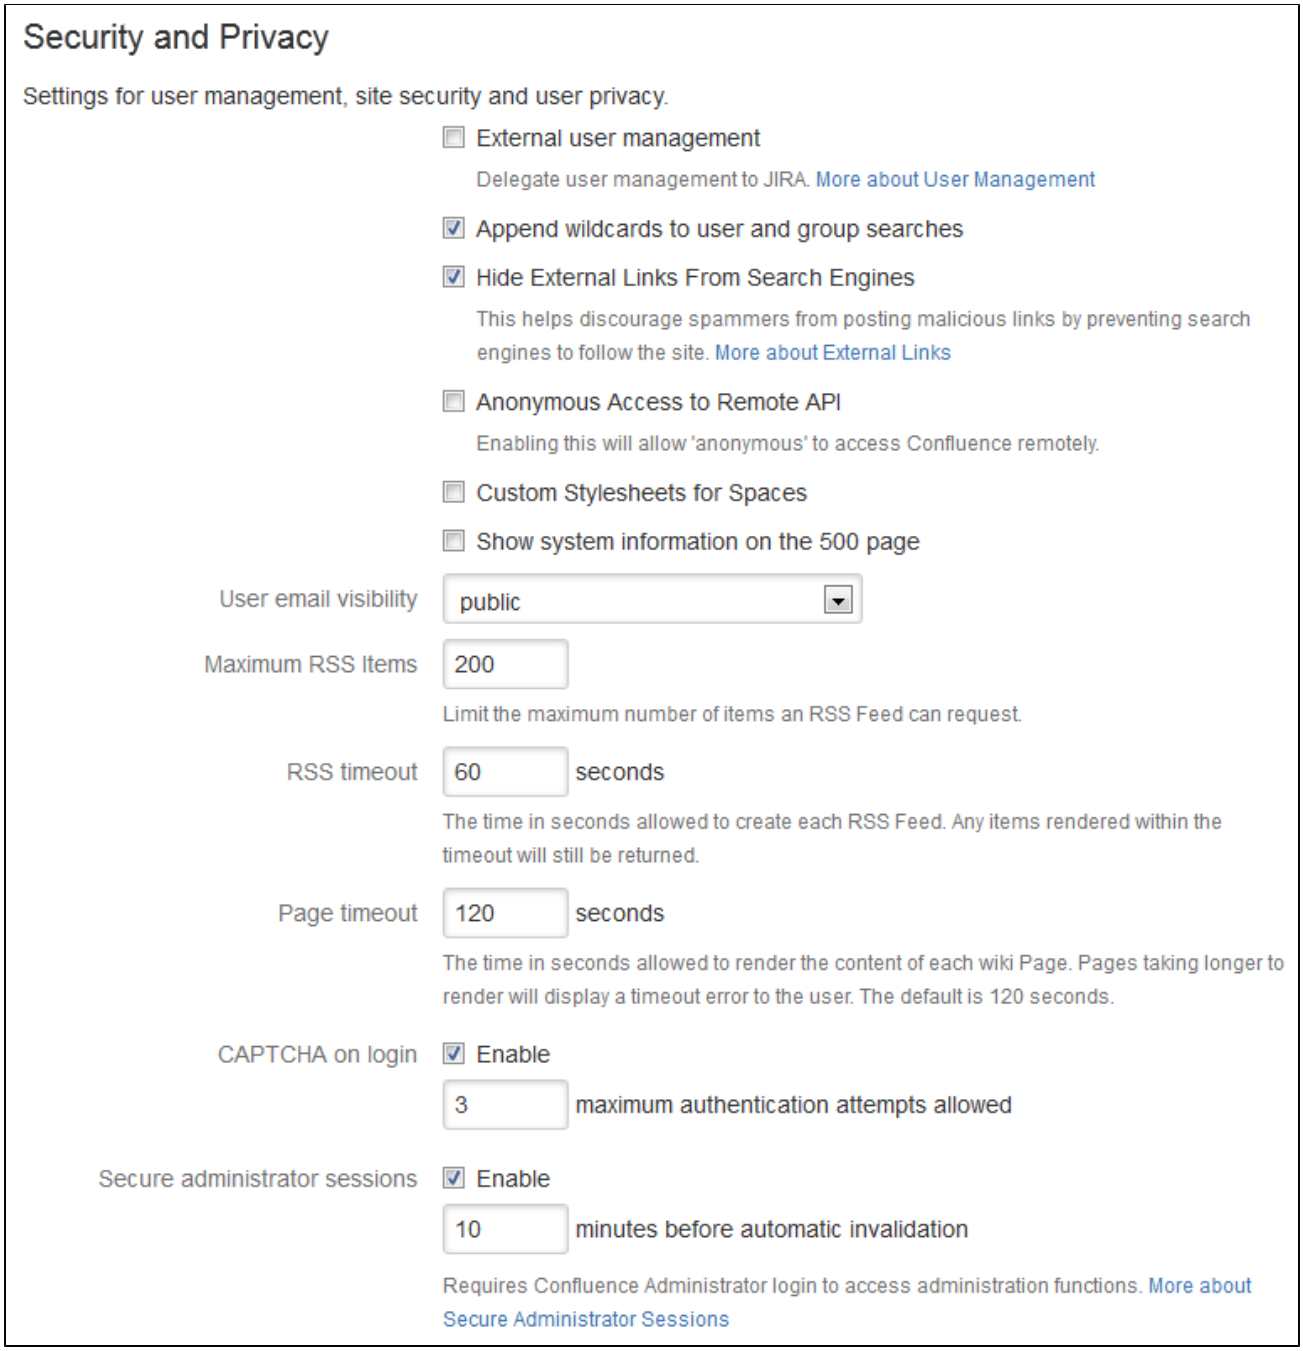

Notes

- **Disabling all password confirmation requests, including Captcha on login.** Confluence installations that use a custom authentication mechanism may run into problems with the Confluence security measure that requires password confirmation. If necessary, you can set the `password.confirmation.disabled` system property to disable the password confirmation functionality on ad ministrative actions, change of email address and Captcha for failed logins. See Recognized System Properties.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Configuring XSRF Protection

Confluence requires an XSRF token to be present on comment creation, to prevent users being tricked into unintentionally submitting malicious data. All the themes bundled with Confluence have been designed to use this feature. However, if you are using a custom theme that does not support this security feature, you can disable it.

⚠️ **Please carefully consider the security risks before you disable XSRF protection for comments in your Confluence installation.**

Read more about XSRF (Cross Site Request Forgery) at cgisecurity.com.

To configure XSRF protection for comments:

1. Choose the**cog icon** , then choose**General Configuration**
2. Choose **Security Configuration** in the left-hand panel.
3. Choose **Edit**.
4. Uncheck the **Adding Comments** checkbox in the **XSRF Protection** section, to disable XSRF protection.
5. Choose **Save**.

Related pages:

- Configuring Confluence Security
- Confluence Administrator's Guide
- Developer documentation on XSRF protection in Confluence

User Email Visibility

Confluence provides three options for email address privacy which can be configured by a Confluence administrator from the **Administration Console**:

- **Public**: email addresses are displayed publicly.
- **Masked**: email addresses are still displayed publicly, but masked in such a way to make it harder for spam-bots to harvest them.
- **Only visible to site administrators**: only Confluence administrators can see the email addresses. Note that, if you select this option, email addresses will not be available in the 'User Search' popup (e.g. when setting Page Restrictions).

To configure user email visibility:

1. Choose the**cog icon** , then choose**General Configuration**
2. Choose '**Security Configuration**'.
3. Choose '**Edit**'. The fields on the '**Security Configuration**' screen will be editable.
4. Select one of the options from the '**User email visibility**' dropdown: '**public**', '**masked**', or '**only visible to site administrators**'.
5. Choose '**Save**'.

Screenshot: Email Visibility

Anonymous Access to Remote API

Administrators may wish to disable anonymous access to the Confluence remote API. to make it harder for malicious users to write 'bots' that perform bulk changes to the site.

To disable anonymous access to the remote API:

1. Choose the**cog icon** , then choose**General Configuration**
2. Choose **Security Configuration** in the left-hand panel. The **Security Configuration** screen will appear.
3. Choose **Edit**.
4. Uncheck the **Anonymous Access to API** check box.
5. Choose **Save**.

Notes

This page is about access to the remote API. If you are looking for information about preventing anonymous users from accessing Confluence, see Global Permissions Overview.

Configuring RSS Feeds

A Confluence System Administrator can configure the following aspects of RSS feeds:

- The maximum number of items that Confluence returns to an RSS feed request.
- The maximum time period that Confluence allows to respond to an RSS feed request.

Both of these are set in the 'Edit Security Configuration' screen.

To configure RSS feeds:

1. Choose the**cog icon** , then choose**General Configuration**
2. Choose**Security Configuration**.
3. Choose**Edit**.
4. Enter a value for**Maximum RSS Items**. The default value is 200.
5. Enter a value for**RSS timeout**.
6. Choose**Save**.

Screenshot: Configuring RSS feeds

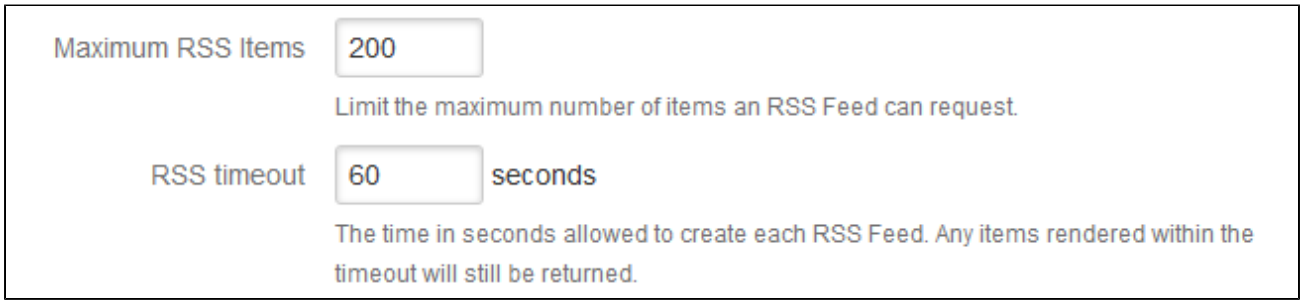

On this page:

- Notes

Related pages:

- The RSS Feed Builder

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:

```
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.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

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.

Related pages:

- Getting Started as Confluence Administrator
- Supported Platforms

Confluence is a J2EE web application. On the client side, users access Confluence primarily via a web browser.

This section contains the following guidelines:

- Confluence Home and other important directories
- Application Server Configuration
- Starting Confluence Automatically on System Startup

Diagram: A Confluence installation

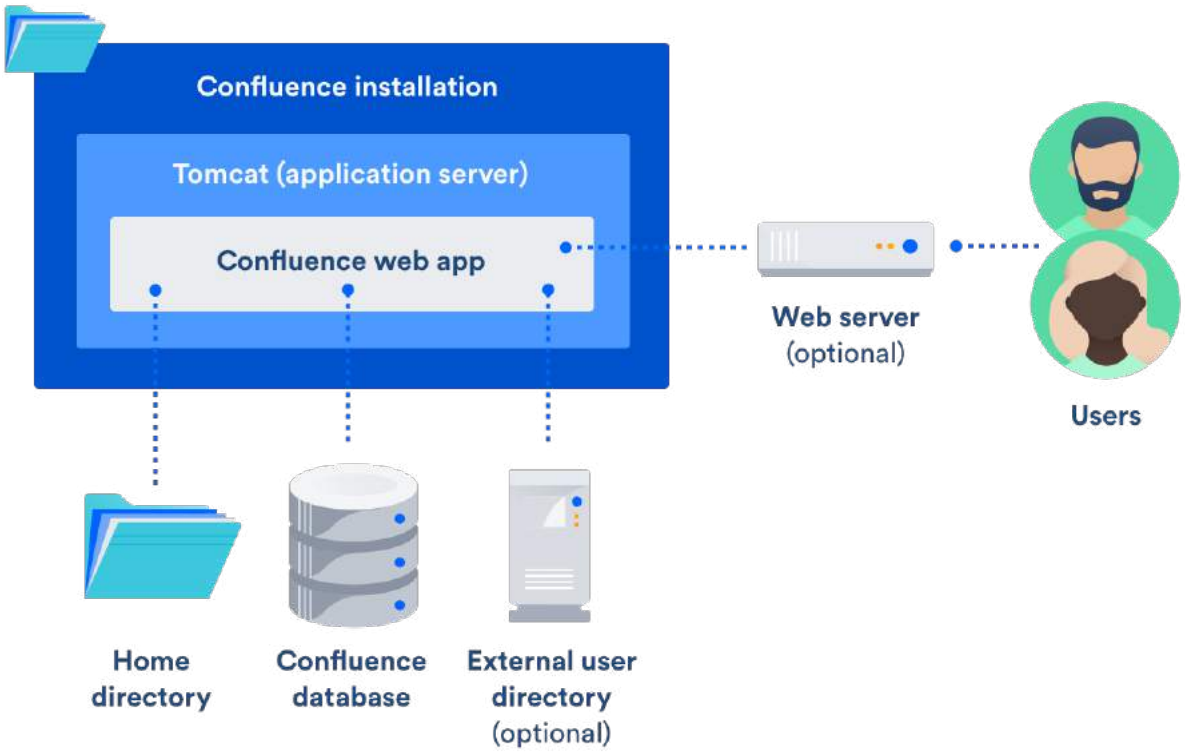

Confluence Home and other important directories

Confluence installation directory

The 'Confluence Installation directory' is the directory where Confluence was installed. This directory is also sometimes called the 'Confluence Install directory'.

Important files in the installation directory:

- `bin/setenv.batorbin/setenv.sh`
  This file is used to edit CATALINA_OPTS memory and garbage collection settings and define system properties.
- `confluence/WEB-INF/classes/confluence-init.properties`
  This file contains the location of the Confluence Home directory.

On this page:

- Confluence installation directory
- Confluence home directory
- Changing the location of the home directory
- Database
- Temp directory (installation directory)

Confluence home directory

The Confluence Home directory is the folder where Confluence stores its configuration information, search indexes and page attachments. Another term for 'Home directory' would be 'data directory'. We also refer to this as the 'local home directory' in Data Center.

Finding the home directory

The location of the Confluence home directory is defined when youinstall Confluence. This location is stored in the`confluence-init.properties`file, which is located in the`confluence/WEB-INF/classes`directory of your Confluence Installation directory.

When Confluence is running you can find the location of the home directory in > **General Configuration**>**SystemInformation**>**Confluence Information - Confluence Home**.

If you're using Confluence Data Center (a clustered instance), you will also have a**shared home**directory which will contain some data (such as attachments and backups) that would otherwise reside in the home directory. The location of your shared home directory can be found in your`<local-home>/confluence.cfg.xml`filein the`confluence.cluster.home`property.

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`. In Data Center this directory is usually found in the Shared Home directory. This directory won't be used if S3 object storage is configured.
`backups/`	Confluence will place its daily backup archives in this directory. 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`. In Data Center this directory is usually found in the Shared Home directory.
`bundled-plugins/`	Confluence includes a set of*bundled*plugins. The`bundled-plugins`directory is where Confluence will unpack its bundled plugins when it starts up. This directory is refreshed on every restart, so removing a plugin from this directory will not uninstall the plugin, as it will be replaced the next time Confluence starts up.
`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-cached`irectory. This directory is updated as plugins are installed and uninstalled from the system and is completely repopulated from the database every time Confluence is restarted. Removing plugins from this directory does not uninstall them.
`temp/`	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. You can specify a different temp directory location, if necessary. Edit `<confluence-home>/confluence.cfg.xml` and set the new location in the `webwork.multipart.saveDir` property. Your new location can't be in the installation directory, as this will cause some functions, such as download, to fail. We recommend you keep the `temp` directory in the local home directory.
`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. In Data Center this directory is usually found in the Shared Home directory.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

`shared-home/`	This sub-directory is created in your home directory when you install Confluence Server. If you choose move to Data Center, you'll move the contents of this directory to a separate shared home directory, that is accessible to all nodes. Cache files for some features, including Office document and PDF previews are also stored in this directory (in both Server and Data Center).
`restore/site/`	Confluence will save site backups generated in the administration console in this directory. See Back up a Site for more information. If you Restore a Site in the admin console, you can select to import a file from this location. In Data Center this directory is usually found in the Shared Home directory.
`restore/space/`	Confluence will save single space and multi-space backups generated from the administration console in this directory. See Back up a space or multiple spaces for more information. Confluence will also save XML backups done via Space tools in this directory. See Export Content to Word, PDF, HTML and XML for more information. If you Restore a Spacein the admin console, you can select to import a file from this location. In Data Center this directory is usually found in the Shared Home directory.

Changing the location of the home directory

When Confluence first starts up, it reads the`confluence-init.properties`file to determine where to look for the Home directory.

To change the location of the home directory edit the`confluence.home`property in the`confluence-init.properties`file as follows:

- **Windows**
  In Windows, the path`C:\confluence\data`would be written as:
  `confluence.home=C:/confluence/data`
  Note that all backslashes (\) are written as forward slashes (/)
- **Linux**
  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

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

All other data, including page content, is kept in the database. If you installed Confluence as a trial, or chose to use theEmbedded H2 Databaseduring setup, the database will store its files under`database/`in theConfluence Home Directory. Otherwise, the database management system you are connecting to is responsible for where and how your remaining data is stored.

Temp directory (installation directory)

The temp directory is configured in the Java runtime and some Confluence components write temporary files or lockfiles into this directory.

The tempdirectoryis located in theinstallationdirectory as /temp.

To change the location of this directory, start the Java Virtual Machine in which confluence is running with the argument:

`-Djava.io.tmpdir=/path/to/your/own/temp/directory`.

⚠Note: this is not the same as the temp directory in Confluence Home where exports, for example, are saved. See the table above to find out how to change the location of the`<confluence-home>/temp` directory.

There's a known issue with setting a temp directory in Confluence. See

CONFSERVER-59613 - java.io.tmpdir has no effect on changing the installation temp directory
GATHERING IMPACT

.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Application Server Configuration

The following pages contain information about configuring your application server for Confluence:

- Managing Application Server Memory Settings

Managing Application Server Memory Settings

The minimum and maximum JVM heap space allocated to the application server affects performance. Confluence administrators may wish to modify this value from the defaults depending on their server load. This document only provides guidelines rather than rules, so administrators optimizing for performance should use this document as a starting point only.

> ⚠ 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 maximized under peak load, and when creating a site XML backup. In many cases, the backup can be the cause of the OOM, so increase -Xmx values and verify if a backup was occurring at the time of OOM. A quick rule of thumb for gauging the success of a memory adjustment is using simple anecdotal evidence from users. Is it snappier? The same? How does it handle while a backup is occurring?

> ⚠ Atlassian recommends in normal use, to disable the XML backup and use a Production Backup Strategy.

- If you normally perform manual XML site backups on your server, test your maximum memory requirements by performing a site XML backup while the server is under maximum load
- If you do not create manual XML site backups, simply monitor the server while under maximum load

Applying Memory Settings

See How to fix out of memory errors by increasing available memory.

Related Topics

- Garbage Collector Performance Issues
- How to fix out of memory errors by increasing available memory
- Server Hardware Requirements Guide
- Performance Tuning
- Troubleshooting Slow Performance Using Page Request Profiling
- Tomcat JVM options and Modify the Default JVM Settings

Starting Confluence Automatically on System Startup

You can configure Confluence to start automatically on system startup, allowing it to recover automatically after a reboot.

- Start Confluence Automatically on Linux
- Start Confluence Automatically on Windows as a Service

Start Confluence Automatically on Linux

On Linux/Solaris, the best practice is to install, configure and run each service (including Confluence) as a dedicated user with only the permissions they require.

To install, configure and run Confluence automatically on Linux/Solaris:

1. Create a `confluence` user for instance, using the following command:

   ```
   sudo useradd --create-home -c "Confluence role account" confluence
   ```

2. Create a directory to install Confluence into.In this example we're using`/usr/local/confluence`.

   ```
   sudo mkdir /usr/local/confluence
   sudo chown confluence: /usr/local/confluence
   ```

3. Log in as the `confluence` user to install Confluence:

   ```
   sudo su - confluence
   cd /usr/local/confluence/
   tar zxvf /tmp/confluence-5.6.4.tar.gz
   ln -s confluence-5.6.4/ current
   ```

4. Edit <<CONFLUENCE_INSTALL_DIRECTORY>>/confluence/WEB-INF/classes/confluence-init.properties file, and set confluence.home=/usr/local/confluence/<Confluence_Data_Home> (ensure you have removed the comment '#')
5. Then back as root, create the file `/etc/init.d/confluence` (code shown below), which will be responsible for starting up Confluence after a reboot (or when manually invoked).
   ⚠️ If you are running Ubuntu Jaunty (or later) do not perform this step. Please use the instructions further down this page.

```
#!/bin/sh -e
# Confluence startup script
#chkconfig: 2345 80 05
#description: Confluence

# Define some variables
# Name of app ( JIRA, Confluence, etc )
APP=confluence
# Name of the user to run as
USER=confluence
# Location of Confluence install directory
CATALINA_HOME=/usr/local/confluence/current
# Location of Java JDK
export JAVA_HOME=/usr/lib/jvm/java-7-oracle

case "$1" in
  # Start command
  start)
    echo "Starting $APP"
    /bin/su -m $USER -c "$CATALINA_HOME/bin/start-confluence.sh &> /dev/null"
    ;;
  # Stop command
  stop)
    echo "Stopping $APP"
    /bin/su -m $USER -c "$CATALINA_HOME/bin/stop-confluence.sh &> /dev/null"
    echo "$APP stopped successfully"
    ;;
   # Restart command
   restart)
        $0 stop
        sleep 5
        $0 start
        ;;
  *)
    echo "Usage: /etc/init.d/$APP {start|restart|stop}"
    exit 1
    ;;
esac

exit 0
```

6. Make this file executable:

```
sudo chmod +x /etc/init.d/confluence
```

7. Set this file to run at the appropriate runlevel. For example, use `sudo chkconfig --add confluence` on Redhat-based systems, `sudo update-rc.d confluence defaults` or `rcconf` on 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)

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

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`:

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

```
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

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Start Confluence Automatically on Windows as a Service

For long-term use, we recommend that you configure Confluence to start automatically when the operating system restarts. For Windows servers, this means configuring Confluence to run as a Windows service.

There are two ways to install the Confluence as a service: using theConfluence installeror manually as described below.

On this page:

- Reasons for starting Confluence as a service
- Changing the user running the service
- Manually installing Confluence as a service
- Managing Confluence as a service
- Upgrading Confluence
- Troubleshooting Confluence while running as a Windows service
- Requesting Support

Problem with 64-bit Windows

If you are running 64-bit Windows, please note that you may encounter problems with Apache Tomcat running as a Windows service if you are using a 64-bit JDK. Refer to our knowledge base article for more information.

Reasons for starting Confluence as a service

Installation as a Windows service offers these advantages:

- Reduced risk of shutting down Confluence by accident (If you start Confluence manually, a console window opens and there is a risk of someone accidentally shutting down Confluence by closing the window).
- Automated Confluence recovery after server restart.
- Improved troubleshooting through logging server output to file.

You can read more about Windows services in the Microsoft Developer Network.

Changing the user running the service

If you wish to run the service as a non-administrator user for security, or if you are using network drives for backups, attachments, or indexes, you can run the service as another user. To change users, open the Apache Tomcat Confluence properties, go to the 'Log On' tab, and enter the required username and password. Go to your Windows Control Panel -> User Accounts and confirm that the user has write permissions for the`<CONFLUENCE-INSTALL>`and `<CONFLUENCE-HOME>`directories and all subfolders. Note that any network drives must be specified by UNC and not letter mappings (eg. `\\backupserver\confluence` not `z:\confluence`).

For more detail, see Creating a Dedicated User Account on the Operating System to Run Confluence.

Manually installing Confluence as a service

In Windows:

1. Open a command prompt and change directory to the `<CONFLUENCE-INSTALL>/bin` directory. You'll need to run the command prompt using 'Run as administrator' so that you can complete some of these steps.
2. Confirm that the JAVA_HOME variable is set to the JDK base directory with the command:

```
echo %JAVA_HOME%
```

If you installed the Java Runtime Environment (JRE) or used the Confluence installer, replace `JAVA_HOME` with `JRE_HOME`. See Setting the JAVA_HOME Variable in Windows for more info.

Note that any directory in the path with spaces (eg. `C:\Program Files` must be converted to its eight-character equivalent (e.g. `C:\Progra~1`).

3. Use the following command to install the service with default settings:

```
service.bat install Confluence
```

The service will be called **Atlassian Confluence** and will be configured to start automatically by default, but will not automatically start up until the next server reboot.

4. If you have a large Confluence installation, you can increase the maximum memory Confluence can use (the default is 1024MB). For example, you can set the maximum memory to 2048MB using:

```
tomcat9w //US//Confluence --JvmMx 2048
```

5. If you don't have any JVM parameters that you pass to Confluence, you can skip this step. If you do, add them to the service using:

```
tomcat9w //US//Confluence ++JvmOptions="-Djust.an.example=True"
```

Alternatively you can use the following command to launch the service properties dialog then navigate to the Java tab to add more JVM parameters.

```
tomcat9w //ES//Confluence
```

For further configuration options, please refer to the Tomcat Windows Service How-To guide.

6. Go to **Control Panel** > **Administrative Tools** > **Services** > **Atlassian Confluence** and right-click **Properties** to verify the settings are correct. Start the Confluence service with the command:
7. Finally, start the Confluence service. From now on this will happen automatically after the a server reboot.

```
net start Confluence
```

Managing Confluence as a service

You can manage the Confluence service from the command prompt.

- Stop Confluence with:

```
net stop Confluence
```

- Uninstall the Confluence service with:

```
service.bat remove Confluence
```

Upgrading Confluence

After upgrading Confluence, you can either uninstall and reinstall the Windows service or change the `StartPath` parameter to your new folder. Refer to the Tomcat documentation for help.

Troubleshooting Confluence while running as a Windows service

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

- Check the Knowledge Base articles:
  - Getting 'The image file tomcat6.exe is valid, but is for a machine type other than the current machine'
  - Confluence Does Not Start Due to Windows Firewall
  - Unable to start Confluence Windows service after allocating JVM memory
  - Unable to Configure Confluence to Run as a Service on Tomcat 5
  - Unable to Install Service on Windows Vista
- If none of the above solves your problem, please refer to the complete list of known issues in our Knowledge Base.
- When investigating memory issues or bugs, it may be useful to view information from Confluence's garbage collection. To turn on the verbose garbage collection seeHow to Enable Garbage Collection (GC) Logging.
- You can use a Sysinternals tool called **`Procmon.exe`** from the The Microsoft Windows Sysinternals Team, to check that the error occurred at the specific time when the Confluence service started. You need to match the time when Tomcat failed, as captured by this tool, against the time in the Windows Event Viewer.

  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:

- Give us the result of running **`java -version`** from Windows command line console.
- A screen shot of your Windows Registry setting for Tomcat.
- If you have modified **service.bat**, please give us a copy of this file for review.
- A support zip, containing your application logs and configuration files.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Performance Tuning

This document describes tuning your application for improved performance. It is not a guide to troubleshooting Confluence outages. Check Trouble shooting Confluence hanging or crashing for help if Confluence is crashing.

Like any server application, Confluence may require some tuning as it is put under heavier use. We do our best to make sure Confluence performs well under a wide variety of circumstances, but there's no single configuration that is best for everyone's environment and usage patterns.

If you are having problems with the performance of Confluence and need our help resolving them, you should read Requesting Performance Support.

On this page:

- Performance Data Collector
- Use the latest version of your tools
- Avoid swapping due to not enough RAM
- Being aware of other systems using the same infrastructure
- Choice of database
- Database connection pool
- Database in general
- Database statistics and query analyzers
- Cache tuning in Confluence and Apache
- Antivirus software
- Enabling HTTP compression
- Performance testing
- Access logs
- Built-in profiler
- Application server memory settings
- Web server configuration
- Troubleshooting possible memory leaks
- Enable faster permissions

Performance Data Collector

ThePerformance Data Collectoris a server-side, standalone application that exposes a number of REST APIs for collecting performance data. It can be used to collect data, such as thread dumps, disk speed and CPU usage information, to troubleshoot performance problems.

SeeHow to use the Performance Data Collectorfor more information.

Use the latest version of your tools

Use the latest versions of your application servers and Java runtime environments. Newer versions are usually better optimized for performance.

Avoid swapping due to not enough RAM

Always watch the swapping activity of your server. If there is not enough RAM available, your server may start swapping out some of Confluence's heap data to your hard disk. This will slow down the JVM's garbage collection considerably and affect Confluence's performance. In clustered installations, swapping can lead to aCluster Panic due to Performance Problems. This is because swapping causes the JVM to pause duringGarbage 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 have configured Confluence to access the database directly, you will need to manually edit the hibernate.c3p0.max_size property and hibernate.hikari.maximumPoolSize property (if present)in the confluence.cfg.xml file in your confluence.home directory. After you have changed the URL in this file, restart Confluence.

To assess whether you need to tune your database connection pool, take thread dumps during different times (including peak usage). Inspect how many threads have concurrent database connections.

Database in general

If Confluence is running slowly, one of the most likely cause is that there is some kind of bottleneck in (or around) the database.

The first item you should check is the **"Database Latency"** field in the System Information tab in the admin console.

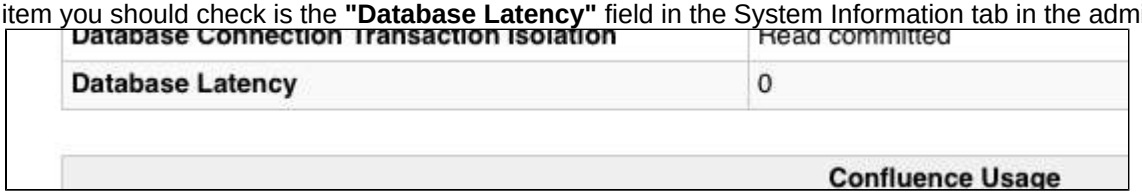

The latency is calculated by sending a trivial request to the database, querying a table which is known to have only one column and one row. ("select * from CLUSTERSAFETY"). Obviously this query should be blazing fast, and return within 1 or 2 milliseconds. If the value displayed is between 3 and 5 milliseconds, you might already have an issue. If the value is above 10ms, then you **definitely** need to investigate and improve something! A few milliseconds may not sound so bad, but consider that Confluence sends quite a few database queries per page request, and those queries are a lot more complex too! High latency might stem from all sorts of problems (slow network, slow database, connection-pool contention, etc), so it's up to you to investigate. Don't stop improving until latency is below 2ms on average.

Obviously, latency is just the very first thing to look at. You may get zero latency and still have massive database problems, e.g. if your tables are poorly indexed. **So don't let a low latency fool you either.**

Database statistics and query analyzers

Modern databases have query optimizers based on collecting statistics on the current data. Using the SQL EXPLAIN statement will provide you information on how well the query optimizer is performing. If the cost estimate is wildly inaccurate then you will need to run statistics collection on the database. The exact command will depend on your database and version. In most cases you can run statistics collection while Confluence is running, but due to the increased load on the database it's best to do this after normal hours or on a week-end.

Cache tuning in Confluence and Apache

To reduce the load on the database, and speed up many operations, Confluence keeps its own cache of data. Tuning the size of this cache may speed up Confluence (if the caches are too small), or reduce memory (if the caches are too big).

Please have a look at our documentation on Cache Performance Tuning for information on how to tune Confluence caches.

Antivirus software

Antivirus software greatly decreases the performance of Confluence. Antivirus software that intercepts access to the hard disk is particularly detrimental, and may even cause errors with Confluence. You should configure your antivirus software to ignore the Confluence home directory, its index directory and any database-related directories.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Enabling HTTP compression

If bandwidth is responsible for bottlenecking in your Confluence installation, you should consider enabling HTTP compression. This may also be useful when running an external facing instance to reduce your bandwidth costs.

Take note of the known issues with HTTP compression in versions of Confluence prior to 2.8, which may result in high memory consumption.

Performance testing

You should try out all configuration changes on a demo system. Ideally, you should run and customize loadtests that simulate user behavior.

Access logs

You can find out which pages are slow and which users are accessing them by enabling Confluence's built-in access logging.

Built-in profiler

You can identify the cause of page delays using Confluence's built-in profiler according to Troubleshooting Slow Performance Using Page Request Profiling.

Application server memory settings

See How to fix out of memory errors by increasing available memory.

Web server configuration

For high-load environments, performance can be improved by using a web server such as Apache in front of the application server. There is a configuration guide to Running Confluence behind Apache.

When configuring your new web server, make sure you configure sufficient threads/processes to handle the load. This applies to both the web server and the application server connector, which are typically configured separately. If possible, you should enable connection pooling in your web server connections to the application server.

Troubleshooting possible memory leaks

Some external plugins, usually ones that have been written a long time ago and that are not actively maintained anymore, have been reported to consume memory and never return it. Ultimately this can lead to a crash, but first this manifests as reduced performance. The Troubleshooting Confluence hanging or crashing guide is a good place to start. Some of the known causes listed there could result in performance issues short of a crash or hang.

Enable faster permissions

Confluence needs to regularly check the current user's permissions in order to determine what to display. The faster permissions service changes the way permissions information is stored in the database to optimize these permissions checks. Although this comes with some overhead, it can provide a significant performance improvement in sites with a lot of content and complex permissions. If you only have a small amount of content or very simple permissions (just a few groups, or few nested page restrictions), this service is unlikely to make your Confluence site significantly faster.

Learn more about the faster permissions service

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Cache Performance Tuning

Confluence performance can be significantly affected by the performance of its caches.

Before you change the size of your caches, it's important to take a baseline so you can measure how effective each individual change is, and decide whether they are needed.

On this page we'll take you through some example statistics and discuss how you might be able to improve Confluence performance by resizing these caches.

If you just want to check your cache statistics, or make a change to your cache config, seeCache Statistics.

On this page:

- Cache tuning example
- Important caches

Cache tuning example

As an example of how to tune Confluence's caches, let's have a look at the following table:

Caches	% Utilization	% Effectiveness	Objects/Size	Hit/Miss/Expiry
Attachments	87%	29%	874/1000	78226/189715/187530
Content Attachments	29%	9%	292/1000	4289/41012/20569
Content Bodies	98%	81%	987/1000	28717/6671/5522
Content Label Mappings	29%	20%	294/1000	4693/18185/9150
Database Queries	96%	54%	968/1000	105949/86889/83334
Object Properties	27%	18%	279/1000	5746/25386/8102
Page Comments	26%	11%	261/1000	2304/17178/8606
Users	98%	5%	982/1000	6561/115330/114279

The maximum size of the caches above is 1000 (meaning that it can contain up to 1000 objects).You can tell when a cache size needs to be increased because the cache has both:

- a high usage percentage (above 75%)
- a low effectiveness percentage.

Check the 'effectiveness' versus the 'percent used'. A cache with a low percent used need not have its size lowered; it does not use more memory until the cache is filled.

Based on this, the sizes of the "Attachments", "Database Queries", and "Users" caches should be increased to improve their effectiveness.

As the stored information gets older or unused it will expire and be evicted from the cache. Cache expiry can be based on time or on frequency of use.

There's not much that you can do with a cache that has both a low percentage of usage and effectiveness. Over time, as the cache is populated with more objects and repeat requests for them are made, the cache's effectiveness will increase.

Important caches

The following suggestions are general guidelines. In cases of large databases, 20-30% of the size of the table may be unnecessarily large. Check the effectiveness and percent used categories in the cache for more specific assessments.

- **Content Objects cache** (`com.atlassian.confluence.core.ContentEntityObject`) should be set to at least 20-30% of the number of content entity objects (pages, comments, emails, news items) in your system. To find the number of content entity objects, use the 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`.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Cache Statistics

Caches help reduce the load on your database, and can make some operations faster.Track the size and hit ratio of each of Confluence's internal caches, and adjust the cache size for better performance.

On this page:

- View cache statistics
  - View cache statistics in a cluster
  - What the statistics mean
  - Cache types
- Change the size of a cache

View cache statistics

To view cache statistics:

1. Go to > **General Configuration**>**Cache Management**.
2. Select**Show Advanced View** to see the full details.

Screenshot: Cache statistics screen showing the utilisation and effectiveness of a selection of caches.

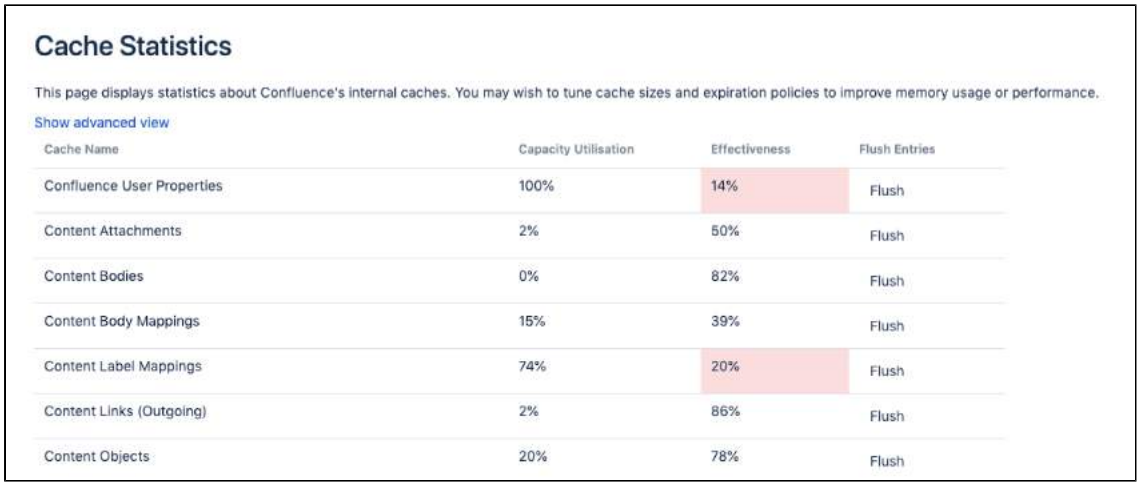

View cache statistics in a cluster

If you're running Confluence in a cluster, this screen shows the statistics for the node you're currently on.

To view cache statistics for another node in the cluster:

1. Go to > **General Configuration**>**Clustering**.
2. Select •••>**Cache statistics**next to the node you want to view.

You will only be able to view the statistics. To flush a cache or adjust the size, you'll need to access the Cache Management screen on each node directly.

What the statistics mean

Here's some information on how eachnumber is generated.

Capacity Utilization	=(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.

Cache types

When running in a cluster, Confluence has three types of caches:

- **local** - cache data is replicated on each node.
- **distributed** -cache data isevenly partitioned across all Confluence nodes in the cluster (known as replicate-via-copy).
- **hybrid** - cache data is replicated on each node, and invalidated remotely by other nodes when things change (known as replicate-via-invalidation).

The cache type is indicated with a lozenge beside the cache name in the advanced view.

Screenshot: Cache statistics advanced view showing the full details of each cache, including the cache type.

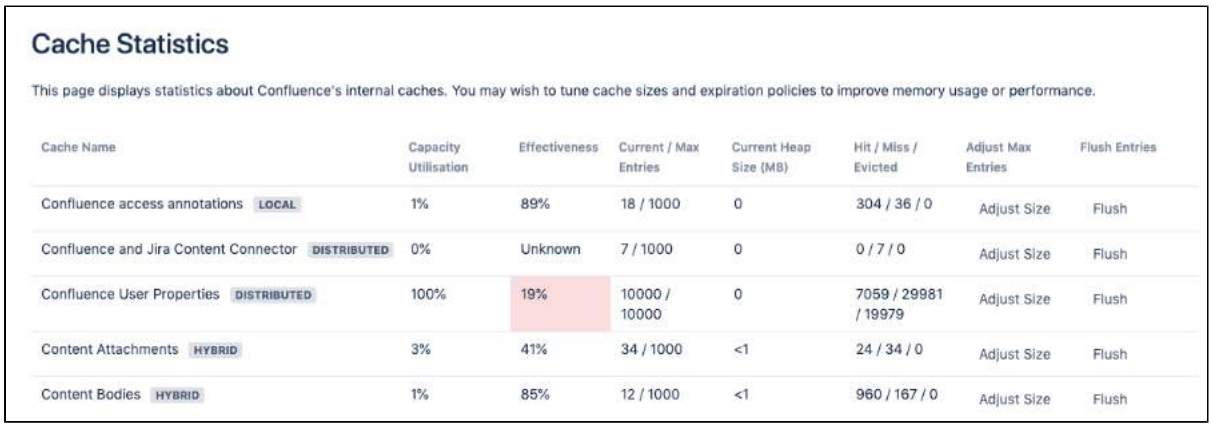

Change the size of a cache

Tuning the size of a cache can speed up Confluence (if the caches are too small), or reduce memory (if the caches are too large).Larger caches will require more memory at runtime, so make sure you review the memory allocation of the Confluence Java process and the physical memory available on your server.

You need **System Administrator** global permission to change the size of a cache.

To change the size of a cache:

1. Go to > **General Configuration**>**Cache Management**.
2. Select**Show Advanced View**.
3. Select **Adjust Size** next to the cache you want to change.
4. Enter the maximum number of items to be stored in the cache and select **Submit**.

The changes will take effect immediately. You don't need to restart Confluence.

Any changes to cache sizes are recorded in:

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

- `<home-directory>/shared-home/config/cache-settings-overrides.properties` if you run Confluence on a single server.
- `<shared-home>/config/cache-settings-overrides.properties`if you run Confluence in a cluster.

To reset the values back to the default, you can delete the`cache-settings-overrides.properties`file and restart Confluence.

SeePerformance Tuning for a more general overview of tuning in Confluence.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

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 customize Confluence's cache to meet your needs, seecache tuning.
To increase the amount of memory available to Confluence, seeHow to fix out of memory errors by increasing available memory.

Mail error queue

Confluence keeps a copy of all emails that it failed to send within an internal error queue. In the event of intermittent failures such as network connectivity issues, the emails in this queue can be manually resent when the problem is fixed. Under certain circumstances, the mail queue can fill up with large objects. The queue is regularly flushed, but if you get a *lot* of mail errors, you might get a spike in memory usage.

Attachments

The indexing of large attachments requires that the attachment be loaded into memory. In the case of large attachments, this can cause a temporary strain on the systems resources, and may result in indexing failing because the attachment could not be fully loaded into memory.

System backup and restore

The Confluence backup and restore process scales linearly with the size of data. This can have a significant impact on large Confluence instances where the amount of data exceeds the amount of available memory. If you are experiencing an `OutOfMemoryError` during either a backup or restore processes, then we strongly recommend that you choose and Production Backup Strategy.

If you encounter an `OutOfMemoryError` while restoring a backup and wish to overcome this issue by increasing memory, how much more will you need to make this process work? A good rule of thumb is to have a look at the size of the `entities.xml` file in your backup. This file contains all of the data Confluence will be loading, so at least that much is required. Add another 64-128Mb to ensure that Confluence has enough memory to load and function and that should be enough. To increase the amount of memory available to Confluence, see How to fix out of memory errors by increasing available memory.

Known issues that we do not have control over

There are also some memory issues we don't have any control over. For example,

- There's a memory leak in the Oracle 10g JDBC drivers. Not much we can do about that.
- One customer found a rather nasty memory leak that appeared to originate inside Tomcat 5, but only using the IBM JDK on PowerPC.

If you are having problems that appear to result from a memory leak, log an issue on http://support.atlassian.com. Our memory profiler of choice is YourKit. It would be helpful to us if you can provide us with a memory dump from that tool showing the leak.

Confluence is taking long periods of time to respond to some actions

A common cause of random pauses in Confluence is the JVM running garbage collection. To determine if this is what is happening, enable verbose garbage collection and look at how long Java is taking to free up memory. If the random pauses match when Java is running its garbage collection, garbage collection is the cause of the pause.

Verbose garbage collection will generate log statements that indicate when Java is collecting garbage, how long it takes, and how much memory has been freed.

To enable gc (garbage collection) logging, start Confluence with the option `-XX:+PrintGCDetails -XX:+PrintGCTimeStamps -verbose:gc -Xloggc:gc.log`. Replace `gc.log` with an absolute path to a `gc.log` file.

For example, with a Windows service, run:

```
tomcat5 //US//Confluence ++JvmOptions="-XX:+PrintGCDetails -XX:+PrintGCTimeStamps -verbose:gc -Xloggc:c:\confluence\logs\gc.log"
```

or in `bin/setenv.sh`, set:

```
export CATALINA_OPTS="$CATALINA_OPTS -XX:+PrintGCDetails -XX:+PrintGCTimeStamps -verbose:gc -Xloggc:${CATALINA_BASE}/logs/gc.log"
```

If you modify `bin/setenv.sh`, you will need to restart Confluence for the changes to take effect.

What can you do to minimize the time taken to handle the garbage collection? See http://java.sun.com/docs/hotspot/gc1.4.2/ for details on tuning the JVM to minimize the impact that garbage collection has on the running application.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Requesting Performance Support

Basic performance troubleshooting steps

Begin with the following procedures:

1. Go through the Troubleshooting Confluence hanging or crashing page to identify the major known performance problems.
2. Proceed with the Performance Tuning tips to help optimize performance.

On this page:

- Basic performance troubleshooting steps
- Requesting basic performance support
- Advanced performance troubleshooting

Related pages:

- Memory Usage and Requirements
- Confluence for Enterprise

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. Asupport zip, containing log files and configuration, ideally with a series ofthread dumpsseparated by 10 seconds.
2. A description with as much detail as possible regarding:
    a. What changes have been made to the system?
    b. When did performance problems begin?
    c. When in the day do performance issues occur?
    d. What pages or operations experience performance issues?
    e. Is there a pattern?

Continue with as much of the advanced performance troubleshooting information as you can.

Advanced performance troubleshooting

Please gather **all** of the information listed below and include it in your support request, even if you think you have a good idea what's causing the problem. That way we don't have to ask for it later.

System information

Confluence server

- Take a screenshot of Confluence's `Administration  System Information` (or save the page as HTML)
- Take a screenshot of Confluence's `Administration  Cache Statistics` (or save the page as HTML)
- Find out the exact hardware Confluence is running on
    - How many CPUs? What make and model? What MHz?
    - How much memory is installed on the machine?
    - How much memory is assigned to Confluence's JVM? (i.e. what are the -Xmx and -Xms settings for the JVM?)
    - What other applications are being hosted on the same box?

Confluence content

- How many *users* are registered in Confluence?
- On average, to how many *groups* does each user belong?
- How many *spaces* (global and personal) are there in your Confluence server?
- How many of those spaces would be viewable by the average user?
- Approximately how many *pages*? (Connect to your database and perform '`select count(*) from content where prevver is null and contenttype = 'PAGE'`')
- How much data is being stored in Bandana (where plugins usually store data)? (Connect to your database and perform '`select count(*), sum(length(bandanavalue)) from bandana`')

The database

- What is the exact version number of Confluence's database server?

- What is the exact version number of the JDBC drivers being used to access it? (For some databases, the full filename of the driver JAR file will suffice)
- Is the database being hosted on the same server as Confluence?
- If it is on a different server, what is the network latency between Confluence and the database?
- What are the database connection details? How big is the connection pool? If you are using the standard configuration this information will be in your confluence_cfg.xml file. Collect this file. If you are using a Data source this information will be stored in your application server's configuration file, collect this data.

User management

- Are you using external user management or authentication? (i.e. Jira or LDAP user delegation, or single sign-on)
- If you are using external Jira user management, what is the latency between Confluence and Jira's database server?
- If you are using LDAP user management:
  - What version of which LDAP server are you using?
  - What is the latency between Confluence and the LDAP server?

Diagnostics

Observed problems

- Which pages are slow to load?
  - If it is a specific wiki page, attach the wiki source-code for that page
- Are they always slow to load, or is the slowness intermittent?

Monitoring data

Before drilling down into individual problems, helps a lot to understand the nature of the performance problem. Do we deal with sudden spikes of load, or is it a slowly growing load, or maybe a load that follows a certain pattern (daily, weekly, maybe even monthly) that only on certain occasions exceeds critical thresholds? It helps a lot to have access to continuous monitoring data available to get a rough overview.

Here are sample graphs from the confluence.atlassian.com system, showing

Load
This graph shows the load for two consecutive days. The obvious pattern is that the machine is under decent load, which corresponds to the user activity, and there is no major problem.

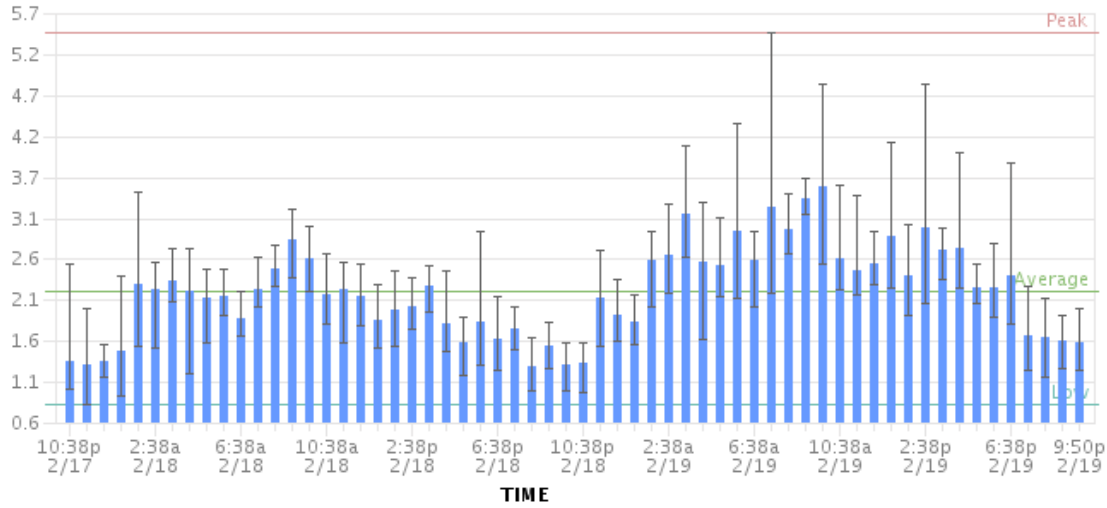

Resin threads and database connections

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

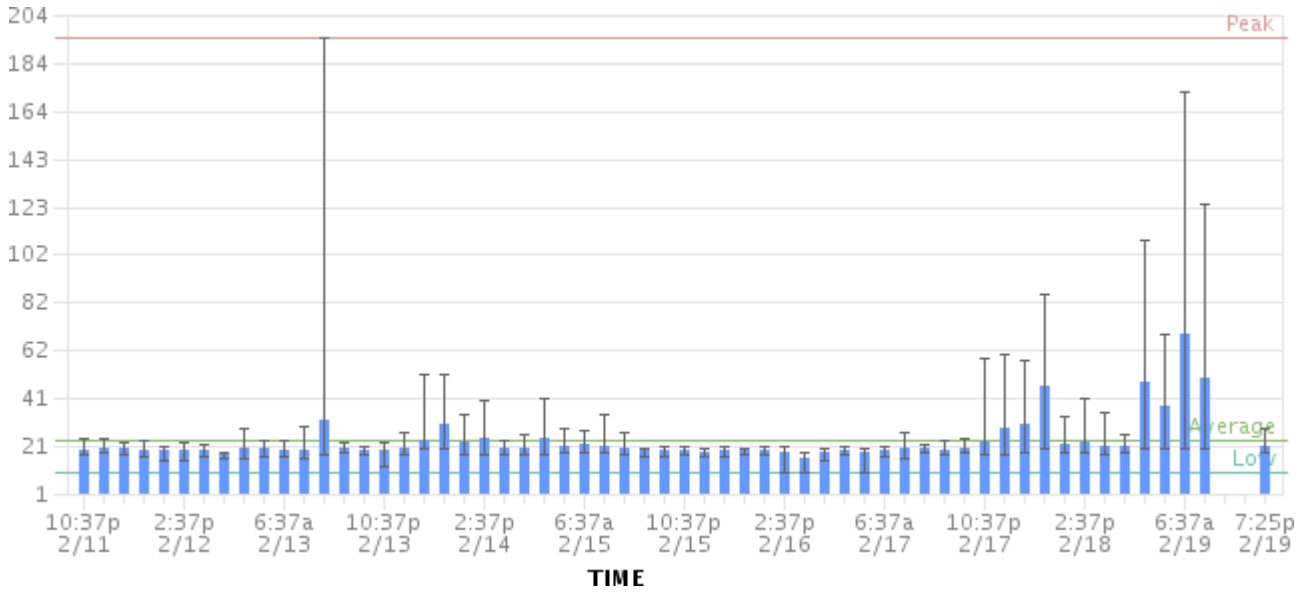

Active number of Java Threads

These two charts show the active threads in the application server (first chart) and the size database connection pool (second chart). As you can see, there was a sudden spike of server threads and a corresponding spike of db-connections.

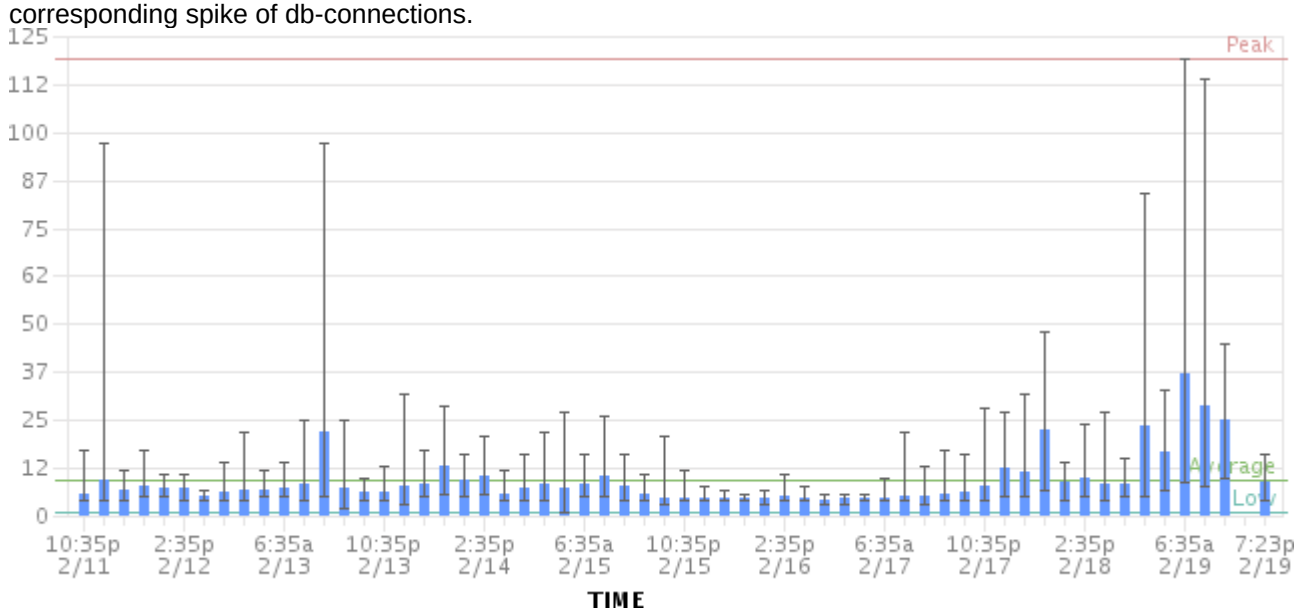

The database connection pool size

The database connection pool size peaked over 112, which happened to be more than the maximum number of connections the database was configured for (100). So it was no surprise that some requests to Confluence failed and many users thought it had crashed, since many requests could not obtain the crucial database connections.

We were able to identify this configuration problem quite easily just by looking at those charts. The next spikes were uncritical because more database connections were enabled.

The bottom line being: it helps a lot to monitor your Confluence systems continuously (we use Hyperic, for example), and it helps even more if you are able to send us graphs when you encounter problems.

Access logs

- Access logging is enabled by default. You can find Confluence access logs at`<confluence install>/logs/conf_access_log.<date>.log`. These logs are configured in the server.xml. Refer to theTomcat Access Log Valve documentation for more information on the attributes that can be logged.
- You can also chose to enable access logging in your application logs via the log4j configuration. SeeHow to Enable User Access Logging, including how to redirect the logs to a separate file

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

You can run these access log files through a log file analyzer such as AWStats, or manually look for pages which are slow to load.

Profiling and logs

- Enable Confluence's built-in profiling for long enough to demonstrate the performance problem using Troubleshooting Slow Performance Using Page Request Profiling.
  - If a single page is reliably slow, you should make several requests to that page
  - If the performance problem is intermittent, or is just a general slowness, leave profiling enabled for thirty minutes to an hour to get a good sample of profiling times
- Find Confluence's standard output logs (which will include the profiling data above). Take a zip of the entire logs directory.
- Take a thread dump during times of poor performance

CPU load

- If you are experiencing high CPU load, please install the YourKit profile and attach two profiler dumps taken during a CPU spike. If the CPU spikes are long enough, please take the profiles 30-60 seconds apart. The most common cause for CPU spikes is a virtual machine operating system.
- If the CPU is spiking to 100%, try Live Monitoring Using the JMX Interface.

Next steps

Open a ticket on https://support.atlassian.com and attach all the data you have collected. This should give us the information we need to track down the source of your performance problems and suggest a solution. Please follow the progress of your enquiry on the support ticket you have created.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

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**
2. Select '**General Configuration**' in the left-hand panel.
3. Enable '**Compress HTTP Responses**'.

It is possible to configure which types of content are compressed within Confluence. By default, the following mime types will be compressed:

- text/htmltext
- javascript
- text/css
- text/plain
- application/x-javascript
- application/javascript

If you wish to change the types of content to be compressed, add a replacement `urlrewrite-gzip-default.xml` file within the `WEB-INF/classes/com/atlassian/gzipfilter/` directory in your Confluence Installation Directory. A sample file is provided as an attachment. It is unlikely that you will need to alter this file.

Garbage Collector Performance Issues

The information on this page relates to memory management with Oracle's Hotspot JVM. These recommendations are based on our support team's successful experiences with customers with large Confluence instances.

Which garbage collector?

Confluence uses the garbage first garbage collector (G1GC) by default. This is the garbage collector we recommend.

SeeGarbage First Garbage Collector Tuningin the Oracle documentation for useful information on tuning this garbage collector.

We have also observed that G1GC performs better with a larger heap (2gb). See the information below about how to increase your heap size gradually.

> ⚠ Don't use the Concurrent Mark Sweep (CMS) Collector with Confluence, unless advised by Atlassian Support. It requires extensive manual tuning and testing, and is likely to result in degraded performance.

Use the right size heap

Keep your heap as small as possible, without the instance experiencing OutOfMemory errors. If you experience OutOfMemory errors and need to increase this, we recommend you do it in 512mb or 1gb allotments, and monitor the instance. If you continue to receive OutOfMemory errors, increase the heap by another 512mb or 1gb, and continue this process until you are operating stably with no OutOfMemory errors. Do not increase the heap further than required, as this will result in longer garbage collections.

Remove any old tuning parameters

On every full GC, the JVM will resize the allocations of Eden, Survivor etc based on the throughput it is actually seeing. It will tune itself based on the real world data of the objects that are being created and collected. Most of the time simply allowing JVM to tune itself will give you better performance.

If you have added JVM parameters in the past and are experiencing difficulties with GC now, we'd recommend you remove all GC related parameters, unless you added them to solve a specific problem, and they did in fact solve that problem. You should also consider re-benchmarking now to ensure that they are still solving that problem, and are not causing you any other issues.

Check your VM resources

If you run Confluence on a VM, check that is it not using the swap file. If it does, when the JVM garbage collects it has to load the objects from the swap file into memory to clean them, and this can cause significantly longer GC pauses. Instead of using swapping, ballooning and bursting, allocate adequate memory to the VM.

Manual Tuning

If you find you are still experiencing difficulties with GC after following these recommendations and you would like to see if you can tune the JVM better to improve performance, see ourGarbage Collection (GC) Tuning Guide.This document was put together a few years ago, but has some useful information on choosing performance goals (throughput/footprint/latency), and how to tune for those goals.

Viewing your GC logs

How to Enable Garbage Collection (GC) Logging, and use a tool likeChewiebug's GCViewerto view the resulting logs.

Troubleshooting Slow Performance Using Page Request Profiling

This page tells you how to enable page-request profiling. With profiling turned on, you will see a record of the time it takes (in milliseconds) to complete each action made on any Confluence page. If Confluence is responding slowly, an internal timing trace of the slow page request can help to identify the cause of the delay.

You will need access to the Confluence server to view a profile.

On this page:

- Profiling an Activity
- Example of a Profile
- Start Confluence with Profiling Enabled

Related pages:

- Requesting Performance Support
- Working with Confluence Logs

Enable page-request profiling

You needSystem Administratorglobal permissions to enable or disable profiling.

To enable or disable page profiling:

1. Go to > **General Configuration**>**Logging and Profiling**.
2. If you run Confluence in a cluster, select a cluster node.
3. Choose **Enable Profiling** or**Disable Profiling**.

Screenshot: Changing Log Levels and Profiling

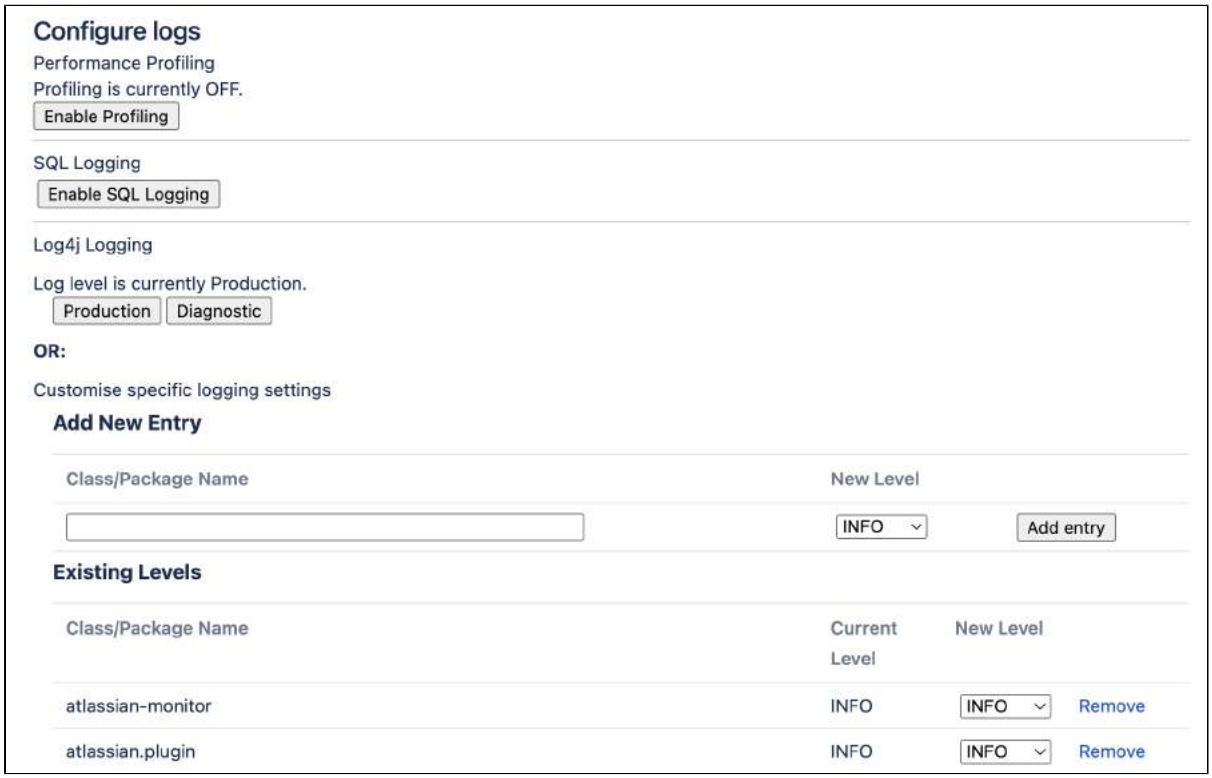

Profiling an Activity

1. Enable profiling, using either of the methods described above.
   Profiles for every page hit, for all users, will now be logged to your application server's default logs until Confluence is restarted. Note that each time a user visits a link, a single profile is printed.
2. Confirm that profiles are being written to the Confluence log file see Working with Confluence Logs for location of the log files and other details.
3. Perform the activity that is resulting in unusually slow response time.

4. Copy the profile for that action. When deciding which profiles to copy, look for the links that took a long time to respond. If a single page is slow, only that profile is necessary. If Confluence is generally or intermittently slow, copy all profiles logged during the slowdown until a reasonable sample has been collected.
5. If you were instructed to profile your instance by Atlassian technical support, attach all relevant profiles to your support ticket.
6. Turn profiling off again, using either of the methods described above.
7. Confirm that profiles are no longer being printed to the Confluence log file.

Example of a Profile

Below are the first few lines of a normal profile for accessing a page called Confluence Overview.

```
[344ms] - /display/ds/Confluence+Overview
  [313ms] - SiteMesh: parsePage: http://localhost:8080/display/ds/Confluence+Overview
    [313ms] - XW Interceptor: Before defaultStack: /pages/viewpage.action (ViewPageAction.execute())
      [0ms] - SpaceAwareInterceptor.intercept()
      [16ms] - PageAwareInterceptor.intercept()
        [0ms] - AOP: PageManager.getPage()
        [16ms] - AOP: PermissionManager.hasPermission()
          [0ms] - AOP: SpacePermissionManager.hasPermission()
          [16ms] - AOP: SpacePermissionManager.hasPermission()
        [0ms] - AOP: SpacePermissionManager.hasPermission()
      [0ms] - AOP: SpacePermissionManager.hasPermission()
      [281ms] - XW Interceptor: After defaultStack: /pages/viewpage.action (ViewPageAction.execute())
        [281ms] - XW Interceptor: After validatingStack: /pages/viewpage.action (ViewPageAction.
execute())
           ...
```

Notice that each indented line is a recursive call that rolls up into the parent line. In the example above, the Confluence Overview page takes 344ms. Part of that, 313ms, is spent in sitemesh.

Start Confluence with Profiling Enabled

There may be some situations where you may wish to have Confluence profiling enabled during startup. This may be useful if you restart often and may forget to enable profiling for Support/Trouble-shooting purposes.

Edit the file *CONFLUENCE_INSTALL\confluence\WEB-INF\web.xml*. You should see a section similar to the one below. Set the parameter value for **autostart** to **true**:

```
    <filter>
       <filter-name>profiling</filter-name>
       <filter-class>com.atlassian.confluence.util.profiling.ConfluenceProfilingFilter</filter-class>
       <init-param>
           <!-- specify the which HTTP parameter to use to turn the filter on or off -->
           <!-- if not specified - defaults to "profile.filter" -->
           <param-name>activate.param</param-name>
           <param-value>profile</param-value>
       </init-param>
       <init-param>
           <!-- specify the whether to start the filter automatically -->
           <!-- if not specified - defaults to "true" -->
           <param-name>autostart</param-name>
           <param-value>true</param-value>
       </init-param>
   </filter>
```

Remember to turn it back to **false** or your logs will grow very large.

Identify 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 to help identify slow performing macros.

1. Go to > **General Configuration**>**Logging and Profiling**.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

2. If you run Confluence in a cluster, select a cluster node.
3. Add the`com.atlassian.renderer.v2.components.MacroRendererComponent` package, and set the level to DEBUG.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Confluence Diagnostics

When investigating a performance problem or outage, it's useful to know as much as possible about what was happening in your site in the lead-up to the problem. This is when diagnostics information can help.

While often not individually actionable, diagnostic alerts can help you build up a detailed picture of your site's behaviour, and identify symptoms that may be contributing to the problem.

> This feature is still experimental in Confluence 6.11. We plan to fine-tune the thresholds and provide a UI for this diagnostic information in an upcoming Confluence release. Stay tuned!

About diagnostic alerts

The purpose of the diagnostics tool is to continuously check for symptoms or behaviours that we know may contribute to problems in your site. An alert is triggered when a set threshold is exceeded.

For example,if the free disk space for your local home (or shared home) directory falls below 8192MB, an alert is triggered. This is useful because if you run out of space, your users may not be able to upload new files, export spaces, or perform other tasks that rely on writing files to that directory.

It's important to note that the thresholds are just the point at which the alert is triggered. It's not the same as a timeout, or other hard limit. For example a long running task may trigger an alert after 5 minutes, and still complete successfully after 8 minutes.

When an alert is triggered a message is written to the `atlassian-confluence.log`file (your application log), and further details provided in the `atlassian-diagnostics.log` file. It's also included in support zips.

Some behaviours trigger a single alert, for others, multiple alerts are possible. Diagnostic information is stored in the database, and retained for 30 days. Old alerts are cleaned up automatically.

Types of alerts

There a several types of alerts.

Alert and KB	Level	Default threshold	Configurable
Low free disk space	Critical	8192 megabytes	Yes
Low free memory	Warn	256 megabytes	Yes
Node left or joined the cluster	Info	N/A	No
Long running task exceeded time limit	Warn	300 seconds	Yes
Garbage collection exceeded time limit	Warn	10% (over the last 20 seconds)	Yes

Availability

Some diagnostic alerts are disabled by default, because they may have a performance impact on your site, or are not designed to run continuously.

Our support team may ask you to enable one of the following alerts when troubleshooting a specific problem. They'll provide you with information on how to do this.

Alert and KB	Level	Default threshold	Configurable
HTTP request exceeded time limit	Warn	60 seconds	Yes
Macro rendering exceeded time limit	Warn	30 seconds	Yes

Thread memory allocation rate exceeded limit	Warn	5% over the last 20 seconds)	Yes
Sandbox crashed or was terminated during document conversion	Info	N/A	No

Alert levels

There are three levels of diagnostic alerts:

- **Info** - information that might be useful when troubleshooting a problem, for example a node joined the cluster
- **Warning** - a problem that may impact performance or availability in future, for example low memory
- **Critical** - a serious problem that is likely to impact system stability or availability, for example low disk space.

Most alerts don't require any immediate action.

Change alert thresholds

Some alert thresholds are configurable.If you find you are seeing too many instances of an alert, you can change the threshold, so it's not triggered so easily.

Head toRecognized System Properties for a list of system properties for each alert. This info can also be found on the knowledge base articlefor each alert.

Change diagnostics behaviour

You can also change the way the diagnostics framework itself behaves. For example, you might change how often checks are performed, or how long diagnostics information is retained.

Head toRecognized System Properties for the full list of system properties.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Faster permissions service

Confluence's default permission checking method is very efficient when checking permissions for a single page, and guarantees strong permission consistency. However, this method can be quite slow and memory intensive when Confluence needs to check many thousands of pages.For example, to render the Task report macro, we need to find all the pages with a task assigned to the user, then check they have permission to see the spaces and pages the task appears on. The faster permissions service allows Confluence to check permissions on a large number of pages more quickly.

On this page:

- Enable the faster permissions service
- Disable the faster permissions service
- Troubleshooting

How it works

The faster permissions service replicates permissions information in a database structure that can be queried more efficiently. It uses database triggers to register changes to pages, spaces and their permissions, and records these in a change log. These change logs are processed regularly, and keep the fast permissions database tables up to date.

When faster permissions is enabled:

- the faster permissions method is used when listing pages in the task report, children display, and page index macros, search requests when the cache is cold, and on the dashboard the first time someone logs in.
- the default permissions checking method is always used when viewing a page, and in all other situations.

When space permissions or page restrictions are changed, it can take a few seconds for the faster permissions service to know about the change. Here's how we handle this delay:

- When a page is created and page restrictions are applied before being published for the first time, the page won't appear to anyone in macros or the dashboard until the fast permissions service has processed the change log.
- When a previously unrestricted page is restricted, it will continue to appear in macros and the dashboard for a few seconds, until the fast permissions service has processed the change log, then it will only appear to the relevant users.
- When a previously unrestricted page is restricted and the page title changed, it will continue to appear with the original title in macros and the dashboard for a few seconds, until the fast permissions service has processed the change log, then it will only appear (with the new title) to the relevant users.

The time it takes for a change to be reflected in faster permissions tables is shown in the faster permissions screen as **Space service lag** and **Content service lag**. If the lag exceeds 60 seconds, Confluence will automatically fall back to the slower default permission checking method until it decreases again. This is to ensure Confluence doesn't rely on out of date permissions information.

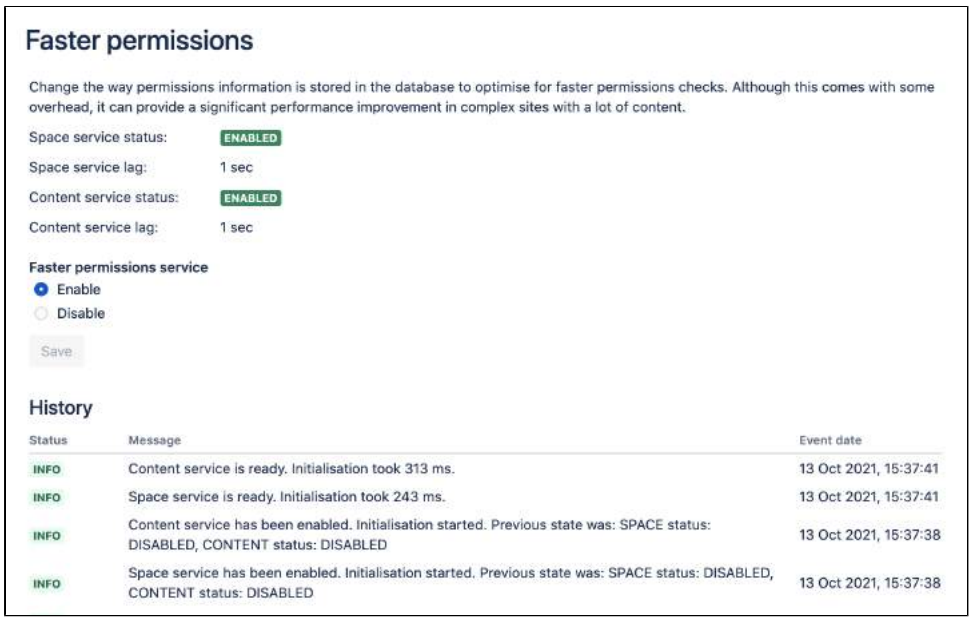

Screenshot of the faster permissions administration screen.

Enable the faster permissions service

The faster permissions service is disabled by default. If you have a small site, enabling the faster permissions service may not be beneficial.

To enable the faster permissions service:

1. Go to > **General Configuration**> **Faster Permissions**.
2. Select**Enable**.

The status of each service will change to initializing as it starts up and begins to populate the new database tables. This may take some time depending on the size of your site, and is done in small batches to avoid any performance impact. Confluence is still completely operational during this time, and will start using each service as soon as the status changes to enabled.

Disable the faster permissions service

To disable the faster permissions service:

1. Go to > **General Configuration**> **Faster Permissions**.
2. Select**Disable**.
3. Wait for the services to shut down.

Confluence will use thedefault permission checking method from this point on.The additional database tables that were added when the service was first initialised will not be removed. If you later re-enable the service, the service will re-populate the tables.

Troubleshooting

Database requirements for faster permissions service

- If you use SQL Server, your database user needs full create, read and write permissions for the database tables. Confluence must be able to create its own schema, and have the ability to create /drop triggers and functions. Refer to your SQL Server documentation for how to do this.
- If you use Oracle, you must use Oracle 12c Release 2, there are known problems with Release 1 (which is no longer supported by Confluence).

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Confluence guardrails

Background

Were committed to supporting the needs of our largest customers, and this includes continually improving the performance and scalability of our products. The amount of data in your instance can be a factor in performance and stability problems. As your instance grows, so does your risk of performance degradation over time. Often this is a gradual degradation and can go unnoticed until you reach a point where it has a significant impact on your team.

In the table below, weve described the performance and stability impacts that weve observed and suggested some actions you can take to reduce your risk. The guardrails are based on real-world experiences with some of our largest customers, but wont necessarily be representative of every organizations experience.

Ways you can reduce the risk of experiencing serious performance and stability problems may include:

- application changes, such as upgrading to a newer application version to get the benefit of performance improvements, or changing the way users are managed.
- infrastructure changes, such as increasing memory, CPU, or running a cluster or mirrors.
- data cleanup activities to reduce your footprint, such as archiving or breaking up monolith sites.

Its important to note that these arent hard limits, and some of your product instances may already exceed these thresholds. There are a number of factors, including the interplay between different data types, and site load, which will influence whether you experience the potential impacts listed below, and to what degree. As with any type of risk, its essential to identify the risk and make a plan, so you can prioritise those actions that will help you reduce the probability of future performance problems.

Definition

Product **Guardrails** are data type recommendations designed to help you identify potential risks and aid you making decisions about next steps in your instance optimization journey.

Confluence guardrails

The following guardrails are provided to help you identify and mitigate scale risks, and make decisions about cleaning-up your instance.

Index size

Content type	Total size of the index files in the home directory
Guardrail	30GB total file size, or 10,000,000 Current versions items
How to find this number	How to check the size of your index
Risks	We've observed these problems when operating above this guardrail: • Slow search results. • Querying data from the index takes a long time. • Reindex takes a long time.

Mitigation options	• Use SSD disks for your local home and shared home directories. This can improve how quickly Confluence can add, update, or retrieve information from the index. • Dont index the content of attachments if you dont need the contents to be searchable. Learn how to disable indexing attachments • Delete spaces that are no longer needed. Learn how to identify and remove unused spaces • Migrate some spaces to Cloud to reduce the size of your Data Center instance. Learn how to migrate individual spaces to Cloud • Consider splitting your site into 2 or more instances. If splitting your instance is a viable option for your organization, we strongly recommend you to get assistance from a Atlassian Partner to successfully execute the split, and get advice on federating your instances. • Permission complexity. When searching, Confluence needs to check whether the user is permitted to see the content. • Database type and performance, This is because permissions are not indexed and must be checked against the database. Weve observed that PostgreSQL databases are more efficient to check permissions than the other supported types. • Network latency between the database and application node affects how quickly the application can write to the index. Hosting the database and application nodes in the same availability zone can help.

Spaces

Content type	Total number of spaces
Guardrail	10,000 spaces
How to find this number	Check the number of spaces in your site
Risks	We've observed these problems when operating above this guardrail: • High memory and CPU consumption whenever Confluence needs to perform permission checks to determine what pages display, for example on the dashboard, and in macros.
Mitigation options	• Enable the faster permissions service. Learn how to enable the faster permissions service (available from Confluence 7.15, or as an experimental feature from Confluence 7.12) • Delete spaces that are no longer needed. Learn how to identify and remove unused spaces • Migrate some spaces to Cloud to reduce the size of your Data Center instance. Learn how to migrate individual spaces to Cloud • Consider splitting your site into 2 or more instances. If splitting your instance is a viable option for your organization, we strongly recommend you to get assistance from a Atlassian Partner to successfully execute the split, and get advice on federating your instances.

Space size (for import)

Content type	Total number of pages, blogs, attached files, version history, and trash in a single space
Guardrail	Total size of 5GB for the entities.xml file within the space export zip file.
How to find this number	Check the size of a space before importing

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Risks	We've observed these problems when operating above this guardrail: • Out of memory errors when importing a space, which could result in application crashes. • High CPU and memory consumption when importing a space, which affects overall performance of the site.
Mitigation options	• Split the space prior to export. Move some parts of the page tree to a new space, import both spaces, then move the pages back to the main space. • Use retention rules to reduce the number of page and attachment versions, and items in the trash (expected to be available from Confluence 7.16) • Consider apps like Better Archiving for Confluence and Script Runner to reduce the size of the space.

LDAP users

Content type	Total number of users synchronised between LDAP and Confluence
Guardrail	If using Microsoft Active Directory: • 100,000 users If using another connector: • 70,000 users
How to find this number	How to get the number of users or groups
Risks	We've observed these problems when operating above this guardrail: • Full sync takes a very long time. • Potential for high CPU and memory consumption when identifying group memberships during permission checks.
Mitigation options	• If you use Microsoft Active Directory, enable incremental synchronization. This fetches changes from LDAP, avoiding the need for a full sync. • Use Crowd to take advantage of features like: • Access Based Synchronisation, which only synchronises users that have access to an application. Learn about access based synchronization • Use the User, Group, and Membership schema configuration filters to restrict the data synchronised with Confluence. Learn how to connect to an LDAP directory

LDAP groups

Content type	Total number of groups synchronised between LDAP and Confluence
Guardrail	If using Microsoft Active Directory: • 30,000 groups If using another connector: • 20,000 groups

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

How to find this number	How to get the number of users or groups
Risks	We've observed these problems when operating above this guardrail: • Instance instability, including performance degradation and potential outages when Confluence is under high load • Directory synchronization takes a long time • User authentication can take longer than expected • Application access and group management admin screens can become unresponsive

Mitigation options	• If you use Microsoft Active Directory, enable incremental synchronization. This fetches changes from LDAP, avoiding the need for a full sync.
	• Use Crowd to take advantage of features like:
	  • Access Based Synchronisation, which only synchronises users that have access to an application. Learn about access based synchronization
	  • Use a Delegated directory so that Crowd can import users' group memberships from LDAP each time they authenticate. Learn how to configure a delegated authentication directory
	• Use the User, Group, and Membership schema configuration filters to restrict the data synchronised with Confluence. Learn how to connect to an LDAP directory

Depth of nested groups

Content type	Number of levels of hierarchy when groups are nested
Guardrail	4 levels deep. We also recommend groups do not contain a mix of users and other groups, as this can also influence performance.
How to find this number	You cant get this number directly from Confluence. Youll need to look at the hierarchies defined in your external directory.
Risks	We've observed these problems when operating above this guardrail: • Instance instability, including performance degradation and potential outages when Confluence is under high load • Directory synchronization takes a long time • User authentication can take longer than expected
Mitigation options	• Change the group structure in your directory to avoid having too many levels of nesting • Change the group structure in your directory so that groups only contain either users or other groups.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

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 inourPrivacy Policy. Please read it - we've tried to avoid legal jargon and made it as straightforward as possible.

To view a sample of data that might be collected from your specific installation, go to > **General Configuration** > **Analytics**.

Data is always collected in Confluence Cloud.

How is data collected from Confluence?

Older versions of Confluence (prior to Confluence 5.6 or Confluence Questions 1.0.618) didn't collect usage data. Analytics are collected using the Atlassian Analytics system app. The app collects analytics events in a log file which is located in `<confluence-home>/analytics-logs`. The logs are periodically uploaded using an encrypted session and then deleted. If Confluence is unable to connect to the Internet, no logs are ever uploaded.

Enabling/disabling data collection in Confluence

You can turn off analytics collection at any time. Go to > **General Configuration** > **Analytics**.

Managing emojis

Users can add custom emojis to create personalized and engaging content. For example, add icons to documentation for greater consistency, or upload logos to use as emojis in internal communications. Learn more about emojis

As a system administrator:

- you can control who can add an emoji either logged-in users or system admins only
- you can delete uploaded emojis
- you can change the maximum upload number
- you can change the maximum upload size

By default, any logged-in user can upload emojis using the emoji menu in the editor.

Manage who can add an emoji

To enable users to add their own emojis:

1. > **General Configuration** > **Emojis**
2. Under**Permissions,** toggle the switch on to allow logged-in users to upload custom emojis

If the switch is off, only system admins will be able to add custom emojis.

Delete uploaded emojis

To view and delete custom emojis uploaded to your Confluence site:

1. > **General Configuration** > **Emojis**
2. Under **Manage custom emojis,** find the emoji you want to delete
3. Select **Delete** in the same table row
4. You will receive a confirmation message, select**Delete** again

If the deleted emoji is in use on pages, blogs, or comments, it will be replaced by its text emoji shortcut.

You can also delete any custom emojis you've personally uploaded from the emoji menu in the editor.

Other emoji configurations

You can change these default emoji settings by configuring system properties.

Setting	Default	System property
File size	1MB	`confluence.emoticons.max.file.size`
Site upload limit	2000	`com.atlassian.confluence.plugins.emoticons.max.allowed.uploads`
Image resizes allowed in parallel	The maximum number of processors available to the virtual machine; never smaller than one	`emoticon.thumbnail.generator.permits.size`

Administering Collaborative Editing

Collaborative editing takes teamwork to the next level. This page covers everything you need to know about administering collaborative editing.

Head to Collaborative editing to find out how your team canwork together in real time on software requirements, meeting notes, retros, and any other Confluence page you can think of.

On this page:

- About Synchrony
- Change the editing mode
  - What happens to existing drafts when the mode changes?
- Maximum editor limit
- Auditing considerations
  - No version history in unpublished drafts
  - Visibility of edits made by anonymous users
- Proxy and SSL considerations
  - SSL
  - Proxies
  - WebSockets
- Change your Synchrony configuration
- Start and stop Synchrony
- Monitor Synchrony
- Accessing Synchrony logs
- Managing Synchrony data

Related pages:

- Troubleshooting Collaborative Editing

About Synchrony

Collaborative editing is powered by Synchrony which synchronizes data in real time. Confluence manages Synchrony, so administrators should rarely need to interact with it directly.

Synchrony runs on port 8091 by default, and an internal Synchrony proxy means that you shouldn't need to open this additional port.

How you connect to Synchrony will depend on your environment, and your Confluence license. See Possible Confluence and Synchrony Configurations.

To see your collaborative editing and Synchrony setup, head to > **General Configuration**>**Collaborative editing**.

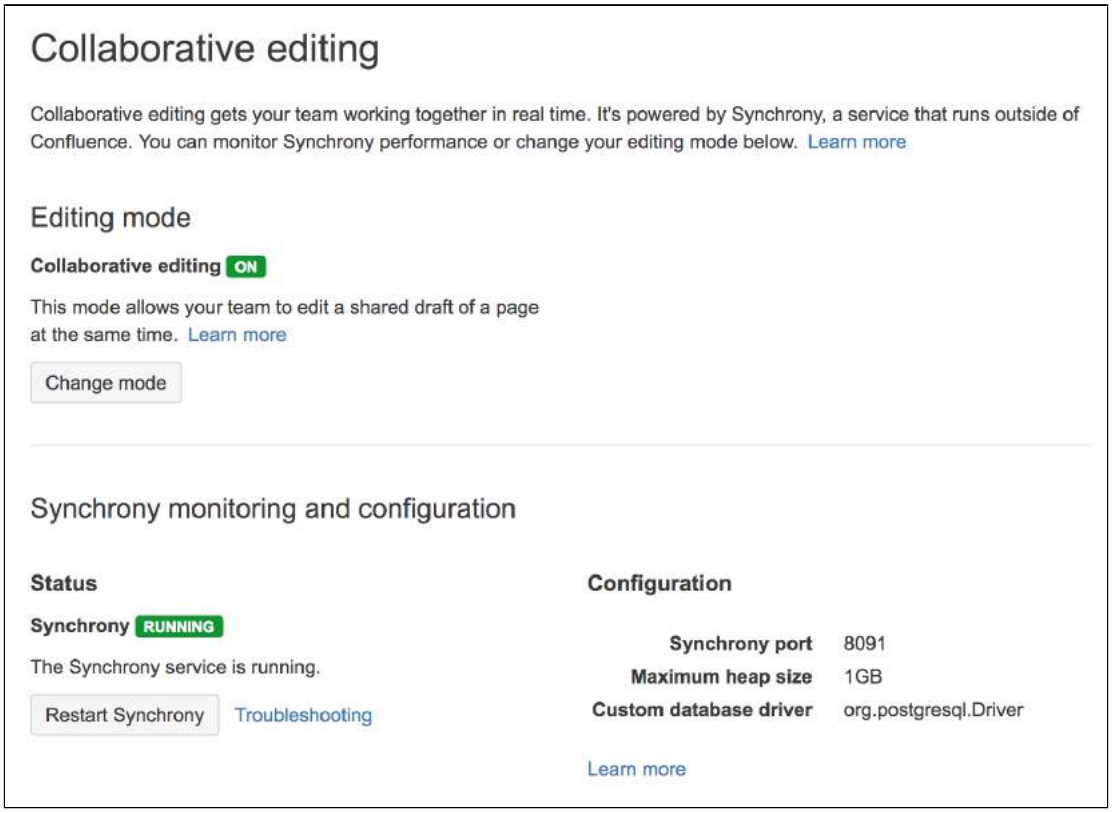

Change the editing mode

The editing mode determines the editing experience for all users in your site. This is how you turn collaborative editing on or off.

To change the editing mode:

1. Go to > **General Configuration**>**Collaborative editing**.
2. Choose**Change mode**.
3. Select either **On** or **Off** and choose**Change**.

Changing the editing mode is not trivial, so it is good to understand the implications of each mode.

Mode	Implications
On	This mode allows your team to edit a shared draft of a page at the same time, and see each others' changes in real time. This is the recommended editing mode.
Off	This mode means that your team can only edit their own personal draft of a page. Confluence will attempt to merge any conflicts on save. Consider turning collaborative editing on for the full experience This mode is useful if you are unable to run Synchrony successfully in your environment, or if you have decided that collaborative editing is not for you (for example if you have auditing requirements that would prohibit using collaborative editing just yet). It's a good idea to prompt your users to publish any shared drafts before you turn collaborative editing off, as they will not be able to resume editing existing shared drafts or unpublished changes.

What happens to existing drafts when the mode changes?

Users can always access any existing personal drafts and shared drafts from the **Drafts** page in their profile. Whether they can resume editing the draft depends on the editing mode.

When collaborative editing is **ON**,users will be able to discard or resume editing any personal or shared drafts. A personal draft will be converted to a shared draft when a user resumes editing.

When collaborative editing is **OFF**,users will be able to discard or resume editing any personal drafts. They can't resume editing existing shared drafts, but can view and copy the contents of those drafts.

Shared drafts will only appear in a user's Drafts page if, when collaborative editing was on, they:

- created a draft, and never published it
- edited a published page, and did not publish their changes.

Maximum editor limit

A maximum of 12 people can edit a page at the same time. This means that people can't enter the editor if there are already 12 other people editing the page, and will need to wait until someone leaves.

Administrators can increase or decrease this limit using a system property. If you experience performance issues when many people are editing, you might want to decrease this limit.

Auditing considerations

We know that auditing is a major consideration for some customers. We don't yet have very granular auditing capabilities with collaborative editing. All page changes are currently attributed to the person that publishes the page, rather than the people who made each specific change.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

If this is going to be a problem in your site, we recommend turning collaborative editing off in your site for now.

No version history in unpublished drafts

We're saving all the time in collaborative editing, but we don't save versions of unpublished changes. When restoring an earlier page version, you can only roll back to an existing published version. Any unpublished changes will be lost when you restore a previous version.

Visibility of edits made by anonymous users

There are some additional things to be aware of if you have granted the **Add**page permission (and **Can use** global permission) to anonymous users.

You won't be alerted, when closing the editor or publishing a page, if anonymous users made the only unpublished changes on the page. This means a logged in user may inadvertently publish changes they were not aware had been made to the page.

The changes themselves are visible in the page, but the usual warning dialog will not appear if the only people to have made changes were not logged in.

If there are unpublished changes from both logged in users and anonymous users, the warning dialog will appear, but only the logged in users will be listed in the dialog. Changes made by all users (including anonymous) will be included if you view the changes from that dialog.

Proxy and SSL considerations

How you connect to Synchrony will depend on your environment. We know that most Confluence sites run behind a reverse proxy, often with SSL. Here's some information to help you identify the right configuration for your environment, and any changes you might need to make to your environment to use collaborative editing in your site.

SSL

Synchrony runs in a separate JVM, and does not support direct HTTPS connections. If you are not using a reverse proxy, SSL should be terminated at Tomcat. If you are using a reverse proxy or load balancer, SSL should be terminated at your reverse proxy or load balancer.

SeePossible Confluence and Synchrony Configurations for detailed diagrams and examples.

Proxies

If you run Confluence behind a reverse proxy, you should take a look at the Possible Confluence and Synchrony configurations for guidance on how your Confluence and Synchrony setup may impact your proxy.

SeePossible Confluence and Synchrony Configurations for detailed diagrams and examples, plus links to example proxy configuration files.

WebSockets

For best results, your load balancer and proxies should allow WebSocket connections. If your users cannot get a WebSocket connection, Confluence will fall back to an XML HTTP Request (XHR), allowing them to edit pages successfully.

XHR fallback is enabled by default, but can be disabled using asystem property(passed to Confluence)if necessary. You shouldn't need to change this.

Change your Synchrony configuration

You can't change your Synchrony configuration through the Confluence UI. In most cases you shouldn't need to make changes to the default configuration.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

If you need to change the port Synchrony runs on or the maximum memory available, for example, you can do this using a system property, or in your start-synchrony script (if you're running your own Synchrony cluster).

SeeConfiguring Synchronyfor more information.

Start and stop Synchrony

If Synchrony is**managed by Confluence**(recommended), Confluence will automatically start Synchrony for you when it starts up. You can also restart Synchrony from the collaborative editing admin screen in Confluence.

If you're running**Synchrony standalone in a cluster**, you'll use the`start-synchrony.sh`or`start-synchrony.bat`. scripts on each Synchrony node. A process ID (PID) file will be created in your synchrony directory.

Stop Synchrony the same way, using`stop-synchrony.sh`or`stop-synchrony.bat`. This will destroy the PID file that the start script created in your Synchrony directory. If you've customized the location for storing the PID file in the`start-synchrony`script, you'll need to also update this in the`stop-synchrony`script.

If you're unable to start Synchrony, check that there isn't an existing PID file in your Synchrony directory.

Monitor Synchrony

To check if Synchrony is running, go to> **General Configuration**>**Collaborative editing**.

If you're running Confluence in a cluster, you can check the status of Synchrony on each node from the clustering screen.

Go to> **General Configuration**> **Clustering**, then on each node choose **Collaborative editing**. You can access all nodes in this way, you don't need to hit a specific node in your browser.

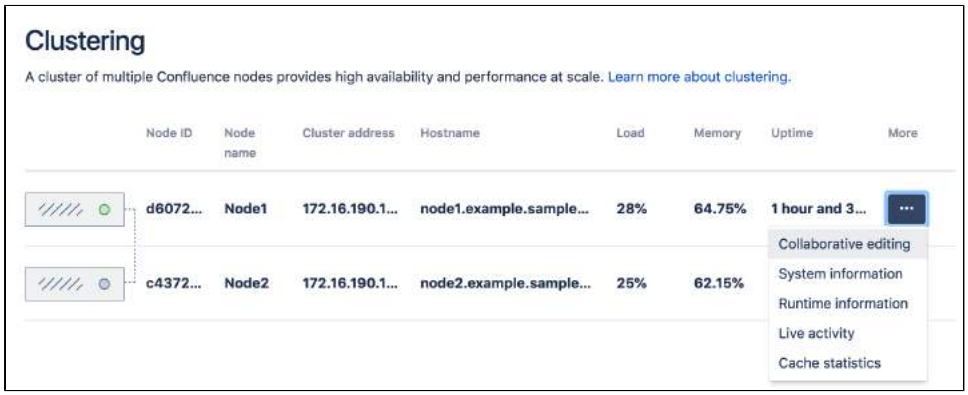

From here you can see the Synchrony status, mode, and URL Confluence is using to connect to it. Here's what it looks like when Synchrony is managed by Confluence.

All Confluence nodes must use the same Synchrony mode. For example, you can't have one node using managed Synchrony, and another node connecting to a standalone Synchrony cluster.

Accessing Synchrony logs

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

If Synchrony is**managed by Confluence**(recommended), Synchrony logs will be stored in your`<local-home>/logs directory`, with the Confluence application logs.

If you're running**Synchrony standalone in a cluster**, your Synchrony logs will be stored in the Synchrony directory on each Synchrony node (wherever you run the start and stop scripts from).

To learn how to change the logging level, seeConfiguring Synchrony.

Managing Synchrony data

Each page and blog post has its own Synchrony change log, which contains a graph of all edits to that page or blog post. In busy Confluence sites the database tables that store the Synchrony change logs can grow very quickly. Because the change logs store all changes as they happen, they may retain personally identifiable information, even after the page they relate to has been deleted.

We provide two scheduled jobs for removing Synchrony data:

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

The soft eviction job runs regularly in the background. The hard eviction job is available for when you need to remove Synchrony data more aggressively, and is disabled by default.

SeeHow to remove Synchrony datafor more information on how these jobs work.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Possible Confluence and Synchrony Configurations

Synchrony is the engine that powers collaborative editing in Confluence.

There are a few different options for running Synchrony, and it is worth taking some time to determine which option will best meet the needs of your organisation.

On this page:

- Possible configurations for Confluence Server
- Possible configurations for Confluence Data Center

Possible configurations for Confluence Server

If you have a Confluence Server license, Synchrony is**managed by Confluence**.

The diagrams below show examples of a common implementation where Confluence is running under the /confluence context path (e.g. www.mysite.com/confluence). The concepts are the same if you use Confluence without a context path (e.g.www.myconfluence.com).

No reverse proxy

If you don't run Confluence behind a reverse proxy, you'll connect to Synchrony via Confluence's internal Synchrony proxy. SSL, if used, is terminated at Tomcat. This is the default configuration, and you shouldn't need to make any additional changes to use collaborative editing.

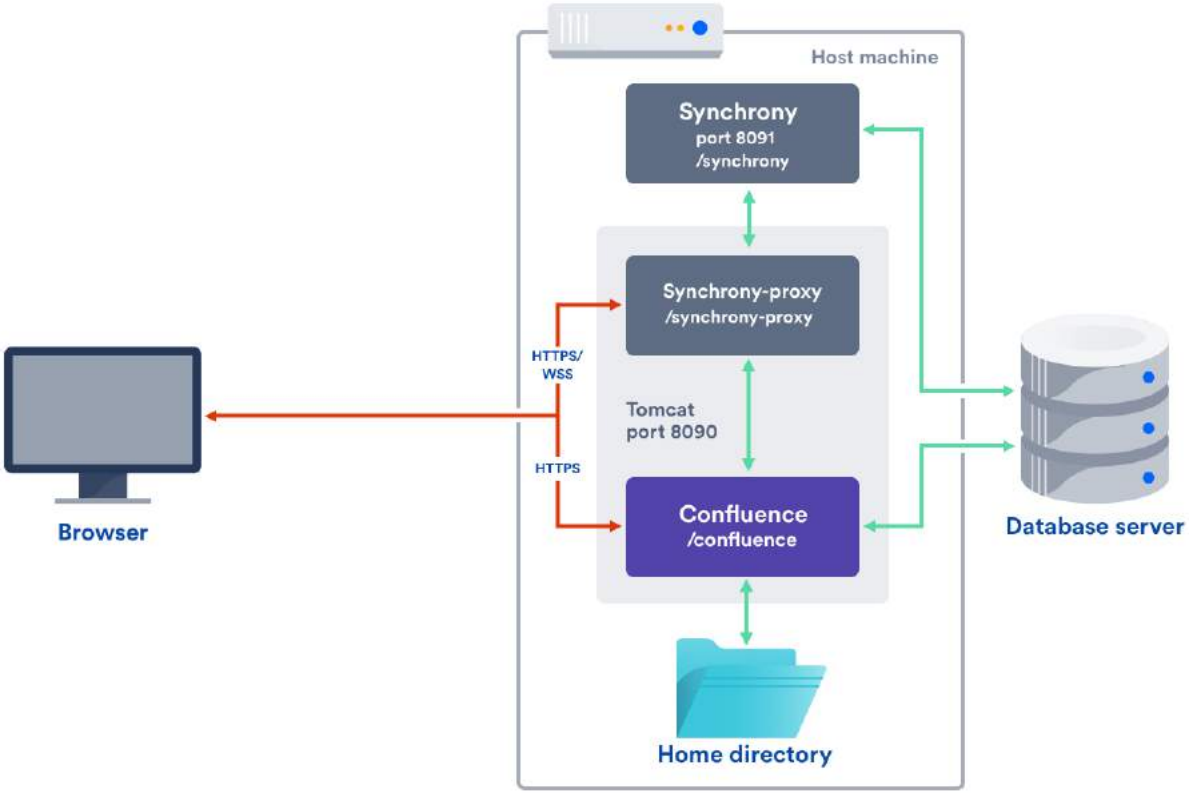

With a reverse proxy

If you run Confluence behind a reverse proxy, you will connect to Synchrony via Confluence's internal Synchrony proxy. This is the default configuration with a reverse proxy, and a good choice if you do not want to open port 8091. SSL should be terminated at your reverse proxy.

You do not need to make any additional changes to your reverse proxy configuration for Synchrony, but for best results your reverse proxy must support WebSocket connections (you may need to manually enable this in your proxy).

To tell Confluence that you want to use the internal proxy, set the`synchrony.proxy.enabled`system propertytotrue. (This is optional, but will prevent Confluence from trying to reach Synchrony via /synchrony first, before retrying via the internal proxy).

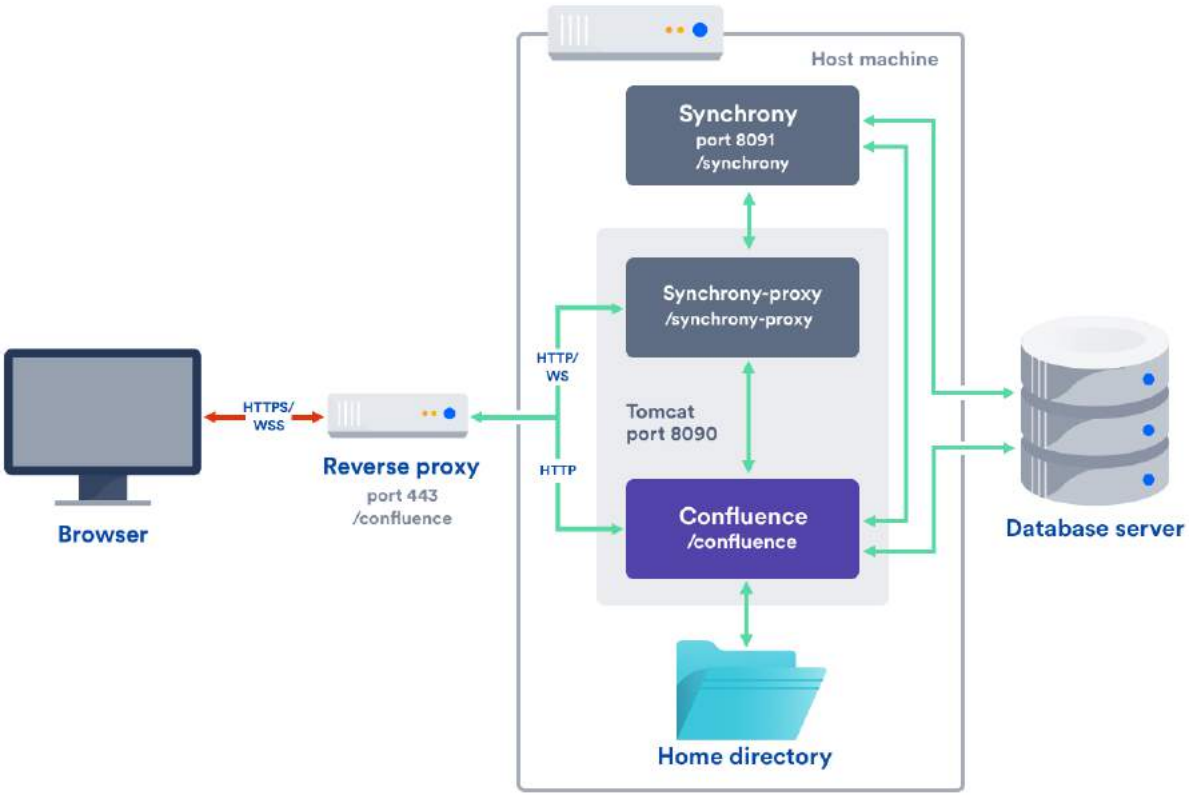

If Synchrony can't be reached via /synchrony-proxy we'll automatically try /confluence/synchrony-proxy (where /confluence is your Confluence context path).

Direct to Synchrony with a reverse proxy

If you run Confluence behind a reverse proxy, and experience latency or other issues connecting to Synchrony via Confluence's internal Synchrony proxy, you can choose to connect direct to Synchrony. This is the optimal setup, but does require some changes to your environment. You will need to open port 8091 and add /synchrony to your reverse proxy configuration.SSL will still be terminated at your reverse proxy, as Synchrony does not accept direct HTTPS connections.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

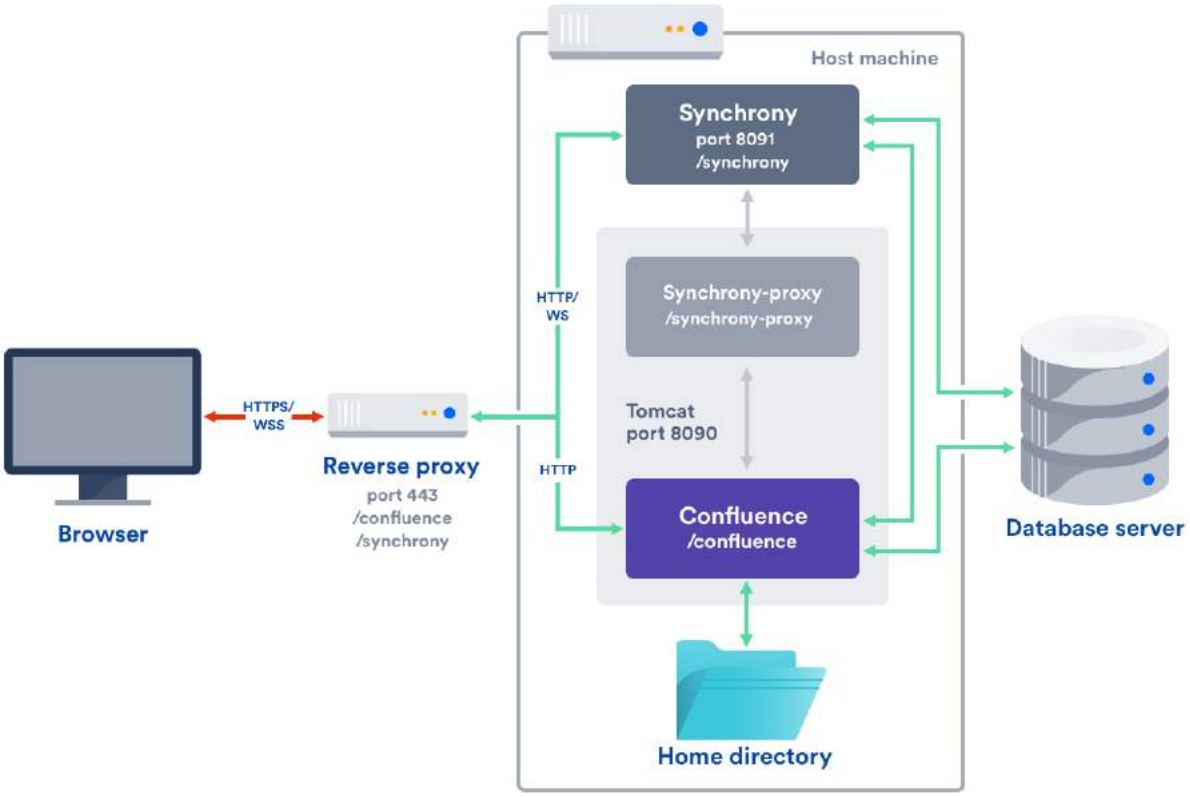

If Synchrony can't be reached via /synchrony we'll automatically try the internal Synchrony proxy via /confluence/synchrony-proxy (where /confluence is your Confluence context path).

See the following guides for example reverse proxy configurations. The order of directives is important, so check our examples.

- Using Apache with mod_proxy
- Running Confluence behind NGINX with SSL
- Proxying Atlassian server applications with Apache HTTP Server (mod_proxy_http)
- Proxying Atlassian server applications with Microsoft Internet Information Services (IIS)
- How to configure Amazon Web Service Application Load Balancer with Confluence

Possible configurations for Confluence Data Center

If you have a Confluence Data Center license, two methods are available for running Synchrony:

- **managed by Confluence**(recommended)
  Confluence will automatically launch a Synchrony process on the same node, and manage it for you. No manual setup is required.
- **Standalone Synchrony cluster (managed by you)**
  You deploy and manage Synchrony standalone in its own cluster with as many nodes as you need. Significant setup is required.During a rolling upgrade, you'll need to upgrade the Synchrony separately from the Confluence cluster.

If you want simple setup and maintenance, we recommend allowing Confluence to manage Synchrony for you. If you want full control, or if making sure the editor is highly available is essential, then managing Synchrony in its own cluster may be the right solution for your organisation.

Managed by Confluence

Here's a simplified view of the architecture when Synchrony is**managed by Confluence**. This is the recommended approach, as no manual set up, or ongoing upgrades are required - it works right out of the box.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

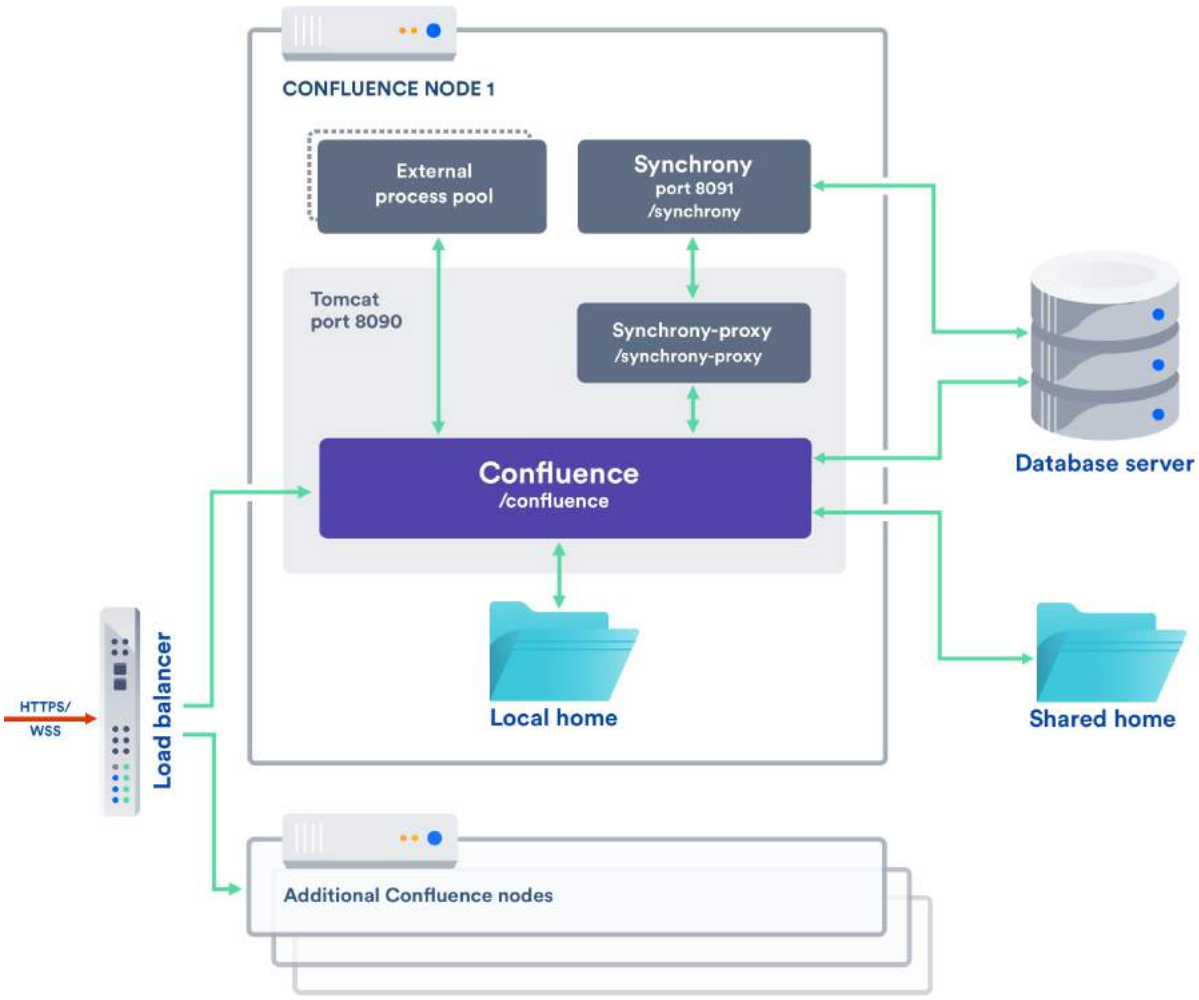

If you choose this approach, the guidance provided for Confluence Server above, also applies to your load balancer and any proxies.

This option is only available in Confluence 6.12 and later.

Standalone Synchrony cluster

If you choose to**manage Synchrony yourself**, the architecture looks more like this. Again the diagram has been simplified, and doesn't show communication between nodes.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

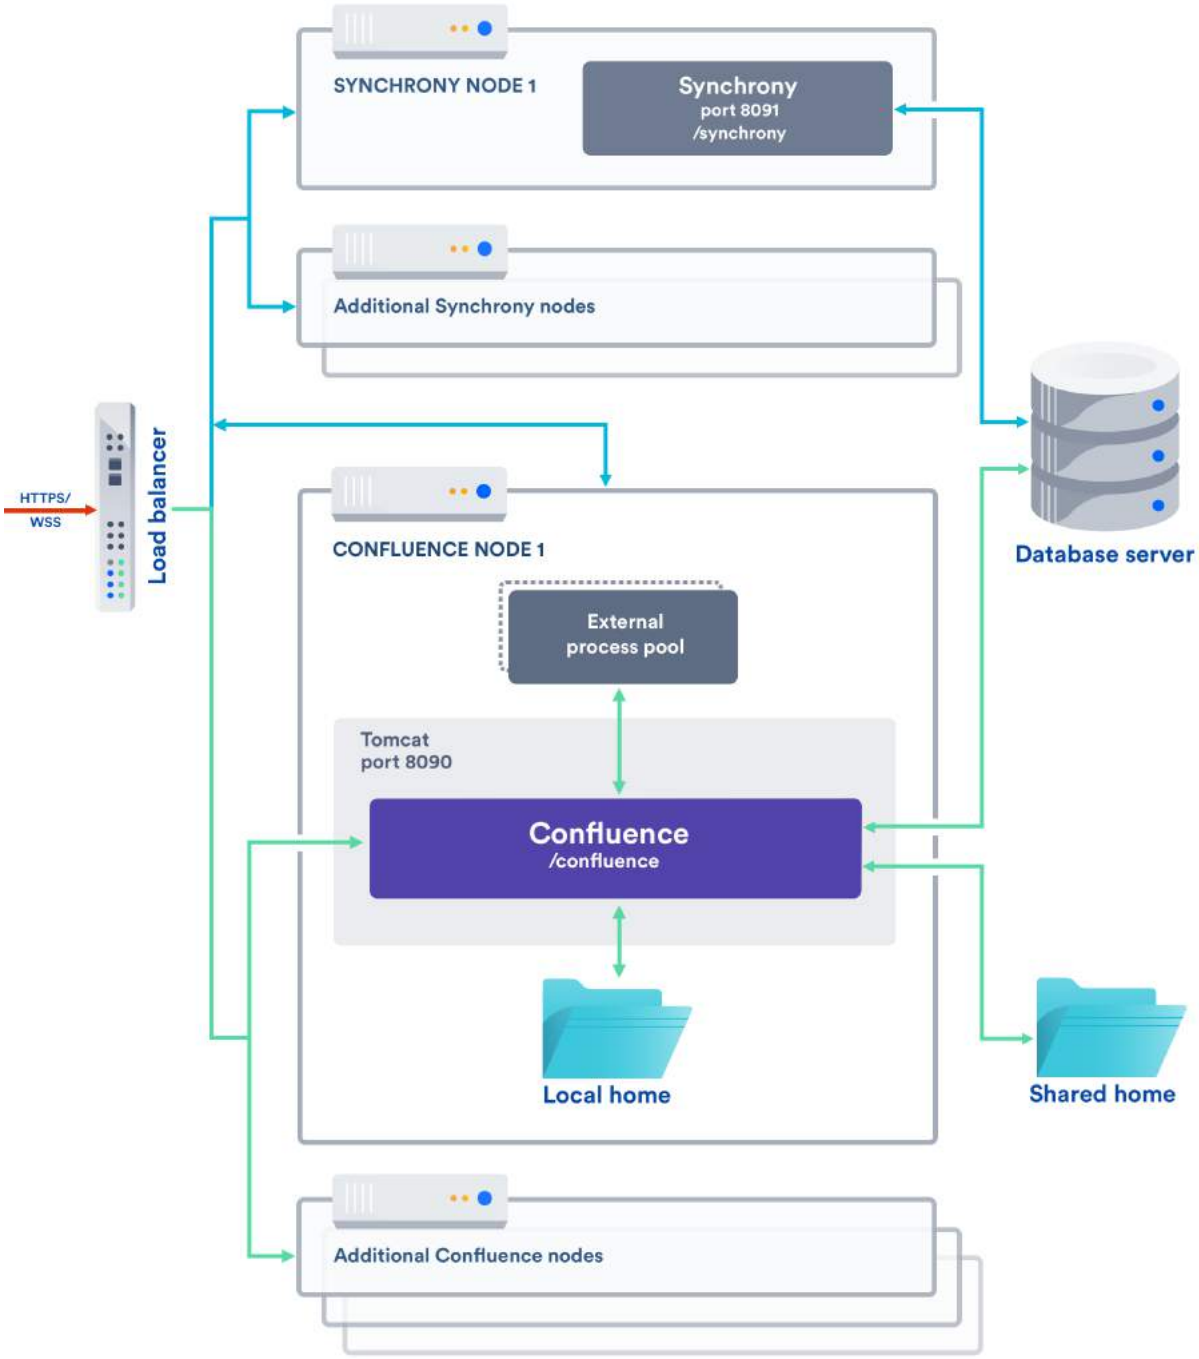

If you choose this approach you will need to:

- Set up and manage multiple Synchrony cluster nodes.
- Always start Confluence with the `synchrony.service.url`system property (this tells Confluence where to find your Synchrony cluster, instead of launching a Synchrony process on the current node).
- Open the Synchrony port (8091), as the Synchrony proxy is never used when you manage Synchrony yourself.
- Terminate SSL at your load balancer. Synchrony cannot accept HTTPS connections.
- Upgrade Synchrony manually, each time you upgrade Confluence.

In most cases, two Synchrony nodes will be adequate for multiple Confluence nodes.

For a step-by-step guide to setting up your Synchrony cluster, seeSet up a Synchrony cluster for Confluence Data Center.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

If you enabled collaborative editing prior to Confluence Data Center 6.12, standalone Synchrony will be your default setup.

If you would prefer a less complex setup, seeMigrate from a standalone Synchrony cluster to managed Synchronyto find out how to allow Confluence to manage Synchrony for you.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Configuring Synchrony

Synchrony is the engine that powers collaborative editing in Confluence.

There's no UI for configuring Synchrony.Configuration changes, such as changing the Synchrony port or memory settings, are made via system properties.How you pass these properties depends on whether Synchrony is managed by Confluence, or deployed as a seperate cluster.

In most cases, Synchrony is **managed by Confluence**.

If you have a Data Center license, you may choose to deploy and manage **Synchrony standalone in a cluster**, instead of allowing Confluence to manage Synchrony for you. SeePossible Confluence and Synchrony Configurations for more information.

On this page:

- Passing recognized system properties to Synchrony
- Common configuration changes
- Change the logging level for managed Synchrony
- Change the logging level for Synchrony standalone
- Troubleshooting

Passing recognized system properties to Synchrony

If Synchrony is **managed by Confluence** (the most common setup), you make changes to Synchrony by passing system properties to Confluence. SeeConfiguring System Properties to find out the best way to do this for your operating system.

You can find a full list of system properties atRecognized System Properties.

If you're running**Synchrony standalone in a cluster**, you pass properties directly to Synchrony via the start-synchrony scripts.

Note that the properties are not always the same as those used when Synchrony is managed by Confluence. A full list of required and optional properties can be found atSet up a Synchrony cluster for Confluence Data Center.

Passing JVM arguments to Synchrony

Sometimes you may want to pass additional arguments, that are not already provided by a system property, directly to Synchrony's JVM.

If Synchrony is**managed by Confluence**, you will need to create a file called `synchrony-args.properties` in your home directory (or shared home if you have a Data Center license) and include the arguments you want to pass to Synchrony, one per line.

For example:

```
synchrony.jvm.arg.0=-Dproperty1=value1
synchrony.jvm.arg.1=-Dproperty2=value2
...
synchrony.jvm.arg.N=-XX:NumberOfGCLogFiles=5
```

For more examples, see Configuring JVM garbage collection logging for Synchrony process.

Remember, you can't use this method for passing any value that is already handled by a Confluence system property, such as synchrony.port, Xmx or Xss.

If you're running**Synchrony standalone in a cluster**, you pass arguments to Synchrony's JVM directly, by adding them to your start-synchrony script, in the**Optional Overrides** section.

Common configuration changes

The two most common changes people make to Synchrony is to change the port that Synchrony runs on, if port 8091 is already in use, and to change the maximum heap memory allocated to Synchrony.

Change the port Synchrony runs on

Synchrony runs on port 8091 by default. If this port is already in use by another application on your server you can use the the`synchrony.port` system property to change it to an available port.

If you're running Confluence 6.0.3 or earlier you'll need to use `reza.port` instead of `synchrony.port`.

To change the maximum heap for Synchrony

Synchrony has a maximum heap size of 2 GB by default.

If you experience out of memory errors related to Synchrony, you can change the heap size allocated to Synchrony using the`synchrony.memory.max`system property.

If you're running Confluence in a cluster, you may want to increase the maximum heap size to 4gb on each node.

Change the logging level for managed Synchrony

The logging level for managed Synchrony is set to `INFO`by default. If you find this too verbose, you can decrease the logging level to `WARN`or `ERROR`.

To change the managed Synchrony logging level:

1. Create a file called `synchrony-log4j.properties`with the following content:

   ```
   log4j.rootLogger=WARN, stdout
   log4j.appender.stdout=org.apache.log4j.ConsoleAppender
   log4j.appender.stdout.Target=System.out
   log4j.appender.stdout.layout=org.apache.log4j.PatternLayout
   log4j.appender.stdout.layout.ConversionPattern=%d %p [%t] [%c{4}] %m%n
   ```

   In this example we'll set the logging level to `WARN`. Replace this with `ERROR`if you only want to log errors.

2. Save the file.You can place the file anywhere, but we recommend your home directory (or shared home) alongside the `synchrony-args.properties`file.
3. Edit your`<home-directory>/synchrony-args.properties`file. If you're running Confluence in a cluster, this will be in your shared home directory.
4. Add the following line to tell Synchrony where to find your log configuration.

   ```
   log4j.configuration=file://<path-to-file>/synchrony-log4j.properties
   ```

   Replace `<path-to-file>`with your file path. In Linux this will be something like`=file:///var/confluence/local-home/synchrony-log4j.properties`, for example.
5. In Confluence, go to > **General Configuration** > **Collaborative editing** and select **Restart Synchrony**to pick up the changes.

Exclude the Confluence DEBUG prefix

Because Synchrony is managed by Confluence, the Synchrony logs include a prefix with information from Confluence itself. You can omit this prefix to make the logs easier to read.

To omit the Confluence DEBUG prefix from the Synchrony logs:

1. Edit the`<install-directory>/confluence/WEB-INF/classes/log4j.properties` file.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

2. Change the `log4j.appender.synchronylog.layout.ConversionPattern` line to remove `%d %p [%t] [%c{4}]` as follows:

```
log4j.appender.synchronylog.layout.ConversionPattern=%m%n
```

3. Save the file, then restart Confluence to pick up the changes.

If you're running Confluence in a cluster, you'll need to repeat this process on each Confluence node.

Change the logging level for Synchrony standalone

If you choose to deploy and manage **Synchrony standalone in a cluster**, you can configure the logging level in your start-synchrony script.

To change the Synchrony standalone logging level:

1. Create a file called `synchrony-log4j.properties` with the following content:

```
log4j.rootLogger=WARN, stdout
log4j.appender.stdout=org.apache.log4j.ConsoleAppender
log4j.appender.stdout.Target=System.out
log4j.appender.stdout.layout=org.apache.log4j.PatternLayout
log4j.appender.stdout.layout.ConversionPattern=%d %p [%t] [%c{4}] %m%n
log4j.category.com.hazelcast=INFO
log4j.category.hazelcast=INFO
```

   In this example we want to set the logging level to `WARN`. Replace this with `ERROR` if you only want to log errors. We keep the Hazelcast logging level at `INFO` so you can still see the Synchrony nodes communicating with each other.

2. Save the file. You can place the file anywhere, but we recommend your Synchrony directory.
3. Edit your `<synchrony-directory>/start-synchrony.sh` or `start-synchrony.bat` file.
4. Add the following line in the Optional Overrides section to tell Synchrony where to find your log config:

```
log4j.configuration=file://<path-to-file>/synchrony-log4j.properties
```

5. Restart Synchrony to pick up the changes.

Repeat this process on each Synchrony node.

Troubleshooting

- If you have a Data Center license, and Synchrony is managed by Confluence, we recommend storing the `synchrony-args.properties` file in the shared home directory, so that all Synchrony processes are started with the same JVM arguments. If you do locate the `synchrony-args.properties` file in the local home, the arguments will only be passed to the Synchrony process on that node.
- Since 7.20, we enabled passing JVM arguments to Synchrony to help with diagnosing Synchrony-related issues. See Configuring JVM garbage collection logging for Synchrony process.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Set up a Synchrony cluster for Confluence Data Center

If you have a Confluence Data Center license, two methods are available for running Synchrony:

- **managed by Confluence**(recommended)
  Confluence will automatically launch a Synchrony process on the same node, and manage it for you. No manual setup is required.
- **Standalone Synchrony cluster (managed by you)**
  You deploy and manage Synchrony standalone in its own cluster with as many nodes as you need. Significant setup is required.During a rolling upgrade, you'll need to upgrade the Synchrony separately from the Confluence cluster.

If you want simple setup and maintenance, we recommend allowing Confluence to manage Synchrony for you. If you want full control, or if making sure the editor is highly available is essential, then managing Synchrony in its own cluster may be the right solution for your organisation.

On this page:

- Architecture overview
- Set up a Synchrony standalone cluster
  - 1 Provision your Synchrony nodes
  - 2 Create the Synchrony home directory
  - 3 Edit the start and stop scripts
  - 4 Add additional Synchrony nodes and configure your load balancer
  - 5 Start Confluence one node at a time
  - 6 Enable collaborative editing
- Required properties for Synchrony standalone
- Optional properties for Synchrony standalone
- Run Synchrony standalone in an IPv6 environment
- Run Synchrony standalone as a service
- Provide credentials to Synchrony standalone using environment variables

On this page we'll guide you through the process of setting up a standalone Synchrony cluster, hosted on your own infrastructure. The ability to run your own Synchrony cluster is only available with a Data Center license.

Architecture overview

Here's a simplified view of the architecture when you manage Synchrony yourself, in a seperate cluster. Note that this diagram doesn't show communication between nodes.

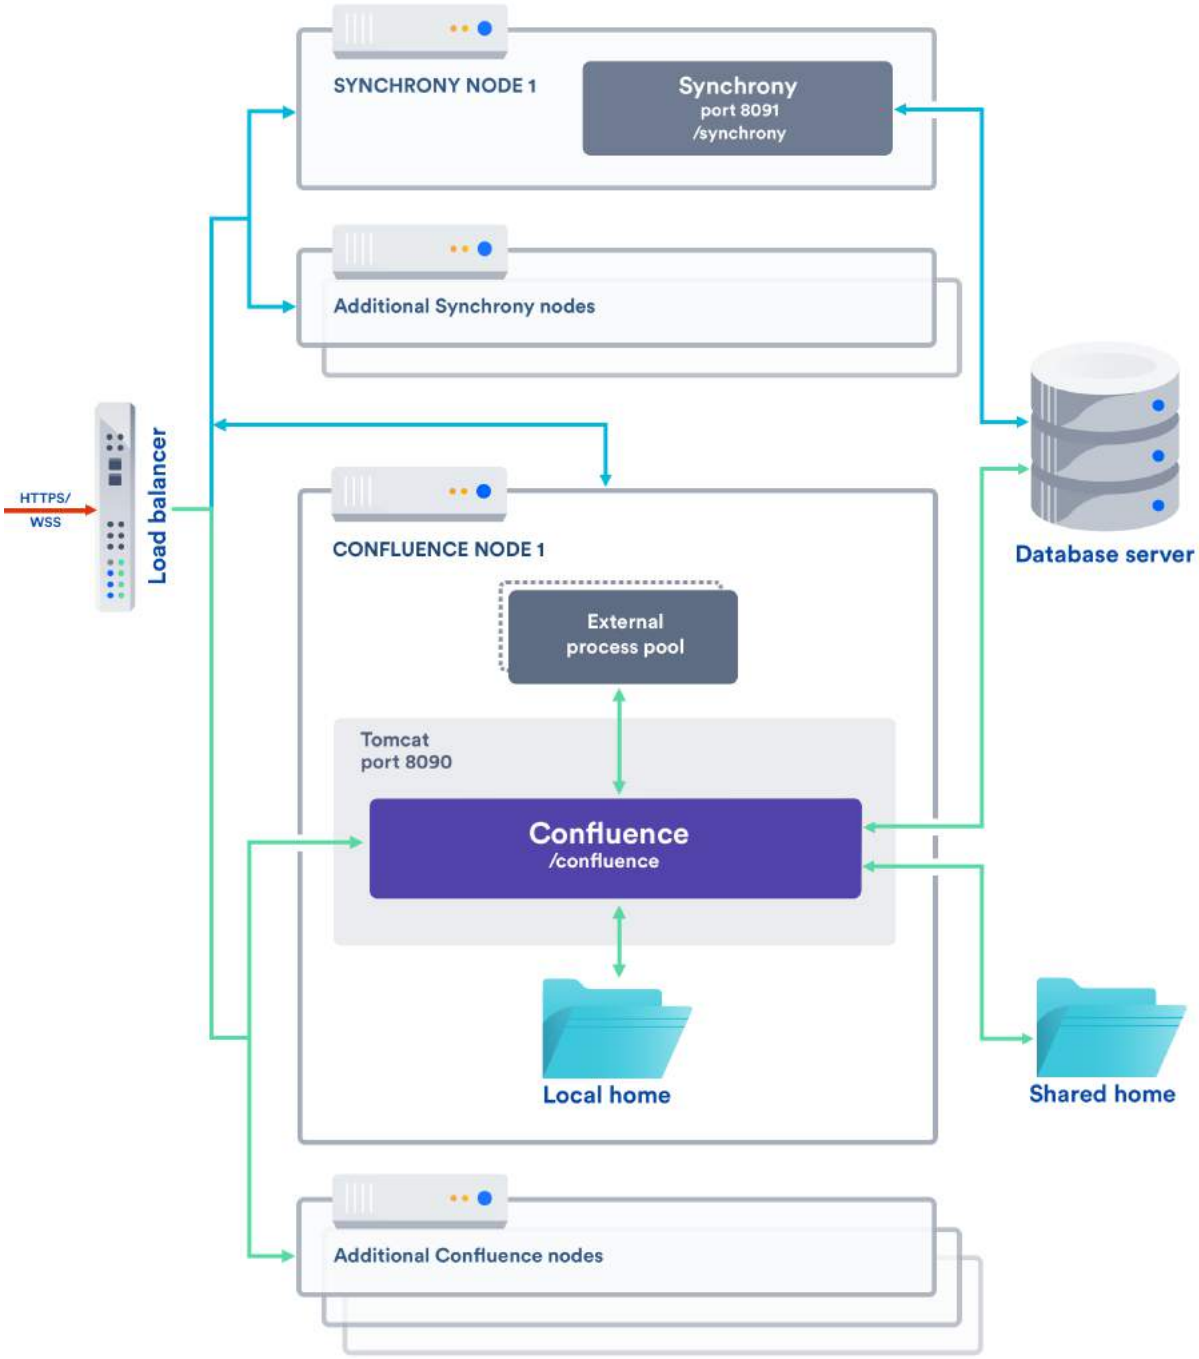

If you enabled collaborative editing prior to Confluence Data Center 6.12, standalone Synchrony will be your default setup.

If you would prefer a less complex setup, seeMigrate from a standalone Synchrony cluster to managed Synchronyto find out how to allow Confluence to manage Synchrony for you.

Set up a Synchrony standalone cluster

This page will guide you through setting up a Synchrony standalone cluster on your own infrastructure.

If you're using AWS or Azure, using one of our templates may be a more efficient way to set up Confluence with a standalone Synchrony cluster.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

1 Provision your Synchrony nodes

For the purposes of this guide, we assume you have already provisioned the hardware or virtual instances for your Synchrony nodes. We recommend starting with 2 Synchrony nodes.

You should allow 2GB memory for Synchrony, and enough disk space for the Synchrony application and logs.

2 Create the Synchrony home directory

To create the Synchrony directory on your first Synchrony node:

1. Grab the`<install-directory>/bin/synchrony`directory from one of your Confluence nodes and move it to your new Synchrony node. We'll call this your`<synchrony-home>`directory.
2. Copy`synchrony-standalone.jar`from your Confluence local home directory to your`<synchrony-home>`directory.
3. Copy your database driver from your Confluence`<install-directory>/confluence/web-inf/lib`to your`<synchrony-home>`directory or other appropriate location on your Synchrony node.

3 Edit the start and stop scripts

We provide scripts to start and stop Synchrony on each node. These need to be edited to add information about your environment:

1. Edit the`<synchrony-home>/start-synchrony.sh`or`start-synchrony.bat`file
2. Enter details for all of the required parameters listed under**Configure parameters**.
   SeeRequired properties below, for a description of each.
3. Enter detail for any optional properties you may want to specify.
   SeeOptional properties below for a description of each.
4. Save the file.
5. Start Synchrony by running the start-synchrony script.
6. Visit`http://<SERVER_IP>:<SYNCHRONY_PORT>/synchrony/heartbeat`to check Synchrony is running.

4 Add additional Synchrony nodes and configure your load balancer

To create your second Synchrony node:

1. Copy your<synchrony-home>directory to the second Synchrony node.
2. Start Synchrony on that node using the start-synchrony script.As each node joins you'll see something like this in your console.

```
Members [2] {
        Member [172.22.52.12]:5701
        Member [172.22.49.34]:5701
}
```

3. Configure your load balancer for Synchrony traffic.
   For best results, your load balancer should allow WebSocket connections. SSL connections must be terminated at your load balancer, as Synchrony can't accept HTTPS requests.

   You can choose to use the same load balancer for both Confluence and Synchrony, or two seperate load balancers. When we refer to the Synchrony load balancer, we mean whichever load balancer is handling Synchrony traffic.

4. Make sure the Synchrony port (8091) is open.Ports used by Atlassian Applicationshas a good summary of all ports Synchrony uses in Data Center. This is the only one that needs to be open.

5 Start Confluence one node at a time

Now that Synchrony is running in a cluster, it's time to get Confluence involved. It is essential that you stop Confluence on **all nodes** before continuing.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

1. Stop Confluence on all nodes.
2. Start Confluence on **one node** with the followingsystem property.
   This property is used to tell Confluence where to find Synchrony, and prevents Confluence fromautomatically launching a Synchrony process on your Confluence node.

   ```
   -Dsynchrony.service.url=http://<synchrony-load-balancer-url>/synchrony/v1
   ```

   For example `http://42.42.42.42/synchrony/v1`or `http://synchrony.example.com/synchrony/v1`
3. Check that Confluence can connect to Synchrony. Head to > **General Configuration**>**Clustering** then choose••• > **Collaborative editing**beside the Confluence node you just started.

   The Synchrony mode should be**Standalone Synchrony cluster**.

   If the mode is 'Managed by Confluence', your Confluence node is not connected to your Synchrony cluster. Make sure you're passing the Synchrony service URL system property correctly.
4. Repeat this process, starting each Confluence node, one at a time, with the `synchrony.service.url`.

SeeHow to check the status of Synchrony for Confluence Data Centerfor more info on how to check Synchrony is running.

6 Enable collaborative editing

If you're installing Confluence for the first time, collaborative editing is enabled by default. If you've upgraded from an earlier Confluence version, or have disabled it in the past, collaborative editing may still be disabled.

To enable collaborative editing:

1. Head to > **General Configuration**>**Collaborative editing**.
2. Choose **Change mode**.
3. Select **On** and choose **Change**.

You can now try editing a page. You'll need to access Confluence via your load balancer. You can't create or edit pages when accessing a node directly.

Any users who had the editor open before you made this change will need to refresh in order to continue editing, as the Synchrony URL they're connected to will have changed.

Required properties for Synchrony standalone

These properties **only apply** when you're running Synchrony standalone in its own cluster. If Synchrony is managed by Confluence (Server or Data Center) these properties don't apply.

The following properties must be provided in the start-synchrony script.

Property name	Description
`SERVER_IP`	Public IP address or hostname of this Synchrony node. It could also be a private IP address - it should be configured to the address where Synchrony is reachable by the other nodes.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

DATABASE_URL	This is the URL for your Confluence database. For example `jdbc:postgresql://yourserver:5432/confluence` . You can find this URL in `<local-home>/confluence.cfg.xml` .
DATABASE_USER	This is the username of your Confluence database user.
DATABASE_PASSWORD	(Optional)This is the password for your Confluence database user. If your password contains special characters, Synchrony may silently fail to connect to the database. Rather than hardcoding your password, we recommend setting your password with the environment variable `SYNCHRONY_DATABASE_PASSWORD`. Any dots (".") in variable names (identifiers) will need to be replaced with underscores ("_").
CLUSTER_JOIN_PROPERTIES	This determines how Synchrony should discover nodes. You'll be prompted to uncomment a set of parameters for either: • TCP/IP • Multicast • AWS Follow the prompts in the script for the values you need to enter for each of these.
DATABASE_DRIVER_PATH	This is the path to your database driver file. If you're running Synchrony on its own node, you'll need to copy your database driver to an appropriate location then provide the path to this location.
SYNCHRONY_JAR_PATH	This is the path to the `synchrony-standalone.jar` file you copied to this node.
SYNCHRONY_URL	This is the URL that the browser uses to contact Synchrony. Generally this will be the full URL of the load balancer Synchrony will run behind plus the Synchrony context path, for example `http://yoursite.com:8091/synchrony` . Note that it *does not* end with `/v1` , unlike the `synchrony.service.url` system property passed to Confluence. If this URL doesn't match the URL coming from a users' browser, Synchrony will fail.
OPTIONAL_OVERRIDES	You can choose to specify additional system properties. See the table below for recognised Synchrony system properties.
FEATURE_AUTH_TOKEN	This is the flag that turns on and off the Synchrony authentication on the Handshaking REST API. Set it to true, and set Synchrony `AUTH_TOKEN` and Confluence `synchrony.service.authtoken` from Recognized System Properties to the same value in order for Confluence calls to be authenticated correctly.
AUTH_TOKENS	This is an authentication token that must be included in the Handshaking REST API call to Synchrony in order for the request to be considered valid. It comes into effect only when `FEATURE_AUTH_TOKEN=true`.

Optional properties for Synchrony standalone

These properties only apply ifyou're running**Synchrony standalone in a cluster**.

When you start Synchrony, we pass default values for the properties listed below. You can choose to override these values by specifying any of these properties when you start Synchrony.

Property name	Default	Description

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

`cluster.listen.port`	5701	This is Synchrony's Hazelcast port. Specify this property if you do not want to use port 5701 or if it is not available. As with the Confluence Hazelcast port (5801) you should ensure that only permitted cluster nodes are allowed to connect to Synchrony's Hazelcast port, through the use of a firewall and or network segregation.
`synchrony.cluster.base.port`	25500	This is the Aleph binding port. Synchrony uses Aleph to communicate between nodes. Specify this property if you don't want to use the default.
`cluster.join.multicast.group`	224.2.2.3	If the cluster join type is multicast, you can specify an IP address for the multicast group if you don't want to use the default.
`cluster.join.multicast.port`	54327	If the cluster join type is multicast, you can specify a multicast port if you don't want to use the default.
`cluster.join.multicast.ttl`	32	If the cluster join type is multicast, this is the time to live threshold. The default, 32, means the scope is restricted to the same site, organization or department. Specify this property if you want to use a different threshold.
`cluster.join.aws.access.key`		If the cluster join type is AWS, this is your AWS access key.
`cluster.join.aws.secret.key`		If the cluster join type is AWS, you can authenticate by IAM role or Secret key. This is your AWS secret key.
`cluster.join.aws.iam`		If the cluster join type is AWS, you can authenticate by IAM role or Secret key. This is your AWS IAM role.
`cluster.join.aws.region`	us-east-1	If the cluster join type is AWS, this is the AWS region your Synchrony nodes will be running in.
`cluster.join.aws.security.group`		If the cluster join type is AWS, and you want to narrow the members of your cluster to only resources in a particular security group, specify the name of your AWS security group.
`cluster.join.aws.tag.key`		If the cluster join type is AWS, and you want to narrow the members of your cluster to only resources with particular tags, specify the AWS tag key.
`cluster.join.aws.tag.value`		If the cluster join type is AWS, and you want to narrow the members of your cluster to only resources with particular tags, specify the AWS tag key value.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

`cluster.join.aws.host.header`		If the cluster join type is AWS, t his is the AWS endpoint for Synchrony to use (the address where the EC2 API can be found, for example ' ec2.amazonaws.com ').
`cluster.join.aws.timeout`	5	If the cluster join type is AWS, this is the joining timeout (in seconds).
`cluster.interfaces`	Defaults to the same value as `SERVER_IP`	This is the network interface Synchrony will use to communicate between nodes. Specify this property if you don't want to use the default, which uses the value of the required property Defaults to the same value as `SERVER_IP` (also known as `synchrony.bind`).
`synchrony.cluster.bind`	Defaults to the same value as `SERVER_IP`	This is the Aleph binding address. This should be set to the same value as `cluster.interfaces`. Specify this property if you did not use the default value for `cluster.interfaces`.
`synchrony.port`	8091	This is the HTTP port that Synchrony runs on. If port 8091 is not available, specify this property to choose a different port.
`synchrony.context.path`	Defaults to the context path of `SYNCHRONY_URL`	This is the context path for Synchrony. There should be no need to change this.
`hazelcast.prefer.ipv4.stack`	True	If you're running Confluence in an IPv6 environment, you will need to set this property to `False`.
`cluster.authentication.enabled`	true	Set this property to false if you don't want to authenticate Synchrony nodes as they join the Synchrony cluster. This is not recommended. This property was added in 7.17.4.
`cluster.authentication.secret`	(automatically generated)	Set this property to change the shared secret used to authenticate Synchrony nodes as they join the Synchrony cluster. The secret must be a string of maximum 40 characters. This property was added in 7.17.4.

Run Synchrony standalone in an IPv6 environment

If you're running a**Synchrony standalone in a cluster**in an IPv6 environment, you will need to start Synchrony with the following JVM argument:

```
-Dhazelcast.prefer.ipv4.stack=false
```

If you're using the start-synchrony scripts, simply uncomment this line in the script.

Run Synchrony standalone as a service

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

If you're running**Synchrony standalone in a cluster**, and you'd prefer to run Synchrony as a service on each node, seeRun Synchrony-standalone as a service on Linux.

It's not possible to run Synchrony standalone as a service on Windows. Consider switching to managed Synchrony instead.

Provide credentials to Synchrony standalone using environment variables

If you're running**Synchrony standalone in a cluster**, and you prefer to store sensitive information in your environment, rather than directly in the Synchrony startup scripts you can create `asynchronyenv`file, and use it to provide your database credentials. This is only available in Linux environments.

SeeProvide credentials to Synchrony standalone using environment variables (Linux)

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Migrate from a standalone Synchrony cluster to managed Synchrony

If you have a Confluence Data Center license, and enabled collaborative editing prior to Confluence 6.12, you will likely be running standalone Synchrony, either in it's own cluster, or manually on each Confluence node.

If you'd prefer a simpler setup, with less ongoing maintenance, you can choose to let Confluence manage Synchrony for you. Confluence will automatically start up a Synchrony process when Confluence is started.

Some Confluence downtime is required for this process.

To switch from managing your own Synchrony cluster to letting Confluence manage Synchrony:

1. Configure your load balancer to direct traffic away from all Confluence and Synchrony nodes.
2. Stop Confluence and Synchrony on all nodes.
3. Remove the the`synchrony.service.url`system property. This property tells Confluence where to find your external Synchrony cluster.

    The way you remove this system property depends on how you run Confluence. Note that this system property is passed to **Confluence**, not Synchrony itself.

    If you start **Confluence manually on Windows**, edit the `<install directory>/bin/setenv.bat`file and remove the following line:

    ```
    set CATALINA_OPTS=-Dsynchrony.service.url=http://example-synchrony.com/synchrony/v1 %CATALINA_OPTS%
    ```

    If you start **Confluence manually on Linux**, edit the `<install directory>/bin/setenv.sh`file and remove the following line:

    ```
    CATALINA_OPTS="-Dsynchrony.service.url=http://example-synchrony.com/synchrony/v1 ${CATALINA_OPTS}"
    ```

    If you're running as a**Confluence as a Windows Service**, you'll need to edit the service and remove the following from the Java options:

    ```
    -Dsynchrony.service.url=http://example-synchrony.com/synchrony/v1
    ```

    SeeConfiguring System Properties for a step-by-step guide to passing system properties to Windows services via the command line or Windows Registry.

4. Set the `synchrony.memory.max`system property to increase the maximum heap memory available to Synchrony to 2gb (or the amount of memory previously allocated to the Synchrony standalone service).

    The way you set this system property depends on how you run Confluence. Note that this system property is passed to **Confluence**, not Synchrony itself.

    If you start **Confluence manually on Windows**, edit the `<install directory>/bin/setenv.bat`file and remove the following line:

    ```
    set CATALINA_OPTS=-Dsynchrony.memory.max=2g %CATALINA_OPTS%
    ```

    If you start **Confluence manually on Linux**, edit the `<install directory>/bin/setenv.sh`file and remove the following line:

    ```
    CATALINA_OPTS="-Dsynchrony.memory.max=2g ${CATALINA_OPTS}"
    ```

    If you're running as a**Confluence as a Windows Service**, you'll need to edit the service and remove the following from the Java options:

```
-Dsynchrony.memory.max=2g
```

SeeConfiguring System Properties for a step-by-step guide to passing system properties to Windows services via the command line, Windows Registry, or in AWS.

5. Make sure all required ports are open, especiallyespecially**5701 and 25500**which are used by the Synchrony cluster. SeeConfluence Server and Data Center portsfor a full list.
6. Start Confluence on one node.
7. In Confluence, edit a page and check that you can successfully make changes.
8. Repeat this process on each Confluence node, starting each node**one at a time**.

Once all nodes are back up and running, and you've confirmed that collaborative editing is working as expected, you can decommission your external Synchrony cluster, including removing any startup scripts or services you may have configured.

Any users who had the editor open before you made this change will need to refresh in order to continue editing, as the Synchrony URL they're connected to will have changed.

You may also need to make some changes to your load balancer configuration. SeePossible Confluence and Synchrony Configurations for more information.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Troubleshooting Collaborative Editing

Collaborative editing is powered by Synchrony which synchronizes data in real time.Under normal circumstances it should not need to be managed manually by an administrator.

This page will help you troubleshoot problems with Synchrony in your instance.

On this page:

- Troubleshooting collaborative editing problems
  - First steps
    - Check Synchrony is running
    - Check you can edit a page
    - Check the logs
    - Restart Synchrony
  - Check port 8091 is available
  - Reverse proxy issues
  - Forward proxy issues
  - Websocket issues
  - SSL issues
  - Memory issues
  - Multiple Synchrony processes
  - Mixed Synchrony modes in cluster
  - Incompatible browser extensions
  - Firewall or anti-virus interference
  - Too many people in the editor
- Create feedback reports to troubleshoot content problems
  - Allow users to self report problems
  - Report an editing problem
  - Report an editing problem by page ID
  - Send a report to Atlassian support

Related pages:

- Administering Collaborative Editing

Troubleshooting collaborative editing problems

First steps

Check Synchrony is running

To check if Synchrony is running, go to > **General Configuration**>**Collaborative editing** .

Note: if you're running Confluence Data Center, this page will only be able to tell you if the current Confluence node is connected to your Synchrony cluster. You may want to use a third party monitoring tool to help you monitor your Synchrony cluster. SeeHow to check the status of Synchrony for Confluence Data Centerfor more info.

Check you can edit a page

If you see an error when you edit a page, but Synchrony is running, something is preventing your browser from connecting to Synchrony.

The most common issue is a misconfigured reverse proxy. See our proxy troubleshooting tips later in this page or head toAdministering Collaborative Editingto find out more about possible proxy and SSL configurations.

Check the logs

You can find the Confluence application logs at `<home-directory>/logs/atlassian-confluence.log`and Synchrony specific logs at `<home-directory>/logs/atlassian-synchrony.log` .

Restart Synchrony

If Synchrony is managed by Confluence, go to > **General Configuration**> **Collaborative editing** and choose**Restart Synchrony**.

If you run your own standalone Synchrony cluster, manually restart Synchrony on each node.

Check port 8091 is available

Synchrony runs on port 8091 by default. If this port is already in use by another application on your server you can use the the`synchrony.port` system property to change it to an available port.

(If you're using Confluence 6.0.3 or earlier you'll need to use `reza.port` instead of `synchrony.port`.)

SeeConfiguring System Propertiesto find out how to change this.

For Confluence Data Center the way you run Synchrony is a little different. SeeConfiguring Synchrony for more information.

Reverse proxy issues

If you have configured your reverse proxy, but can't edit pages, here's some things to check in your configuration:

- Go to `installation-directory>/conf/server.xml` and check the `Connector`directive. Make sure that you have correct values for `<protocol>`and `<proxyName>`. See the examples in the guides below for more information.
- The`http`connector always needs to be present in the`<installation-directory>/conf/server.xml`file, even if you're configuring SSL or using the AJP connector. The Synchrony health check uses HTTP and will fail if this connector is not present. Alternatively, if you do not want to include the http connector, you can use the`synchrony.proxy.healthcheck.disabled` system property to disable the health check.
- If you're using Apache, make sure you're using Apache 2.4 (with WebSockets support) and all required modules have been enabled (`mod-proxy`,`mod_rewrite`,`proxy_wstunnel`).
- If you're using Apache and want to connect directly to Synchrony, in your proxy config file, make sure you've included `/synchrony`and that the order of the Confluence and Synchrony directives and location blocks is correct. See the examples in the guides below for more information.

SeeAdministering Collaborative Editingto find out more about possible proxy and SSL configurations then check out the following guides for more information on how to include Synchrony in your reverse proxy config, if you want to connect direct to Synchrony:

- Using Apache with mod_proxy
- Running Confluence behind NGINX with SSL
- Proxying Atlassian server applications with Apache HTTP Server (mod_proxy_http)
- Proxying Atlassian server applications with Microsoft Internet Information Services (IIS)
- How to configure Amazon Web Service Application Load Balancer with Confluence

Forward proxy issues

If you're using a forward or outbound proxy, you will need to add the IP that Synchrony listens on to your config to ensure it is bypassed. SeeConfiguring Web Proxy Support for Confluencefor more info.

By default, the IP is 127.0.0.1, or it will be the value of the`synchrony.host`system property, if you've customized the hostname or IP that Confluence uses to connect to Synchrony.

Websocket issues

Collaborative editingworks best with a WebSocket connection. If one can't be established due to a timeout, or a proxy server or firewall that doesn't allow WebSocket connections, the editor will attempt to connect via an XML HTTP Request (XHR).

You can usehttp://websocket.org/echo.htmlto perform a quickHTML5 WebSocket test against an echo server.

It's possible to verify the websocket connection using the developer tools in your browser. In the following example, we'll use Chrome to check the websocket connection:

1. Open the Chrome**Developer Tools**(**Shift**+**CTRL**+ J)
2. Select the**Network**tab
3. Select the**WS**filter to show only WebSocket traffic
4. Edit a Confluence page
5. Select the websocket connection:`ws://<confluence-url>/synchrony-proxy/v1/bayeux-sync1`
6. Select the Frames/Messages tab to see the traffic between browser and the WebSocket.

SSL issues

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Synchrony cannot accept direct HTTPS connections, so you will need to terminate SSL at your reverse proxy or load balancer, or at Tomcat if you are not using a reverse proxy.

Memory issues

If you experience out of memory errors related to Synchrony, you can change the heap size allocated to Synchrony using the `synchrony.memory.max` system property.

If you're Confluence 6.0.3 or earlier you'll need to use `reza.memory.max` instead of `synchrony.memory.max`.

See Configuring System Properties to find out how to change this.

For Confluence Data Center the way you run Synchrony is a little different. See Configuring Synchrony for more information.

Multiple Synchrony processes

If you see an error immediatley in the editor, but Confluence reports that Synchrony is running, check to make sure that you only have one Synchrony process running.

If you do have multiple Synchrony processes running, stop Confluence, kill the additional Synchrony processes and then restart Confluence.

You can avoid this problem by always using `stop-confluence.sh` / `stop-confluence.bat` to stop Confluence, rather than simply closing the Tomcat window.

Mixed Synchrony modes in cluster

If you're running Confluence in a cluster, all of your Confluence nodes must connect to Synchrony in the same way.

If users are able to use collaborative editing on one Confluence node, but not on another Confluence node, go to > **General Configuration** > **Clustering**, then on each node choose **Collaborative editing**.

Make sure all of your Confluence nodes are reporting the same Synchrony mode - either Managed by Confluence, or Standalone Synchrony cluster.

You can access all nodes in this way, you don't need to hit a specific node in your browser.

Incompatible browser extensions

Some third party browser extensions that interact with the editor, such as Grammarly,may not function correctly with collaborative editing. See (Archived) Confluence Collaborative Editing blocks Grammarly Extension to find out how to disable Grammarly for just your Confluence site.

Firewall or anti-virus interference

We've had a few reports of firewalls or anti-virus software blocking some requests to the server, resulting in unexpected behavior in the editor. You may need to add Confluence to your whitelist / trusted URLs if you experience issues. See Weird Page or Editor Behaviors with Kaspersky Internet Security for more information.

Too many people in the editor

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

A maximum of 12 people can edit a page at the same time. This means that people can't enter the editor if there are already 12 other people editing the page, and will need to wait until someone leaves.

Administrators can increase or decrease this limit using a the`confluence.collab.edit.user.limit`system property.

Create feedback reports to troubleshoot content problems

If you experience problems such as data duplication, our support team may ask you to create a feedback report to help us troubleshoot the problem.

When someone creates a feedback report, we collect all the changes that have happened to that page or blog post, including the history, any page events, and Synchrony data from the editor (for the past 72 hours). This helps us construct a complete snapshot of the lifecycle of the page.

⚠ Note: Synchrony captures every key stroke made in the editor when editing a page. For example if you type "The best ice cream is Vanilla" and then backspace and change it to "Chocolate", both "Vanilla" and "Chocolate" will be included in the Synchrony data.

The user reporting the problem will also be prompted to describe the issue in their own words.

This data is exported as a zip file, and saved to the following directory:

- `<shared-home>/collab-data/`if you run Confluence in a cluster
- `<local-home>/shared-home/collab-data/`you are using non-clustered Confluence.

Because the collected data includes user generated content, generated reports are only kept for 5 days, then automatically deleted by the Clean up collaborative editing feedback reports scheduled job. A maximum of 200 reports are kept at any one time. You can change both of these values using a system property. See Recognized System Properties.

Allow users to self report problems

By default, only people with system administrator or Confluence administrator global permissions can create a feedback report report problems through the administrative interface. You can choose to allow any user with**Add page** space permission to create a report.

To allow users to self-report problems:

1. Go to > **General Configuration** >**Collaborative editing feedback**.
2. Under **Who can create feedback reports**, select the **Anyone with edit permissions** radio button.

This will add the**Report editing problems** item to the help menu in the header, when viewing or editing a page.

Report an editing problem

To report an editing problem from a page or blog post:

1. Navigate to the problematic page or blog post.
2. Select **Help** > **Report editing problems** from the header.
3. Enter any relevant information, such as what you were doing when the problem occurred.
4. Select Send feedback.

The report will be saved to the shared home directory, and will appear in the list of available reports (only accessible to system administrators).

Report an editing problem by page ID

System administrators can also generate reports for problematic pages using the page ID. This is useful if you don't want to allow users to self-report problems.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

To generate a report for a particular page or blog post:

1. Go to > **General Configuration**> **Collaborative editing feedback**.
2. Enter the page or blogpost ID in the **Content ID** field.
3. Select **Create report** button.

The report file will appear in the list of available reports.

Send a report to Atlassian support

If you choose to share this data with our Support team for analysis, you can transfer the files through a support ticket.

Feedback reports should **ONLY** be provided when requested by a Support team member. Don't share this data in any public locations, such as in Atlassian bug reports.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Using read-only mode for site maintenance

This feature is available with a Confluence Data Center license.

If you need to perform maintenance while Confluence is still running, or if you're preparing to migrate to a new site, you can put your site into read-only mode to limit what users can do. Your users will be able to view pages, but not create or change them.

Turn on read-only mode

You need System Administrator global permissions to do this.

To enable read-only mode:

1. Go to > **General Configuration**>**Maintenance**
2. In the Read-only mode section, choose**Edit**.
3. Select**Read-only mode**.
4. Update the wording of the**banner message**, if you'd like to provide a customised message.
5. Choose**Save**.

The banner message will display above the header on all pages in your site. It's not possible to disable this banner while read-only mode is enabled, but you can customise the message, for example to let your users know when you expect the maintenance to be complete.

It's also possible to turn on the banner*before* you enable read-only mode. This can be helpful if you want to warn users that you'll be doing some maintenance later that day.

Impact of read-only mode on your site and database

Read-only mode limits the actions that an end user can perform. Some operations may still write to your database, but for the most part people will be unable to make any changes.

While read-only mode is on, youwon't be able to:

- Create, edit, rename, move, delete or otherwise interact with pages.
- Create, delete or rename spaces.
- Access most space tools, including reorder pages, make changes to the look and feel, or add integrations.

Here's how a page looks when read-only mode is enabled:

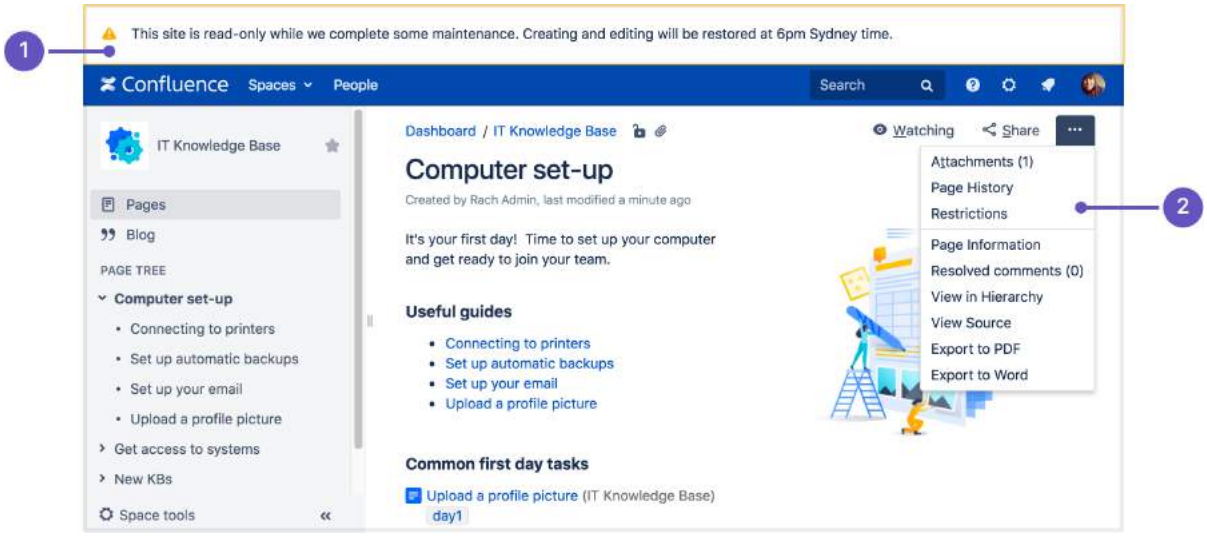

1. **Customizable banner** - the banner appears on all pages in your site. Admins can customize the message to let you know when the site will be available again.
2. **Options are limited** - we hide buttons and menu items that are not available, including create, edit, move, and delete.

If you happen to be in the editor at the point read-only mode is enabled, you'll be able to keep typing, but any further changes won't be saved.

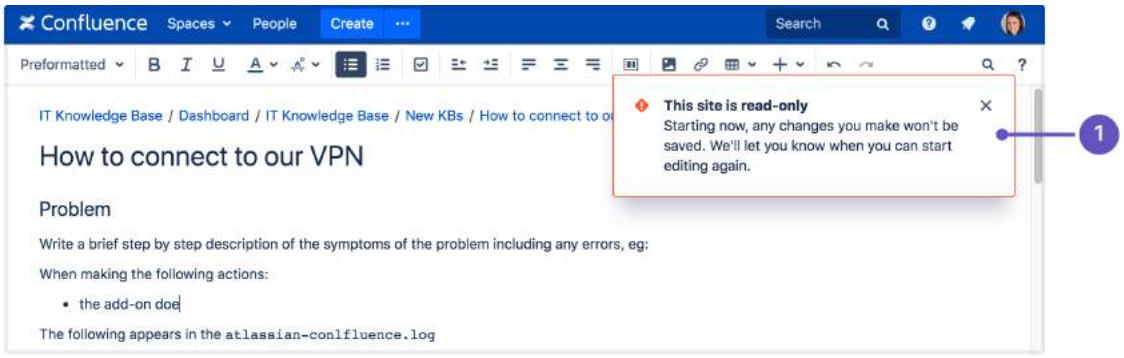

1. **Read-only warning** - although you can keep typing in the editor (including comment fields), changes you make after read-only mode is enabled won't be saved. It's best to stop editing at this point.

While read-only mode is on, people with **system administrator** global permissions will be able to perform some administrative functions, such as:

- Install, uninstall, enable, disable system and user installed apps
- Manage users, groups, and permissions
- Change the site appearance
- Export and import spaces
- Change logging levels, and other configuration.

Not all admin features will be available, and just like end-users, admins won't be able to create, edit, or delete any content.

While on read-only mode, people with **Confluenceadministrator**global permissions will also be able to perform some administrative functions, but they won't be able to make changes to space permissions.

Confluence's scheduled jobs will continue to run normally in read-only mode.

It's important to note that read-only mode**does not prevent data from being written to the database**, but will significantly limit the changes that can be made.

If you're doing database maintenance, and need to make sure that*absolutely nothing*is written to the database during that time, it may be best to stop Confluence, rather than using read-only mode.

User-installed app compatibility

Not all apps (also known as plugins or add-ons) are compatible with read-only mode, and may continue to allow users to create or update content while read-only mode is enabled.

To check if your apps are compatible:

- Go to > **General Configuration**>**Maintenance**
- Check whether any of your user-installed apps are listed as incompatible.

If an app is incompatible, you may want to disable it while you perform maintenance, to avoid users being able to create content via the app.

If you've developed your own custom apps, seeHow to make your app compatible with read-only modeto find out how to test your app and mark it as compatible.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Ways you might use read-only mode

If you're excited by the possibilities of read-only mode, but not sure when you might use it, here are some examples.

Upgrading Confluence

The way you upgrade Confluence hasn't changed, but read-only mode can help you minimize the impact on your organization.

If some downtime is acceptable, the simplest option is to enable read-only mode while you perform the pre-upgrade steps, such as checking Marketplace app compatibility and backing up your file system and database (if your database supports online backups). This helps you keep the overall downtime to a bare minimum, as users can view pages right up to the point you need to stop Confluence.

If you need to provide uninterrupted access, the approach you take may depend on whether Confluence is running on virtualized or physical hardware.

- If virtualized, you might want to take a 'move forwards' approach. You could enable read-only on your production site, clone your database, install, and home directories, then upgrade the clone. Once the upgrade is complete and you've validated that everything is working fine, you candirect traffic to the upgraded site, and tear down the old site.
- If you're running Confluence on physical hardware it might be more appropriate to create a temporary read-only site. You couldclone your production database, install, and home directories to create a temporary read-only site (similar to the process involved in creating a staging site), and direct traffic to that sitewhile you upgrade your production Confluence site in place.

You should also always test the upgrade on a staging or test instancefirst.As when creating astaging site, it's essential to make sure Confluence is always pointing to the correct database and home directory.

Upgrading your infrastructure

Need to move Confluence to another server, or provision more space for your shared home directory? The approaches outlined above for upgrading Confluence can also be useful when upgrading parts of your infrastructure.

Note that some data may still be written to the database while read-only mode is enabled, so if you're doing database maintenance of any sort, directing your users to a secondary site (with a copy of your database) that has read-only enabled, may be a good approach. You can't, for example, upgrade your production database while Confluence is still running, even if read-only mode is enabled.

Again, always make sure Confluence is pointing to the correct database!

Consolidating multiple confluence sites

It's quite common for multiple Confluence sites to pop up in big organisations. If you're consolidating or merging sites, read-only mode can help limit changes to content while you work through the process of exporting spaces and importing them into your new site.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Administering the Atlassian Companion App

This page applies to Confluence Server and Data Center

The Companion app was removed from Confluence Cloud at the end of March 2022. Read announcement

We have no plans to end support for the Companion app in Confluence Data Center and Server.

On this page:

- Download and install the Atlassian Companion app
- Compatibility with virtual desktop environments
- Recover edited files
- How to delete the Cache folder
- Disable file editing
- Alternatives to the Atlassian Companion app

The Atlassian Companion app enables users to edit Confluence files in their preferred desktop application, then save the file back to Confluence automatically.

The download and re-upload of files is managed by the Atlassian Companion app, which needs to be installed on each user's machine (not in the Confluence installation directory) to enable file editing.

Download and install the Atlassian Companion app

To edit files, users need to install the Atlassian Companion app and have it running in the background. The first time a user clicks the Edit button in file preview, we prompt them to download and install the app. SeeEdit Filesfor details.

If your users aren't able to install applications themselves, you may want to distribute the app to them or deploy using the Microsoft Installer.

Download the latest Companion version

Download the Atlassian Companion app for Mac or Windows.

Single sign-on considerations

If you've configured single sign-on (SSO) in such a way that your reverse proxy redirects the requests to your SSO gateway, and only successfully authenticated requests ever reach Confluence, your users won't be able to edit files using the Atlassian Companion app.This is because the Atlassian Companion app uses JWT tokens to authenticate requests, and only Confluence can authenticate these requests, not your SSO authenticator.

To make sure requests from the Atlassian Companion app can be authenticated, you should configure your reverse proxy to always allow requests from the following URLs:

- `<base-url>/rest/token-auth/api/*`
- `<base-url>/download/token-auth/attachments/*`
- `<base-url>/plugins/servlet/imgFilter*`
- `<base-url>/rest/analytics/1.0/publish/bulk` (only necessary if you haveopted in to data collection)

If an unauthenticated user tries to access these URLs directly, they would be redirected to the Confluence login screen. The wouldn't be able to access any content or download files while unauthenticated.

There's a known issue where the `token-auth` path is not included in the download URL that Confluence provides Companion. See

~~CONFSERVER-63189~~ - Companion App does not perform download request using token-auth endpoint
CLOSED

Content security policy (CSP) considerations

If you have a restrictive content security policy, your browser will refuse to launch companion, and you'll see a content security policy error in the browser console. This error occurs because Confluence 7.3 and later uses a hidden iframe to attempt to launch Companion's custom protocol (`atlassian-companion`). To resolve this problem you will need to add `atlassian-companion:` to the `default-src` or `frame-src` list. For example:

```
frame-src atlassian-companion:;
```

The content security policy is most commonly configured in your reverse proxy.

Install the Companion app via Microsoft Installer (MSI)

We also provide a Microsoft Installer package (.msi file) to deploy the Atlassian Companion app for Windows across multiple users or machines. By default, the Companion app installs to the Program Files directory, but you can customize this.

Download the Atlassian Companion MSI (69 MB)

If the link above downloads an .exe file instead of the MSI, copy the URL below into your browser to download the file.

```
https://update-nucleus.atlassian.com/Atlassian-Companion/291cb34fe2296e5fb82b83a04704c9b4/latest/win32/ia32/Atlassian%20Companion.msi
```

Use the Microsoft Installer to install the Companion app for all users on a given computer:

```
msiexec /i Atlassian Companion.msi COMPANION_TRUSTED_DOMAINS=https://confluence.atlassian.com; https://support.atlassian.com; /qb ALLUSERS=1"
```

If you deploy using the Microsoft Installer, the Companion app wont automatically get the latest updates, including security and bug fixes, so some maintenance is required.

We may update the Companion app before or after we release a new version of Confluence. Check the Atlassian Companion app release notes to make sure you're on the latest version.

Standard install switches

The Microsoft Installer supports standard install switches. For example, install with the `TARGETDIR` and `APPLICATIONROOTDIRECTORY` parameters to change the installation directory:

```
msiexec /i "Atlassian Companion.msi" TARGETDIR="C:\Users\Emma\AppData\Local\Companion" APPLICATIONROOTDIRECTORY="C:\Users\Emma\AppData\Local\Companion" /qb
```

Set trusted domains

In Companion 1.2.0 and later, set your Confluence URL as a trusted domain so users dont have to select 'Trust this domain' when they edit a file for the first time.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

System administrators have two options for setting trusted domains/sites before rolling out the Companion app to all users. Either set an environment variablecalled `COMPANION_TRUSTED_DOMAINS`on each user's computer, or pass the parameter `COMPANION_TRUSTED_DOMAINS`to the Microsoft Installer (MSI).Set multiple trusted domains by using semicolons (;) as separators.

To set trusted domains when installing using the MSI:

```
msiexec /i "Atlassian Companion-1.1.0.msi" COMPANION_TRUSTED_DOMAINS=https://confluence.atlassian.com;
https://support.atlassian.com;"
```

Compatibility with virtual desktop environments

From Confluence 7.3 onwards, Atlassian Companion app should work in most session-based virtual desktops.

Recover edited files

When a user edits a file, that file is also downloaded and saved to the Atlassian Companion folder on their computer. Files modified more than 60 days ago are automatically cleared when the Companion app restarts.

Follow our guide toaccessing Confluence files edited with the Atlassian Companion app.

How to delete the Cache folder

If youd like to free up disk space, its safe to manually delete the cache folder. Deleting individual files in the cache folder may cause errors, so you should delete the entire Cache folder. If the cache folder is locked while Companion App is running, quit Companion, delete the cache folder, then open Companion.

For Windows, Companion 1.0.0, go to C:\Users\admin\AppData\Roaming\Atlassian Companion\Cache.

For Windows, Companion 1.1.0, go to C:\Users\admin\AppData\Local\Atlassian Companion\Cache.

For Mac, Companion 1.0.0 and 1.1.0, go toHome/Library/Application Support/Atlassian Companion /Cache.

Disable file editing

From Confluence 7.3 onwards, it is not possible to disable file editing completely. However, you can choose to revert to the previous Edit in Office functionality, which disables the Companion app integration.

Alternatives to the Atlassian Companion app

In some versions of Confluence, you can revert to the previousEdit in Officefunctionality. This is a workaround for customers who are unable to use the Companion app in their environment.

To enable the legacy Edit in Office functionality:

1. Go to > **General Configuration**>**Office Connector**.
2. Choose**Enable Edit in Office for all users**and save your changes.

This will disable Companion app functionality for all users in the site.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Notifications from Atlassian

The Atlassian Notifications system app provides targeted in-app notifications in Confluence.

These notifications are mostly directed at administrators and provide information about things like new Confluence versions, upcoming license renewals, and important security announcements.

To disable these notifications:

1. Go to > **Manage apps**.
2. Search for **Atlassian Notifications**.
3. Expand the listing choose **Disable**.

Administer analytics

This feature is available with a Confluence Data Center license.

If you have Confluence Administrator or System Administrator global permissions you can configure Confluence analytics to suit the needs of your organisation.

To find out what data is collected, seeAnalytics.

On this page:

- Disable analytics
- Permissions
- Data retention
- Rate limiting
- Increased privacy mode
- Known issues
- There are some known issues with analytics that you should be aware of.
- Analytics for Confluence marketplace app

This page is not about the analytics data which is sent to Atlassian to help us improve the product. See Data Collection Policy to find out what is collected, and how you can opt in or out of this data collection.

Disable analytics

If you don't want Confluence to collect user activity and page view data, you can disable the system app that provides the Analytics feature.

To disable analytics:

1. Go to > **Manage apps**.
2. Search for**Analytics for Confluence**.
3. Expand the listing then select**Disable**.

Disabling the system app won't remove any existing analytics data from your database, but will stop Confluence collecting data.

Permissions

By default, all logged in users can view the Analytics option in the header, and access analytics data for spaces they have permissions to see.Anonymous users can't view analytics.

Limit who can view analytics reports

If you don't want everyone to be able to view analytics, you can chose to limit access to specific groups.

To limit analytics tospecific groups:

1. Go to> **General Configuration**>**Analytics Permissions**.
2. Search for a group.
3. Select **Add**.
4. Repeat this process for each group.

Only people who are a member of these groups will see the Analytics option in the header, or in a space.

To revert to the default, and allow all users to view analytics:

1. Go to> **General Configuration**>**Analytics Permissions**.

2. Remove all groups from the list.

You don't need to add every group individually to give everyone access.

Limit who can view analytics reports for specific spaces

You can further limit who can view analytics reports for specific spaces. You need Space Administrator permission to do this.

To change who can see analytics reports in a space:

1. Go to the space and choose **Space tools** > **Permissions** from the bottom of the sidebar.
2. Select the **Analytics Permissions** tab.
3. Select **Viewing Analytics Restricted** from the drop down.
4. Enter the users and/or groups you want to allow to view analytics reports.
5. Save your changes.

Good to know:

- If a Confluence administrator has denied a group permission to view analytics, adding the group at the space level will not grant this permission.
- The space will still appear in the site analytics report (if the user has permission to see the space, and use analytics globally), but they will be prevented from viewing the space analytics report if they don't have space permissions.

Data retention

If you have a very big, busy site, the amount of data collected can have an impact on your database size and general site performance. To avoid any problems, you can choose how long to retain events for, and change the maximum number of events that can be stored at any time.

If you have a very large, active site, we recommend keeping the data retention settings quite low, and using rate limiting.

Set a data retention period

By default we store analytics events in the database for 12 months. You can reduce or increase this limit if you have different requirements. A long retention period can affect the size and performance of your database, so if you have a big, busy site you may need to choose a shorter retention period.

To change the data retention period:

1. Go to > **Manage apps**.
2. Search for **Analytics for Confluence.**
3. Expand the listing then select **Configure**.
4. Under **Data Retention Period**, select **Edit**.
5. Choose either 6, 12, 18, or 24 months.
6. Save your changes.

Your change will take effect after 24 hours. This is to prevent data being immediately deleted if you accidentally choose the wrong retention period.

Confluence regularly deletes any events older than the chosen retention period (every 60 seconds). Note that the event retention limit also applies, so events exceeding the maximum number of events will be deleted, regardless of whether they fall within the data retention period.

Set an event retention limit

By default, we store a maximum of 20,000,000 analytics events in the database. This is to ensure analytics queries are reasonably performant in all databases. You can choose to further reduce this limit.

To change the maximum number of events to retain:

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

1. Go to > **Manage apps**.
2. Search for**Analytics for Confluence.**
3. Expand the listing then select**Configure**.
4. Under**Event retention**, select**Edit**.
5. Enter the maximum number of events.
6. Save your changes.

Confluence regularly deletes any events exceeding limit (every 60 seconds), starting with the oldest events.

We don't recommend increasing this limit beyond 20 million events, as this will result in very large database tables, and eventually cause performance issues.

Rate limiting

Too many people accessing analytics data at the same time can slow down your site. To avoid performance bottlenecks:

- You can limit the number of **concurrent sessions** across your instance. This helps you control the number of analytics sessions that are open across your site.
- You can limit the number of **concurrent reports per session**. This allows you to set the number of reports that can be generated at the same time in a session.
- You can set a **timeout for reports** that are taking too long to load. This helps you manage overall site performance.

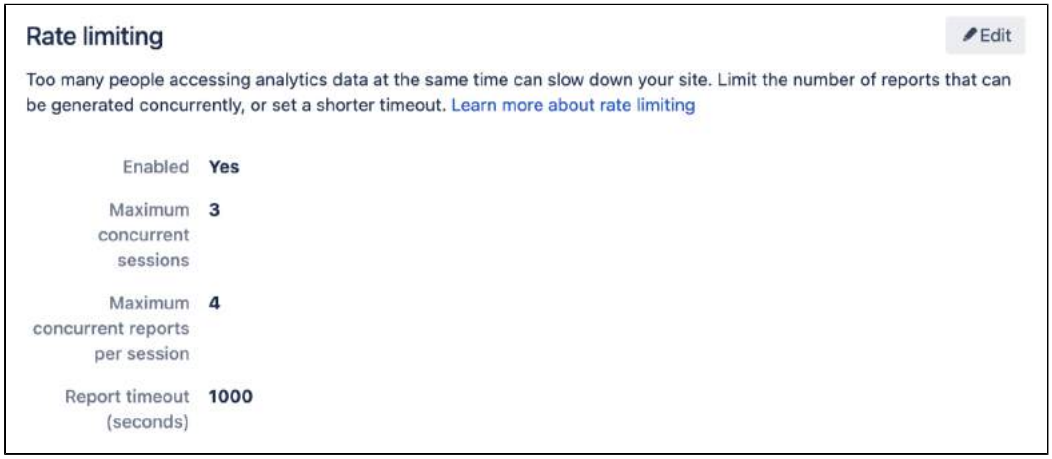

Screenshot: Configure rate limiting for analytics in the admin console.

To change the maximum number of concurrent sessions, or concurrent reports per session:

1. Go to > **General Configuration**.
2. In the menu items (left panel), find the **Confluence Analytics** section,then select **Configuration**.
3. Next to **Rate limiting,** select **Edit**.
4. Enter the maximum number of current sessions.
5. Enter the maximum number of concurrent reports per session.
6. Select **Save**.

If someone tries to view an analytics report, and the limit has been reached, they'll see a message to try again in a few minutes.

To change the timeout:

1. Go to > **General Configuration**.
2. In the menu items (left panel), find the **Confluence Analytics** section,then select **Configuration**.
3. Next to **Rate limiting,** select **Edit**.
4. Enter the new timeout value in seconds.
5. Select **Save**.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

If someone tries to view a report that can't be loaded within the time limit, they'll see a timeout error. When this happens, the best option is to change the report filters, such as reducing the date range.

Increased privacy mode

Privacy requirements can differ greatly between organisations and between different jurisdictions. In some circumstances it will not be appropriate for people to see the names of users who have viewed content in your site, or for you to be collecting this data in an identifiable format.

Turn on Increased privacy mode to minimise the amount of personally identifiable information (PII) that is collected and displayed to users.

To turn on increased privacy mode:

1. Go to > **Manage apps**.
2. Search for**Analytics for Confluence.**
3. Expand the listing then select**Configure**.
4. Under**Increased Privacy Mode**, select**Edit**.
5. Select the **Enabled** toggle.
6. Save your changes.

From this point onwards, space or page activity is no longer linked to individuals, but instead attributed to an anonymised user. In analytics reports, people will be represented as "User 12345" with an anonymised avatar. This means you still get an accurate picture of the engagement with your content, but without revealing user information.

It is very important to note that changing this setting does not affect any previously collected data. This means:

- You may need to manually remove any previously collected data after turning on increased privacy mode.
- Unique user counts will be temporarily inaccurate, as we do not attempt to connect a named user with their anonymised alias, so unless you have manually removed the named user data, both will exist for a time.

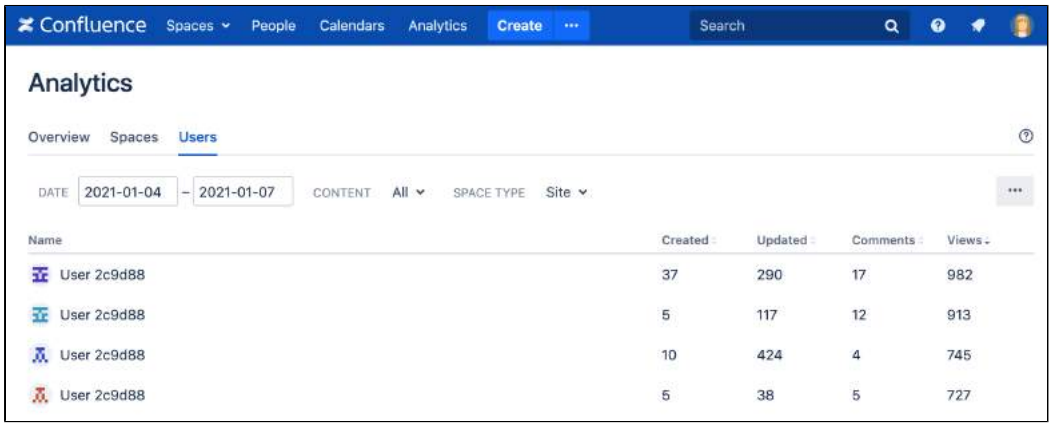

Screenshot: the site analytics users tab, with increase privacy mode turned on.

Known issues

There are some known issues with analytics that you should be aware of.

Site analytics reports may take a long time to load

There are a few situations where the site analytics reports (accessed from Analytics in the header) take longer than expected to load, including:

- if your spaces have complex permissions, or most spaces in your site are restricted
- if you're running Confluence on MySQL.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Some reports show data for deleted content

Some analytics reports continue to show aggregate data for content or user accounts that have subsequently been deleted.

The following reports continue to show views data for pages that have been deleted:

- Site analytics report - spaces tab
- Site analytics report - users tab
- Space analytics report - users tab

The following reports continue to show views data for user accounts that have been deleted:

- Site analytics report - spaces tab
- Space analytics report - content tab

Comments column always shows data for pages and blogs

In the Users tab of any analytics report, the comments column always lists the total number of comments on pages and blog posts, regardless of whether the content filter is set to just pages, or just blog posts.

Analytics for Confluence marketplace app

The Good Software Analytics for Confluence was previously available on Marketplace. We ended new sales and renewals for this app in 2019, and subsequently made the app's features part of Confluence Data Center (version 7.11 and later).

If you have a Confluence Server license, you can continue to use the app until it reaches end of life, but features in your legacy version will differ from features available in Confluence Data Center.

If you have a Confluence Data Center license, the Good Software Analytics for Confluence marketplace app will be automatically disabled when you upgrade to Confluence Data Center 7.11 or later. Previously collected analytics data will be preserved.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Monitor application performance

App monitoring can give you a deeper insight into what apps are doing in your instance. This can be useful when troubleshooting issues with a specific app, or to help you determine whether an app may have contributed to a drop in overall performance or stability.

On this page:

- Set up monitoring
- Identify the app name
- Enable optional tags
- Troubleshooting
- Next steps

Set up monitoring

Before you can connect your APM to Confluence, you need to:

- configure a JMX exporter, and
- make sure that JMX and App metrics are enabled in your site.

The instructions on this page assume youll be using Prometheus. You can use any Application Performance Monitoring (APM) solution, the steps will be very similar for each.

Configure a JMX Exporter

The exporter takes the JMX MBeans and transforms them into the right format for Prometheus. It also hosts a HTTP endpoint which Prometheus will connect to. Learn more about the Prometheus JMX exporter

If you don't plan to use Prometheus, you'll need to check which exporter or agent is required for your APM solution. For example, this Java agent for NewRelic.

To install the exporter:

1. Download the Prometheus JMX exporter jar file from theGitHub repository.

   ```
   $ curl -L https://repo1.maven.org/maven2/io/prometheus/jmx/jmx_prometheus_javaagent/0.16.1
   /jmx_prometheus_javaagent-0.16.1.jar > jmx-exporter.jar
   ```

2. Create a configuration file for the java exporter. We recommend you use the jmx-exporter-config.yml configuration fileprovided in our repository.
3. Copy the jar file and configuration file to each application node (the local home directory is a good option).
4. Stop Confluence on one node.
5. Add the following system properties to tell Confluence where to find the JMX exporter.See Configuring System Properties to check how to do this for your site.

   ```
   -javaagent:<full-path-to-jmx-exporter-jar>=<port>:<full-path-to-jmx-exporter-config.yml>
   ```

   The JMX exporter defaults to port 8080. You'll need to specify a different port for the exporter if 8080 is in use by another application.
6. Start Confluence.
7. To check that the exporter is working, go to `localhost:<jmx-exporter-port>`. You should see the metrics output.

Repeat these steps for all remaining nodes, if you run Confluence in a cluster. You can perform a rolling restart to avoid any downtime.

Make sure the JMX exporter endpoint is not exposed outside your network, or take appropriate steps to secure it.

Check that app monitoring is enabled

Application monitoring uses JMX (Java Management Extensions), so you'll need to check that both JMX monitoring and App monitoring are enabled. These are both enabled by default.

To confirmapp monitoring is enabled:

1. Go to > **General Configuration**> **Monitoring**.
2. Check that **JMX monitoring** is enabled.
3. Check that**App monitoring** is enabled.

If you have previously set up JMX monitoring for Confluence, theres nothing else you need to do. The additional application monitoring metrics will be exposed in the same way as existing application metrics.

Identify the app name

App metrics include the plugin key rather that the app's display name. For example, `com.atlassian.troubleshooting.plugin-confluence` is the plugin key for the Troubleshooting and Support Tools system app for Confluence.

To find the app name:

1. Go to `<base-url>/plugins/servlet/upm/osgi`
2. Enter the plugin key in the **Search bundled metadata** field
3. The plugin details will be returned, including the name and vendor.

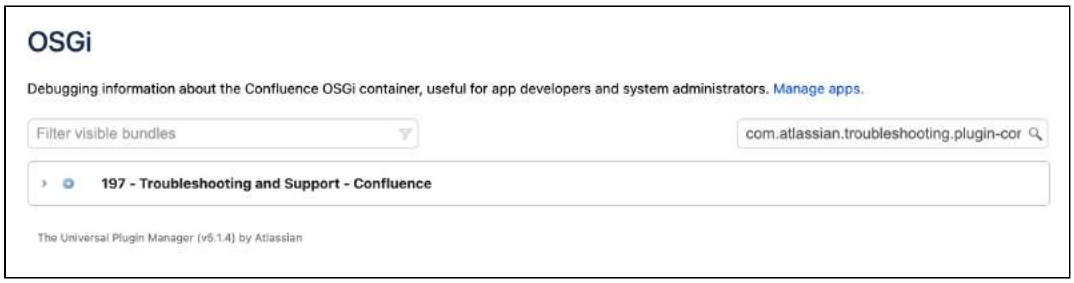

OSGi admin screen showing search results for a plugin key

Enable optional tags

App vendors can choose to include additional metadata which can help when troubleshooting a performance issues. These tags are not included by default.

You can use the`atlassian.metrics.optional.tags` system property to show additional tags for a metric.

```
atlassian.metrics.optional.tags.<metric-name>=<tag-key1>,<tag-key2>
```

For example, if the full metric name is `sampleApp.asset.loadtime` and the app vendor included a tag to output additional information about the content type.

```
atlassian.metrics.optional.tags.sampleApp.asset.loadtime=sampleApp-type
```

The app vendor will be able to tell you the exact metric and tag names.

Disable app monitoring

To disable app monitoring:

1. Go to > **General Configuration**> **Monitoring**.
2. Disable**App monitoring**.

Once disabled, Confluence will no longer emit app-specific metrics, or write them to logs. If you want to disable JMX altogether, you can also disable **JMX monitoring**.

Troubleshooting

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

JMX disabled via system property

JMX is enabled in Confluence by default, and previously could only be disabled using a system property. Youll see a warning on the Monitoring page if the `confluence.jmx.disabled` property is set on any of your nodes.

You wont be able to toggle JMX monitoring on or off through the Monitoring screen until you have removed the system property.

Out of memory errors

Because the monitoring is happening outside your application, we dont expect there to be a significant impact on your instance performance or stability.

In the event you do notice increased memory usage, or out of memory errors (OOME) caused by the monitoring agent, you may want to increase the minimum heap size (Xms) in the setenv file. SeeHow to fix out of memory errors by increasing available memory.

Next steps

Next, configure your APM tool to point to the JMX exporter endpoint. If you don't have an APM, check out guide to setting up Prometheus and Grafana.

SeeApp metrics reference for a full list of app metrics, and recommended alerts.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Monitor Confluence with Prometheus and Grafana

This page will guide you through how to install and connect Prometheus and Grafana. This is optional, but may be useful if you don't already have an APM, or would like to use our templates and sample queries.

Use Prometheus to monitor app performance metrics

To set up Prometheus to monitor app metrics:

1. Download and install Prometheus.
   For installation options and detailed instructions see the Prometheus documentation.
2. Edit the prometheus.yaml file and add the following scrape configuration to the bottom of the file.

```
# A scrape configuration containing exactly one endpoint to scrape:
scrape_configs:
  - job_name: 'Confluence app metrics'
    scheme: http
    metrics_path: '/metrics'
    static_configs:
     - targets: ["<jmx-exporter-host>:<port>"]
```

   - The target is the JMX exporter, not Jira. For example - targets: ["localhost:8060"]
   - If you deploy Prometheus in Kubernetes, you'll need to use a pipe to indicate the multi-line YAML string, as in the example below.

```
extraScrapeConfigs: |
  - job_name: 'Confluence app metrics'
    scheme: http
    metrics_path: '/metrics'
    static_configs:
     - targets: ["10.23.45.678:8080"]
```

   - See Configuration in the Prometheus documentation for more configuration options.
3. Start Prometheus. How you do this will depend on the way you run Prometheus.
4. Access the Prometheus UI at `http://localhost:9090/`.
5. Go to **Status** > **Targets** to check that Prometheus is successfully connected to the JMX exporter.

Perform a simple query

You can confirm that Prometheus is receiving app metrics with a simple test.

Go to > **Manage apps**and temporarily disable an app (such as the Confluence Migration Assistant, don't disable anything that will interrupt your users).

In Prometheus, run the following query:

```
com_atlassian_confluence_metrics_Count
  {
   category00="plugin",
   category01="disabled"
  }
```

This will return the number of times an app has been disabled since monitoring was turned on.

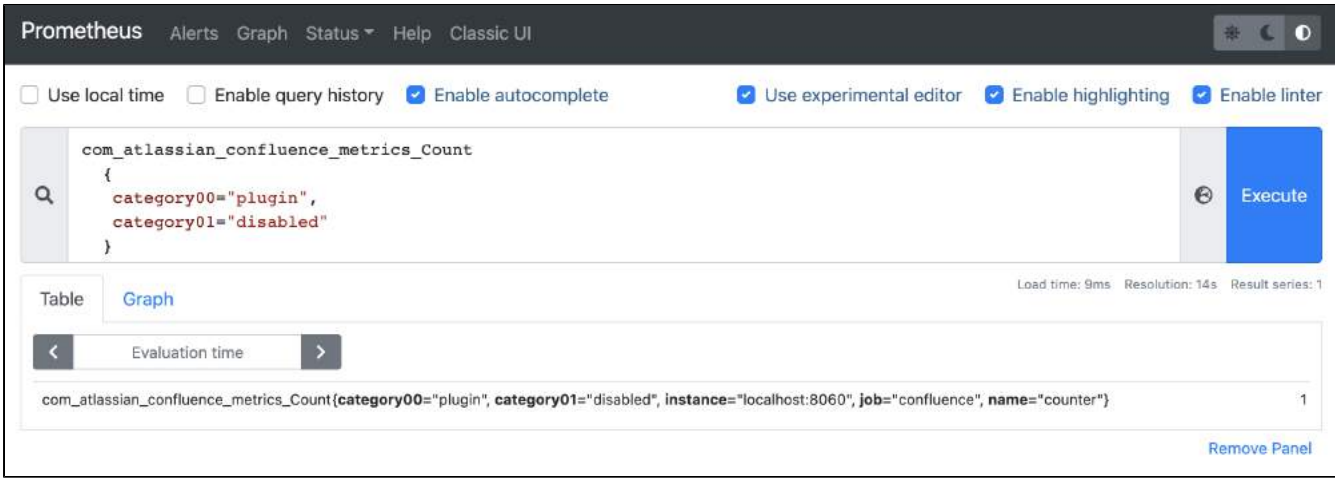

Use Grafana to visualize metrics

While you can use Prometheus to create graphs of your data, if you want to take it to the next level, you can use a tool like Grafana to create more detailed charts and dashboards.

To get you started, we've created some sample dashboards which tracks several important metrics. You can access the JSON for these dashboards in our App monitoring dashboards repository.

To set up Grafana and import the sample dashboard:

1. Download and install Grafana.
   For installation options and detailed instructions see the Grafana documentation.
2. Create a Prometheus data source in Grafana.
   For detailed instructions see the Prometheus documentation.
3. Select **Create (+)** > **Import**.
4. Paste the JSON sample provided in the repository into the **Import via panel json** field. Remember to update the Unique identifier (only required if you already have a dashboard with the same ID).
5. Select **Load**.

Here's an example of a dashboard in Grafana showing response times for various plugins.

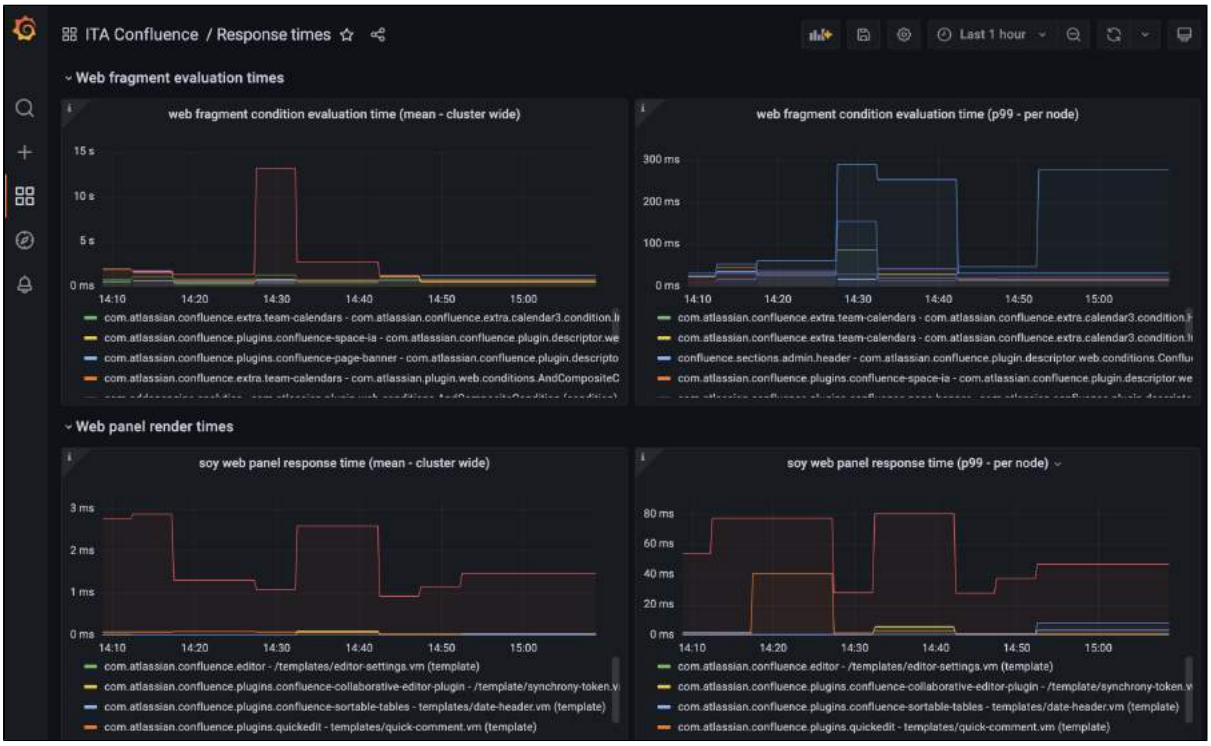

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

App metrics reference

On this page:

- Full list of app performance metrics
- Recommended alerts

App monitoring can give you a deeper insight into what apps are doing in your instance. This can be useful when troubleshooting issues with a specific app, or to help you determine whether an app may have contributed to a drop in overall performance or stability.

Learn how to set up app monitoring

Full list of app performance metrics

This is the full list of metrics that are exposed by the app monitoring agent. This is in addition to any JMX beans that are exposed by the application.

`index.reindex`

This metric indicates that information has been reindexed.

The metric consists of a number of tags:

- If `indexContent` is true, content was reindexed such as page content
- If `indexAttachments` is true, attachments were reindexed
- if `indexUsers` is true, users were reindexed
- If `limitedWithQuery` is false then everything was re-indexed.
- If `active` is 1, a reindex is in progress, and the duration will indicate how long it has been running for.

Action

Reindexing can degrade your site's performance. Ideally, you would reindex during off-peak times.

The `invokerPluginKey` will indicate which app kicked off the reindexing. If the key starts with `com.atlassian` then its likely to be something within Confluence.

Sample query

```
com_atlassian_confluence_metrics_99thPercentile
  {
   category00="index",
   name="reindex"
  }
```

`search.manager`

Measures how long a search request takes.

- `methodName` indicates the api invoked
  - `search`
  - `searchWithToken`
  - `searchWithRequestedFields`
  - `searchEntities`
  - `explain`
  - `searchCategorised`
  - `convertToEntities`
  - `scanWithSearchFilter`
  - `scanWithSearchQuery`
  - `scanWithIndexesAndSearchQuery`
  - The metrics with method name as one of [`scanWithSearchFilter`, `scanWithSearchQuery`, `scanWithIndexesAndSearchQuery`] are expected to take longer time.
- `usedFilter` indicates if a filter is used.
- `searchType` represents the type of search, can be one of [SiteSearch, ContentSearch, CQLSearch].
- `resultSize` the number of documents matches the search, only applicable when `methodName=convertToEntities`

Action

Use the `pluginkey` to identify which app is calling the search API (`com.atlassian.confluence.search.v2.SearchManager`).

If you notice an app is making a lot of searches, or consistently takes a long time to process search results, reach out to the app vendor.

Sample query

```
com_atlassian_confluence_metrics_99thPercentile
  {
   category00="search",
   name="manager"
  }
```

`db.ao.upgradetask`

Measures how long an app is taking to upgrade a part of the data it stores in the database.

Upgrade tasks can happen when an app is updated or enabled. During this time the app functionality will be unavailable, and may temporarily increase load on the database and the node the upgrade task is running on.

Action

If an app stores a lot of data in database consider scheduling any updates when Confluence is less busy.

Samplequery

```
com_atlassian_confluence_metrics_Value
  {
   category00="db",
   category01="ao",
   name="upgradetask"
  }
```

`db.ao.executeInTransaction`

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Measures how long an Active Objects (AO) transaction takes when executed inside the `TransactionCallBack`. This is mainly used by Confluence plugins.

Action

The transaction can have many AO operations. The problem may be that there are too many operations, the query is long running, or the database is under load.

Sample query

```
com_atlassian_confluence_metrics_Value
  {
   category00="db",
   category01="ao",
   name="executeInTransaction"
  }
```

`db.ao.entityManager`

Measures how long an Active Objects (AO) operation (create, find, delete, deleteWithSQL, get, stream, count) that uses the `entityManager` takes.

Action

The operation query may be long running, or the database is under load.

Sample query

```
com_atlassian_confluence_metrics_95thPercentile
  {
   category00="db",
   category01="ao",
   category02="entityManager"
  }
```

Can be filtered further by adding a `name="<operation>"` attribute, for example `name="find"`.

`db.cluster.lock.held.duration`

Measures how long a database cluster lock was held. Used by Confluence in a clustered environment.

Action

Lock contention can lead to performance degradation. It may be normal for a thread to hold on to a lock for a long time, if there aren't any threads waiting for the lock.

See `db.cluster.lock.waited.duration` to find out if there are any threads waiting for the lock.

Sample query

```
com_atlassian_confluence_metrics_Value
  {
   category00="cluster",
   category01="lock",
   category02="held"
  }
```

`db.cluster.lock.waited.duration`

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Measures how long a database cluster lock was waited for. Used by Confluence in a clustered environment.

Action

If many threads are waiting for the same lock, it can lead to performance degradation.

Sample query

```
com_atlassian_confluence_metrics_Value
  {
   category00="cluster",
   category01="lock",
   category02="waited"
  }
```

`db.sal.transactionalExecutor`

Measures how long a Shared Application Layer (SAL) transaction takes, when executed inside the `DefaultTransactionalExecutor`.

Action

The transaction can have many SAL operations, it can be either there are too many operations or the query is long running, or the database is under load.

Sample query

```
com_atlassian_confluence_metrics_Value
  {
   category00="db",
   category01="sal",
   name="transactionalExecutor",
   statistic="active"
  }
```

`web.resource.condition`

Measures how long a web resource condition will take to determine whether a resource should be displayed or not.

Action

Slow web resource conditions can lead to slow page load times especially if they are not cached.

Sample query

```
com_atlassian_confluence_metrics_95thPercentile
  {
   category00="web",
   category01="resource",
   name="condition"
  }
```

`plugin.disabled.counter`

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Measures how many times an app was disabled since uptime.

Action

Some caches are cleared when an app is disabled or enabled. This can have performance impact. If this number increases, check UPM or the application logs to investigate which app is contributing to this number.

Sample query

```
com_atlassian_confluence_metrics_Count
  {
   category00="plugin",
   category01="disabled"
  }
```

`plugin.enabled.counter`

Measures how many times an app was enabled since uptime.

Action

Some caches are cleared when an app is disabled or enabled. This can have a performance impact. If this number increases, check UPM or the application logs to investigate which app is contributing to this number.

Sample query

```
com_atlassian_confluence_metrics_Count
  {
   category00="plugin",
   category01="enabled"
  }
```

`soyTemplateRenderer`

Measures how long a Soy Template web panel takes to render.

Action

The template renderer might be long running.

Sample query

```
com_atlassian_confluence_metrics_95thPercentile
  {
   name="webTemplateRenderer",
   templateRenderer="soy"
  }
```

`webTemplateRenderer`

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Measures how long an Atlassian Template web panel takes to render.

Action

The template renderer might be long running.

Sample query

```
com_atlassian_confluence_metrics_95thPercentile
  {
   name="webTemplateRenderer",
   templateRenderer="velocity"
  }
```

`web.fragment.condition`

Measures how long web fragment condition will take to determine whether a web fragment should be displayed or not.

Action

Web fragments conditions determine whether a link or a section on a page should be displayed. Slow web fragment conditions lead to slow page load times especially if they are not cached.

Sample query

```
com_atlassian_confluence_metrics_95thPercentile
  {
   category00="web",
   category01="fragment",
   name="condition"
  }
```

`cacheManager.flushAll`

Indicates that all caches are being flushed by an app. This operation should not be triggered by external apps and can lead to product slowdowns.

Action

Use the `invokerPluginKey` tag to determine which app invoked the flush.

Sample query

```
com_atlassian_confluence_metrics_Count
  {
    category00="cacheManager",
    name="flushAll"
  }
```

`cache.removeAll`

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Indicates that a single cache has had all of its entries removed. This may or may not cause slowdowns in products or apps.

Action

Check how often these cache removals occur, and from which product. Use the `pluginKeyAtCreation` tag to determine which app created the cache.

Sample query

```
com_atlassian_confluence_metrics_Count
  {
   category00="cache",
   name="removeAll",
   invokerPluginKey!="undefined"
  }
```

`cachedReference.reset`

Indicates that a single entry in a cache has been reset. This may or may not cause slowdowns in products or apps.

Action

Check how often these cache resets occur, and from which product. Use the `pluginKeyAtCreation` tag to determine which app created the cache.

Sample query

```
com_atlassian_confluence_metrics_Count
  {
   category00="cachedReference",
   name="reset",
   invokerPluginKey!="undefined"
  }
```

`rest.request`

Measures HTTP requests of the REST APIs that uses the atlassian-rest module.

Action

Check the frequency and duration of the rest requests.

Sample query

```
com_atlassian_confluence_metrics_95thPercentile
  {
   category00="http",
   category01="rest",
   name="request"
  }
```

Recommended alerts

Automated alerts help you identify issues early, without needing to wait for an end-user to bring problems to your attention. Most APM tools provide alerting capabilities.

The following alerts are based on our research into common issues with apps. We've used Prometheus and Grafana, but you may be able to adapt these rules for other APM tools.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

To find out how to set up alerting in Prometheus, see Alerting overview in the Prometheus documentation.

Heap memory usage

Excessive Heap memory consumption often leads to out of memory errors (OOME). While fluctuations in Heap memory consumption are expected and normal, a consistent increase or failure to release this memory, can lead to issues. We suggest creating an alert which is triggered when there is less than 10% free Heap memory left on a node for an amount of time, such as 2 minutes.

```
- alert: OutOfMemory
  expr: 100*(jvm_memory_bytes_used{area="heap"}/jvm_memory_bytes_max{area="heap"}) > 90
  for: 2m
  labels:
    severity: warning
  annotations:
    summary: Out of memory (instance {{ $labels.instance }})
    description: "Memory is filling up (< 10% left)"
```

CPU utilisation

Consistently high CPU usage can be caused by numerous issues such as process intensive jobs, inefficient code (loops), or too little memory.

We recommend creating an alert that is triggered when CPU load exceeds 80% for an amount of time, such as 2 minutes.

```
- alert: HighCpuLoad
  expr: (java_lang_OperatingSystem_ProcessCpuLoad * 1000 > 80
  for: 2m
  labels:
    severity: warning
  annotations:
    summary: High CPU load (instance {{ $labels.instance }})
    description: "CPU load is > 80%"
```

Full GC

Full garbage collection (GC) occurs when both young and old Heap generations are collected. This is time consuming and pauses the application. Full GC can happen for a number of reasons, but a sudden spike may happen when too many large objects are loaded into memory.

We recommend monitoring any significant increase in the number of full GCs. How you do this will vary depending on the type of Collector being used. For the G1 Garbage Collector (G1GC), monitor the `java_lang_G1_Old_Generation_CollectionCount` metric.

Blocked threads

A high number of blocked or stuck threads means there are fewer threads available to process requests. An increase in blocked threads could indicate a problem.

We recommend creating an alert that is triggered when the number of blocked threads exceeds 10%.

```
- alert: BlockedThreads
  expr: avg by(instance) (rate(jvm_threads_state{state="BLOCKED"}[5m])) * 100 > 10
  for: 0m
  labels:
    severity: warning
  annotations:
    summary: Blocked Threads (instance {{ $labels.instance }})
    description: "Blocked Threads are > 10%"
```

Database connection pool

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

The database connection pool should be tuned for the size of the instance (such as the number of users and plugins). It also needs to match what the database allows.

We recommend creating an alert that is triggered when the number of connections is consistently near the maximum for an amount of time.

Example alert:

```
  - alert: DatabaseConnections
    expr: 100*(<domain>_BasicDataSource_NumActive{connectionpool="connections"}/
<domain>_BasicDataSource_MaxTotal{connectionpool="connections"}) > 90
    for: 5m
    labels:
      severity: warning
    annotations:
      summary: Database Connections (instance {{ $labels.instance }})
      description: "Database Connections are filling up (< 10% left)"
```

Reacting to alerts

Some issues are transient, or may resolve themselves, while others could be a warning sign of a major performance degradation.

When investigating the source of the problem, the app specific metrics below can help. If it's clear from the metrics that one particular app is spending more time or calling an API more frequently, you could try disabling that app to see whether performance improves. If it's a critical app, raise a support ticket, and include any relevant data extracts from your monitoring with the support zip.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Live Monitoring Using the JMX Interface

JMX (Java Management Extensions API) allows you to monitor the status of your Confluence instance in real time. JMX uses objects called MBeans (Managed Beans) to expose data and resources from your application, providing useful data such as the resource usage of your instance and its database latency, allowing you to diagnose problems or performance issues.

On this page, we'll guide you through how to use JConsoleto monitor Confluence locally and remotely. JConsole is included in the Java Development Kit (JDK), but you can use any JMX client.

This page also contains information about In-product diagnostics available through JMX.

This guide provides a basic introduction to the JMX interface and is provided as is. Our Support team can help you troubleshoot a specific Confluence problem, but aren't able to help you set up your monitoring system or interpret the results.

Monitor Confluence using JMX

Monitor Confluence remotely using your APM

To monitor Confluence in yourApplication Performance Monitoring (APM) tool, you'll need to install a JMX exporter to transform theJMX MBeans into the right format for your tool.SeeMonitor application performance to find out how to do this.

If you don't have an Application Performance Monitoring (APM) system,we've created a guide to get you started with Prometheus and Grafana, including some template dashboards that you can use as a jumping off point. See Monitor Confluence with Prometheus and Grafana.

Monitor Confluence remotely using JConsole

Remote monitoring is recommended for production systems, as it does not consume resources on your Confluence server.

To monitor remotely:

1. Add the following properties to your`setenv.sh`/`setenv.bat`file. The port can be any port that is not in use.

   ```
   set CATALINA_OPTS=-Dcom.sun.management.jmxremote %CATALINA_OPTS%
   set CATALINA_OPTS=-Dcom.sun.management.jmxremote.port=8099 %CATALINA_OPTS%
   ```

2. Decide how you will secure your remote connection. SeeRemote Monitoring and Managementfor more information.
   Although it is possible to disable authentication, we do not recommend doing this on a production system.
3. Start JConsole (you'll find it in the`bin`directory of the JDK installation directory).
4. Select**Remote Process**.
5. Enter your hostname and port (this is the port you specified earlier, not the Confluence port).
6. Click**Connect**.

SeeUsing JConsolefor more information on remote monitoring.

Monitor Confluence locally using JConsole

If you are troubleshooting a particular issue, or only need to monitor Confluence for a short time, you can use local monitoring. Local monitoring can have a performance impact on your server, so its not recommended for long-term monitoring of your production system.

To monitor locally:

1. Start JConsole (you'll find it in the`bin`directory of the JDK installation directory)
2. Select**Local Process**.

3. Select the Confluence process. It will be called something like`org.apache.catalina.startup.Bootstrap start`

SeeUsing JConsolefor more information on local monitoring.

Write JMX metrics to a log file

You can also choose to write the following JMX metrics to a log file. This is useful when you are troubleshooting a problem.

- CacheStatistics
- IndexingStatistics
- MailTaskQueue
- RequestMetrics
- SystemInformation
- ThreadPool
- OS
- GC
- Threading
- TomcatManager
- RequestProcessor

To write JMX metrics to a log file:

1. Go to > **General Configuration**> **Scheduled jobs**
2. Enable the **Log JMX Metrics** job.

This job runs once per minute by default, and writes metrics to the `<local-home/logs/atlassian-confluence-jmx.log` file. See Working with Confluence Logs

Disable JMX monitoring

To diable JMX monitoring:

1. Go to> **General Configuration**> Monitoring.
2. Deselect **JMX monitoring**.

This will also disable App monitoring, as it requires JMX to be enabled.

Confluence MBeans

You can use the following Confluence MBeans to see live information about your Confluence instance.

CacheStatistics

This MBean shows information about Confluence caches. This info can also be found on theCache Statisticspag e.

IndexingStatistics

This MBean shows information related to search indexing. Here's some useful attributes.

Property name	Function	Values
Flushing	Indicate whether the cache is currently flushing	True/False
LastElapsedMilliseconds	Time taken during last indexing	Milliseconds
TaskQueueLength	Shows number of tasks in the queue	Integer
ReIndexing	Indicates whether Confluence is currently reindexing	True/False

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

SystemInformation

This MBean shows information such as the Confluence version and uptime. This info can also be found on theSystem Informationpage.

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	The number of requests in the last ten seconds.	Integer

MailServer-SMTPServer

This MBean shows information related to email dispatch attempts and failures. There will be an MBean for every SMTP Mailserver that has been configured in the Confluence instance.

Property name	Function	Values
EmailsAttempted	The number of email messages Confluence has tried to send.	Integer
EmailsSent	The number of email messages sent successfully.	Integer

MailTaskQueue

This MBean shows information related to the email workload.

Property name	Function	Values
ErrorQueueSize	Number of errors in the queue.	Integer
Flushing	Shows state (i.e. flushing, or not)	True/False
FlushStarted	Time that operation began.	Time
RetryCount	The number of retries that were performed.	Integer
TaskSize	Number of email messages queued for dispatch.	Integer

SchedulingStatistics

This MBean shows information related to current jobs, scheduled tasks and the time that they were last run.

Property name	Function	Values
AllJobNames	Shows information on current scheduled jobs including the time they were last run	String

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

CurrentlyRunningJobNames	Lists the scheduled jobs that are currently running	List

App-specific metrics

Enable app monitoring to expose additional metrics that are useful when troubleshooting issues with Marketplace and custom-built apps.

SeeApp metrics reference for a full list of app-specific metrics.

Additional MBeans

To also monitor Hibernate and Hazelcast (Confluence Data Center only) you will need to add the following properties to your`setenv.sh`/`setenv.bat`file first.

```
set CATALINA_OPTS=-Dconfluence.hazelcast.jmx.enable=true %CATALINA_OPTS%
set CATALINA_OPTS=-Dconfluence.hibernate.jmx.enable=true %CATALINA_OPTS%
```

This will make the Hibernate and Hazelcast MBeans available in your JMX client.

Monitoring high CPU consuming threads

TheTop Threads Pluginfor JConsole is useful for monitoring whether the CPU is spiking. Use the following command to start JConsole with this plugin:

```
JConsole -pluginpath /pathto/topthreads.jar
```

In-product diagnostics available through JMX

We've introduced a set of database connectivity, HTTP connection, and indexing metrics for in-product diagnostics available through JMX.

In-product diagnostics (IPD) provides greater insights for you and our Support into how running instances are operating.

IPD uses additional metrics handling Confluences interactions with its database. For example, by using database connectivity metrics, youll efficiently identify what in your environment or infrastructure might cause the performance issues.

The feature is enabled by default. Live metrics are available in the following formats:

- as new JMX MBeans
- as a history of snapshots of the JMX values in the new IPD log file `atlassian-confluence-ipd-monitoring.log`.

The log file is also included in theSupport Zip file created in the Atlassian Troubleshooting and Support app. If needed, send the zip file to Atlassian Support who has the internal tools to interpret it.

The log file is available in the`{confluence_home}\logs`folder where you can find all the existing log files. The log file is also included in the**Support Zip** file created in the ATST plugin. If needed, you can generatethe Support Zip filein the**Atlassian troubleshooting & support tools**plugin and send the file to Atlassian Support, where we have internal tools to interpret it.Learn more about the plugin

Communication

The feature communicates in the following ways:

- JMX: JMX MBeans are updated periodically based on an internal schedule.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

- The log file `atlassian-confluence-ipd-monitoring.log`: JMX values are snapshotted and recorded to the log file on a configurable schedule. By default, the JMX values are polled and written to the log file every 60 seconds.

In-product diagnostics metrics

Expand the following sections to learn more about the metrics available for in-product diagnostics.

> To use the metrics, make sure youve first enabled JMX.

MBean ObjectName	Metric Description
`com.atlassian.confluence:type=metrics,category00=db,category01=connection,category02=latency,name=value`	• db.connection.latency.value The latest measure of latency when querying the database Set to -1 when database connectivity is lost
`com.atlassian.confluence:type=metrics,category00=db,category01=connection,category02=latency,name=statistics`	• db.connection.latency.statistics Aggregated statistics of latency since the last restart
`com.atlassian.confluence:type=metrics,category00=db,category01=connection,category02=pool,category02=numActive,name=statistics`	• db.connection.pool.numActive.statistics Aggregated statistics of the number of active connections in the database connection pool since the last restart
`com.atlassian.confluence:type=metrics,category00=db,category01=connection,category02=pool,category02=numActive,name=value`	• db.connection.pool.numActive.value The latest measure of the number of active connections in the database connection pool. Set to -1 when database connectivity is lost
`com.atlassian.confluence:type=metrics,category00=db,category01=connection,category02=pool,category02=numIdle,name=statistics`	• db.connection.pool.numIdle.statistics Aggregated statistics of the number of idle connections in the database connection pool since the last restart

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

com.atlassian.confluence:type=metrics,category00=db,category01=connection,category02=pool,category02=numIdle,name=value	• db.connection.pool.numIdle.value The latest measure of the number of idle connections in the database connection pool. Set to -1 when database connectivity is lost
com.atlassian.confluence:type=metrics,category00=db,category01=connection,category02=state,name=value	• db.connection.state.value The latest indicator, 0 for a failed connection or 1 for a working connection, of the state of the connection to the database
com.atlassian.confluence:type=metrics,category00=db,category01=connection,category02=failures,name=counter	• db.connection.failures.counter The count of database connection failures since the last restart

Bean ObjectName	Metric Description
com.atlassian.confluence:type=metrics,category00=http,category01=connection,category02=pool,category03=numActive,name=value	• http.connection.pool.numActive.value The latest measure of the number of active connections in the HTTP connection pool. If more than one pool is defined, the sum of active connections from all pools is returned.
com.atlassian.confluence:type=metrics,category00=http,category01=connection,category02=pool,category03=numIdle,name=value	• http.connection.pool.numIdle.value The latest measure of the number of idle connections in the HTTP connection pool. If more than one pool is defined, the sum of idle connections from all pools is returned.
com.atlassian.confluence:type=metrics,category00=http,category01=connection,category02=pool,category03=numMax,name=value	• http.connection.pool.numMax.value The maximum number of threads to be created by the HTTP connectors and made available for requests.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

`com.atlassian.confluence:type=metrics,category00=http,category01=connection,category02=sessions,category03=active,name=value`	• http.connection.sessions.active.value The latest measure of the number of active user sessions
`com.atlassian.confluence:type=metrics,category00=http,category01=connection,category02=sessions,category03=active,name=statistics`	• http.connection.sessions.active.statistics Aggregated statistics of the number of active user sessions
`com.atlassian.confluence:type=metrics,category00=http,category01=connection,category02=sessions,category03=recent,name=value`	• http.connection.sessions.recent.value The latest measure of the number of recent user sessions. Recent session is the session that has been active in the last one hour
`com.atlassian.confluence:type=metrics category00=http,category01=requests,name=value`	• http.requests.value The latest measure of the total number of HTTP requests per minute
`com.atlassian.confluence:type=metrics,category00=http,category01=requests,name=statistics`	• http.requests.statistics Aggregated statistics of the total number of HTTP requests per minute

Bean ObjectName	**Metric Description**

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

`com.atlassian.confluence:type=metrics,category00=index,category01=rebuild,category02=totaltimemillis,name=value`	• com.atlassian.confluence.metrics.index.rebuild.totaltimemillis.value The duration of full reindexing. For multiple nodes, metrics are emitted only on the node where the reindexing is taking place.
`com.atlassian.confluence:type=metrics,category00=index,category01=rebuild,category02=totaltimemillis,name=statistics`	• com.atlassian.confluence.metrics.index.rebuild.totaltimemillis.statistics Aggregated statistics of the duration of full reindexing
`com.atlassian.confluence:type=metrics,category00=index,category01=queue,category02=size,name=value,tag.queueName=main`	• com.atlassian.confluence.metrics.index.queue.size.value `tag.queueName: main` The number of items left in the `main` queue. The increasing number of items left in the `main` queue might indicate that more items are being added to the queue than processed within the same period of time.
`com.atlassian.confluence:type=metrics,category00=index,category01=queue,category02=size,name=value,tag.queueName=change`	• com.atlassian.confluence.metrics.index.queue.size.value `tag.queueName: change` The number of items left in the `change` queue. The increasing number of items left in the `change` queue might indicate that more items are being added to the queue than processed within the same period of time.
`com.atlassian.confluence:type=metrics,category00=index,category01=queue,category02=size,name=value,tag.queueName=edge`	• com.atlassian.confluence.metrics.index.queue.size.value `tag.queueName: edge` The number of items left in the `edge` queue. The increasing number of items left in the `edge` queue might indicate that more items are being added to the queue than processed within the same period of time.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

com.atlassian.confluence:type=metrics,category00=index,category01=queue,category02=batches,category03=processingTimeMillis,name=value,tag.queueName=change	• com.atlassian.confluence.metrics.index.queue.batches.processingTimeMillis.value tag.queueName: change The latest measure of the total processing time of all batches during the IPD measurement interval for the change queue. The default IPD measurement interval is one minute.
com.atlassian.confluence:type=metrics,category00=index,category01=queue,category02=batches,category03=processingTimeMillis,name=value,tag.queueName=edge	• com.atlassian.confluence.metrics.index.queue.batches.processingTimeMillis.value tag.queueName: edge The latest measure of the total processing time of all batches during the IPD measurement interval for the edge queue. The default IPD measurement interval is one minute.
com.atlassian.confluence:type=metrics,category00=index,category01=queue,category02=batches,category03=processingTimeMillis,name=value,tag.queueName=main	• com.atlassian.confluence.metrics.index.queue.batches.processingTimeMillis.value tag.queueName: main The latest measure of the total processing time of all batches during the IPD measurement interval for the main queue. The main index queue is also known as the content Index queue. The default IPD measurement interval is one minute.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

`com.atlassian.confluence:type=metrics,category00=index,category01=queue,category02=batches,category03=processingTimeMillis,name=statistics,tag.queueName=change`	com.atlassian.confluence.metrics. index.queue.batches. processingTimeMillis.statistics `tag.queueName: change` Aggregated statistics of the batch processing time for the `change` queue.
`com.atlassian.confluence:type=metrics,category00=index,category01=queue,category02=batches,category03=processingTimeMillis,name=statistics,tag.queueName=edge`	• com.atlassian.confluence.metrics. index.queue.batches. processingTimeMillis.statistics `tag.queueName: edge` Aggregated statistics of the batch processing time for the `edge` queue.
`com.atlassian.confluence:type=metrics,category00=index,category01=queue,category02=batches,category03=processingTimeMillis,name=statistics,tag.queueName=main`	• com.atlassian.confluence.metrics. index.queue.batches. processingTimeMillis.statistics `tag.queueName: main` Aggregated statistics of the batch processing time for the `main` queue.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

com.atlassian.confluence:type=metrics,category00=index,category01=queue,category02=items,category03=added,name=custom,tag.queueName=change	• com.atlassian.confluence.metrics.index.queue.items.added.custom tag.queueName: change The number of items added to the change queue within one minute.
com.atlassian.confluence:type=metrics,category00=index,category01=queue,category02=items,category03=added,name=custom,tag.queueName=edge	• com.atlassian.confluence.metrics.index.queue.items.added.custom tag.queueName: edge The number of items added to the edge queue within one minute.
com.atlassian.confluence:type=metrics,category00=index,category01=queue,category02=items,category03=added,name=custom,tag.queueName=main	• com.atlassian.confluence.metrics.index.queue.items.added.custom tag.queueName: main The number of items added to the main queue within one minute.
com.atlassian.confluence:type=metrics,category00=index,category01=queue,category02=items,category03=processed,name=custom,tag.queueName=change	• com.atlassian.confluence.metrics.index.queue.items.processed.custom tag.queueName: change The number of items from the change queue processed within one minute.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

`com.atlassian.confluence:type=metrics,category00=index,category01=queue,category02=items,category03=processed,name=custom,tag.queueName=edge`	• com.atlassian.confluence.metrics.index.queue.items.processed.custom `tag.queueName: edge` The number of items from the `edge` queue processed within oneminute.
`com.atlassian.confluence:type=metrics,category00=index,category01=queue,category02=items,category03=processed,name=custom,tag.queueName=main`	• com.atlassian.confluence.metrics.index.queue.items.processed.custom `tag.queueName: main` The number of items from the `main` queue processed within oneminute.

To get more details on cross-product metrics, check the article Interpreting cross-product metrics for in-product diagnostics.

Enabling in-product diagnostics monitoring

IPD monitoring is enabled by default. To manage it:

1. Go to > **General Configuration**
2. In the left panel, select**Monitoring**.
3. Use the **Enable in-product diagnostics** toggle to enable or disable IPD monitoring.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

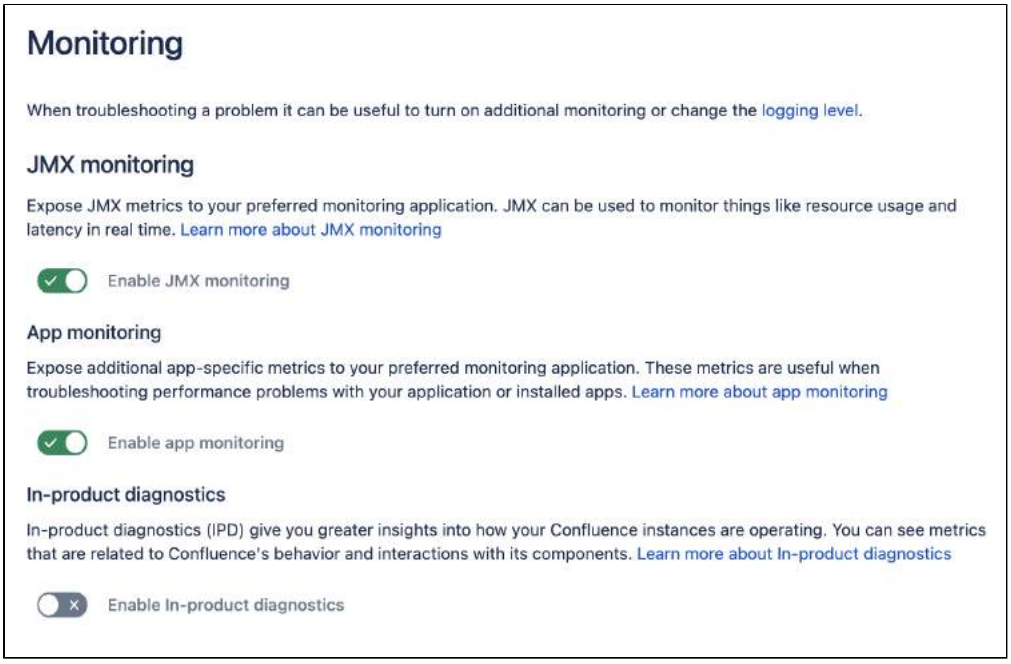

Screenshot: JMX monitoring settings with in-product diagnostics disabled

Log formatting

Writing to `atlassian-confluence-ipd-monitoring.log` is done via log4j. Its configuration is managed in `log4j.properties`.

```
#####################################################
# In-product diagnostics monitoring logging
#####################################################

log4j.appender.ipdLogAppender=com.atlassian.confluence.logging.ConfluenceHomeLogAppender
log4j.appender.ipdLogAppender.LogFileName=atlassian-confluence-ipd-monitoring.log
log4j.appender.ipdLogAppender.MaxFileSize=20480KB
log4j.appender.ipdLogAppender.MaxBackupIndex=5
log4j.appender.ipdLogAppender.layout=com.atlassian.logging.log4j.NewLineIndentingFilteringPatternLayout
log4j.appender.ipdLogAppender.layout.ConversionPattern=%d %m%n

log4j.logger.ipd-monitoring = INFO, consolelog
log4j.additivity.ipd-monitoring = false
log4j.logger.ipd-monitoring-data-logger = INFO, ipdLogAppender
log4j.additivity.ipd-monitoring-data-logger = false
```

Log contents

By default, a concise set of data is included in each log entry. An extended set of data can be logged by enabling the `confluence.in.product.diagnostics.extended.logging` feature flag.

To enable the extended data:

1. Go to `<CONFLUENCE_URL>/admin/darkfeatures.action`, where `<CONFLUENCE_URL>` is the base URL of your Confluence instance.
2. In the **Enable dark feature** field, enter `confluence.in.product.diagnostics.extended.logging`
3. Select **Submit**. Learn how to manage dark features
    a. To disable the extended data, in the **Site Dark Features** section, find `confluence.in.product.diagnostics.extended.logging` and select **remove**.

In the following tables, see the structures of the concise and extended logging formats.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

The metrics in JMX always go in the extended format.

Learn more about metric attributes

Concise data

MBean Type	Properties	Attributes
Counter	timestamp	_count
Value	label	_value
Statistics	attributes	_99thPercentile _max _min _mean

```
2023-01-13 11:51:13,106 IPDMONITORING {"timestamp":"1673610673","label":"DB.CONNECTION.POOL.NUMACTIVE.
STATISTICS","attributes":{"_max":"2.0","_mean":"1.2436699769063984","_99thPercentile":"2.0","_count":"
5","_min":"1.0"}}
```

Extended data

MBean Type	Properties	Attributes
Counter	timestamp label attributes objectName	_count _fifteenMinuteRate _fiveMinuteRate _meanRate _oneMinuteRate _rateUnit
Value		_value _number

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Statistics		_50thPercentile _75thPercentile _95thPercentile _98thPercentile _99thPercentile _999thPercentile _count _min _max _mean _stdDev _durationUnit _fifteenMinuteRate _fiveMinuteRate _meanRate _oneMinuteRate _rateUnit


```
2022-09-06 18:38:48,015 IPDMONITORING {"timestamp":"1662453528","label":
"DB.CONNECTION.LATENCY.STATISTICS","objectName":
"com.atlassian.confluence:category00\u003ddb,category01\u003dconnection,category02\
u003dlatency,name\u003dstatistics,type\u003dmetrics",
"attributes":{"_oneMinuteRate":"0.02012497818617073","_50thPercentile":"0.0",
"_mean":"1.9379304604014412E-25","_max":"1.0","_stdDev":"4.40219315841711E-13",
"_98thPercentile":"0.0","_meanRate":"0.003612560785169162","_rateUnit":
"events/second","_99thPercentile":"0.0","_count":"16","_durationUnit":
"milliseconds","_75thPercentile":"0.0","_fiveMinuteRate":
"0.005912972095043379","_fifteenMinuteRate":"0.0037696657500141968",
"_999thPercentile":"0.0","_95thPercentile":"0.0","_min":"0.0"}}
```


Definitions of metric attributes

Expand the following sections to learn more about metric attributes.

Attribute	Definition
`_count`	The number of occurrences of a metric within the current time window
`_fifteenMinuteRate`	The number of occurrences of a metric over the last 15 minutes
`_fiveMinuteRate`	The number of occurrences of a metric over the last five minutes
`_meanRate`	The **mean** rate at which events have occurred since the meter was created
`_oneMinuteRate`	The number of occurrences of a metric over the last one minute
`_rateUnit`	The unit of measure used for rates

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Pay attention to the following attributes: `_oneMinuteRate`, `_fiveMinuteRate`, and `_fifteenMinuteRate`.

The **_count** gives no indication of how the measurements have changed over time. A sense of recency is provided with the minute rates.

Attribute	Definition
`_value`	A most recently sampled value of the metric
`_number`	Contains the same value as the `_value` attribute

The metrics of the statistics MBean type are also known as aggregated values. They provide statistics for the items that have been subjected to any changes over a period of time. For example, for the items that have been processed in a mail queue or added to an error mail queue.

Time window

Unless stated, aggregated values are calculated over a sliding time window. It covers the last five minutes, approximately.

Percentile values are calculated using a reservoir sampling technique. This technique uses a small, manageable set of values that is statistically representative of the data stream as a whole, hence reducing the quantity of data that must be held in memory.

Resets

Outside of the sliding time window, all aggregated values are reset:

- After each system restart.
- After each time JMX monitoring or in-product diagnostic metrics are enabled.

Learn more about JMX monitoring and in-product diagnostic in other Data Center products:

- Live monitoring using the JMX interface in Jira
- Enabling JMX counters for performance monitoring in Bitbucket

In the following table, find the definitions of statistics metric attributes.

Attribute	Definition
`_50thPercentile`	A measured value below which 50% of all measurements can be found within the current time window; also referred to as the **median** value. This attribute provides an alternative to the **mean** as a representation of the middle measurement. The **median** is less likely to be skewed by outlier values than the **mean**.
`_75thPercentile`	The measured value below 75% of all measurements that can be found within the current time window; the **third quartile** value
`_95thPercentile`	The measured value below 95% of all measurements that can be found within the current time window
`_98thPercentile`	The measured value below 98% of all measurements that can be found within the current time window
`_99thPercentile`	The measured value below 99% of all measurements that can be found within the current time window

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

`_999thPercentile`	The measured value below 999% of all measurements that can be found within the current time window
`_count`	The number of occurrences of a metric within the current time window
`_min`	The minimum measured value within the current time window
`_max`	The maximum measure value within the current time window; the **statistical range** between `_max` and `_min` provides a measure of values variability
`_mean`	The average value within the current time window. This attribute can be skewed by large outlier measurements. In such cases, the `_50thPercentile` provides a better measure of the middle value.
`_stdDev`	A measure of the data variability. A low standard deviation indicates that the values tend to be close to the **mean** of the set, while a high standard deviation indicates that the values are spread out over a wider range of values.
`_durationUnit`	The unit of measure used for durations
`_fifteenMinuteRate`	The number of occurrences of a metric over the last 15 minutes
`_fiveMinuteRate`	The number of occurrences of a metric over the last five minutes
`_meanRate`	The **mean** rate at which events have occurred since the meter was created
`_oneMinuteRate`	The number of occurrences of a metric over the last one minute
`_rateUnit`	The unit of measure used for rates

Pay attention to the following attributes: `_oneMinuteRate`, `_fiveMinuteRate`, and `_fifteenMinuteRate`.

The **_count** gives no indication of how the measurements have changed over time. A sense of recency is provided with the minute rates.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

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 theConfluence Release Notes.

For information on using and administering Confluence refer to the Confluence Documentation.

Long Term Support releases

A Long Term Support release is a feature release that gets backported critical security updates and critical bug fixes during its entire two-year support window. If you can only upgrade once a year, consider upgrading to a Long Term Support release.Learn more

Long Term Support releases were originally referred to as Enterprise Releases.

- System Requirements
  - Server Hardware Requirements Guide
  - Running Confluence in a Virtualized Environment
- Confluence Installation Guide
  - Installing Confluence
  - Installing Confluence Data Center
  - Installing Java for Confluence
  - Creating a Dedicated User Account on the Operating System to Run Confluence
- Confluence Setup Guide
  - Configuring Jira Integration in the Setup Wizard
- Upgrading Confluence
  - Upgrading Beyond Current Licensed Period
  - Confluence Post-Upgrade Checks
  - Migration from Wiki Markup to XHTML-Based Storage Format
  - Migration of Templates from Wiki Markup to XHTML-Based Storage Format
  - Upgrading Confluence Manually
  - Create a staging environment for upgrading Confluence
  - Upgrade Confluence without downtime
- Supported Platforms
  - End of Support Announcements for Confluence
  - Bundled Tomcat and Java versions
  - Supported Platforms FAQ

Downloads

Download the Confluence documentation in PDF format.

Other resources

Confluence Release Notes

Confluence administrator's guide

Confluence Knowledge Base

Atlassian Answers

System Requirements

Confluence can run on a wide range of operating systems and databases, on physical or virtualized servers.

SeeSupported Platformsfor the full list of platforms that we support in this version of Confluence orSupported Platforms FAQfor details on our support handling procedures.

On this page:

- Software requirements
  - Operating systems
  - Application server
  - Databases
  - Java
  - Antivirus considerations
- Hardware requirements
- Hosted solutions Confluence Cloud

Software requirements

Operating systems

Atlassian supports the operating systems listed on theSupported Platformspage.

If you would like to run Confluence on virtualized hardware, please read ourRunning Confluence in a Virtualized Environmentdocument first.

Application server

We only support running Confluence on the version of Apache Tomcat that is bundled with the Confluence distribution.

Databases

You'll need an external database to run Confluence. See the Supported Platforms page for a list of all the databases we support.

When evaluating Confluence, you can use the embedded H2 database included in the Confluence installation, but you will need to migrate to a supported external database once you're ready to roll Confluence out to your team.

Java

The Java Runtime Environment (JRE) is packed up and ready to go when you install Confluence using the Windows or Linux installer. You don't need to install Java yourself.

If you choose to install Confluence from an archive file, you'll need asupportedJRE or JDK, and your JAVA_HOME variable set correctly. SeeInstalling Java for Confluencefor more information.

Antivirus considerations

Antivirus software on the operating system running Confluence can greatly decrease the performance of Confluence. Antivirus software that intercepts access to the hard disk is particularly detrimental and may even cause errors in Confluence.This is particularly important if you are running Confluence on Windows. No matter how fast your hardware is, antivirus software will almost always have a negative impact on Confluence's performance.

You should configure your antivirus software to ignore the following directories:

- Confluence home directory
- Confluence's index directory
- All database-related directories

Hardware requirements

Please be aware that while some of our customers run Confluence on SPARC-based hardware, Atlassian only officially supports Confluence running on x86 hardware and 64-bit derivatives of x86 hardware.

See Server Hardware Requirements Guide for more information.

You may also like to check out our tips on reducing out of memory errors, in particular the section on Permanent Generation Size.

Hosted solutions Confluence Cloud

If you do not have the resources to set up and maintain a Confluence installation locally, consider trying Confluence Cloud. Atlassian can run and maintain your installation of Confluence, handling all the testing, monitoring and upgrading processes for you.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

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.

> **Enterprise-scale Confluence sites**
>
> These recommendations are designed for small or medium sized Confluence sites. For guidance on large or extra large sites, read our enterprise-scale infrastructure recommendations.
>
> We've also created load profiles to help you determine the size of your site.

On this page:

- Minimum hardware requirements
- Example hardware specifications
- Server load and scalability
- Maximum reported usages
- Hard disk requirements
  - Private and public comparison
- Professional assistance
- Example site

Related pages:

- Confluence Installation Guide
- Managing Application Server Memory Settings
- Running Confluence in a Virtualized Environment

Minimum hardware requirements

The values below refer to the minimum available hardware required to run Confluence only; for example, the minimum heap size to allocate to Confluence is 1 GB and 1 GB for Synchrony (which is required for collaborative editing). You'll need additional physical hardware, of at least the minimum amount required by your Operating System and any other applications that run on the server.

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.

Here is our minimum hardware recommendation:

- **CPU:**Quad core 2GHz+ CPU
- **RAM:**6GB
- **Minimum database space:** 10GB

Note: Please be aware that while some of our customers run Confluence on SPARC-based hardware, we only officially support Confluence running on x86 hardware and 64-bit derivatives of x86 hardware. Confluence typically will not perform well in a tightly constrained, shared environment - examples include an AWS micro.t1 instance. Please be careful to ensure that your choice of hosting platform is capable of supplying sustained processing and memory capacity for the server, particularly the processing-intense startup process.

Example hardware specifications

These are example hardware specifications for non-clustered Confluence instances. It is not recorded whether the amount ofRAM refers toeither the totalservermemoryor memoryallocated 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.

Server load and scalability

When planning server hardware requirements for your Confluence deployment, you will need to estimate the server scalability based on peak visitors, the editor to viewer ratio and total content.

- The editor to viewer ratio is how many visitors are performing updates versus those only viewing content
- Total content is best estimated by a count of total spaces

Confluence scales best with a steady flow of visitors rather than defined peak visitor times, few editors and few spaces. Users should also take into account:

- Total pages is not a major consideration for performance. For example, instances hosting 80K of pages can consume under 512MBof memory
- Alwaysuse an external database, and check out theperformance tuning guides.

Maximum reported usages

These values are largest customer instances reported to Atlassian or used for performance testing. Clustering, database tuning and other performance tuning is recommended for instances exceeding these values.

Most Spaces	1700
Most Internal Users	15K
Most LDAP Users	100K
Most Pages	80K

Hard disk requirements

All page content is stored in the database, while attachments are stored in the file system.The more attachments you have, the more disk space you will require.

Private and public comparison

Private instances manage their users either internally or through a user repository such as LDAP, while online instances have public signup enabled and must handle the additional load of anonymous internet visitors. Please keep in mind that these are examples only, not recommendations:

Use Case	Spaces	User Accounts	Editors	Editor To Viewer Ratio	Pages	Page Revisions	Attachments	Comment

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

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 anAtlassian Solution Partnerfor expertise on hardware sizing, testing and performance tuning.

Example site

Here's a breakdown of the disk usage and memory requirements of a large documentation site as at April 2013:

Database size	2827 MB
Home directory size	116 GB
Average memory in use	1.9 GB

Size of selected database tables

Data	Relevant Table	Rows	Size
Attachment metadata	`attachments`	193903	60 MB
Content and user properties	`os_propertyentry (?)`	639737	255 MB
Content bodies (incl. all versions of blogs, pages and comments)	`bodycontent`	517520	1354 MB
Content metadata (incl. title, author)	`content`	623155	459 MB
Labels	`label (5982, 1264 kB),` `content_label (134151, 46 MB)`	140133	47.2 MB
Users	`users`	38766	6200 kB

Note: not all database tables or indexes are shown, and average row size may vary between instances.

Size of selected home directory components

Data	Files	Size
Attachments (incl. all versions)	207659	105 GB

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

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.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Running Confluence in a Virtualized Environment

This page provides pointers for things to look at when running Confluence on virtualized hardware.

On this page:

- Summary
- Recommendations
- Further help

Related pages:

- Server Hardware Requirements Guide
- Confluence Installation Guide
- Confluence Data Center documentation
- AWS Quick Start (Data Center only)

Summary

Running Confluence in a virtual machine (VM) requires specialized skills to set up and manage the virtualized environment. In particular, the performance of Confluence can be affected by the activity of other VMs running on the same infrastructure, as well as how you configure the Confluence VM itself.

Atlassian supports running Confluence and Confluence Data Center in a virtualized environment, but we cannot offer support for problems which are related to the environment itself.

Recommendations

The following recommendations come from our experience in running and testing Confluence in virtualized environments like VMWare and KVM, and our experience in working with customers running on these platforms.

- **Know your platform.** Consult the documentation for your operating system and your chosen virtualization technology, for details on setting up a reliable VM (virtual machine) image.
- **Allocate enough memory.** As a Java web application, Confluence requires a relatively large memory allocation, compared to some other web technologies. Ensure that your VM images have enough physical memory allocated to run Confluence without swapping.
- **Handle high I/O.** Under normal usage, Confluence requires a significant number of input/output (I/O) operations to the database and home directory for each web request. Ensure that you use the correct drivers and consider how you make storage available to your VMs to optimize this access.
- **Handle peak CPU and memory usage.** For certain operations (including PDF export, Office document processing, and displaying large pages) Confluence requires a significant amount of CPU and memory. Ensure that your virtualization infrastructure has the flexibility and capacity to deal with peak load, not just idle load.
- **Synchronize time correctly.** Some customers have had problems with time synchronization between the VM and the host system. This causes problems in Confluence due to irregularities in the execution of scheduled tasks. We strongly recommend checking your VM time sync if you have issues with scheduled tasks in a virtualized environment.

Further help

For further assistance in setting up a virtualized environment for running Confluence, you may want to consult an Atlassian Solution Partner. Several experts have experience with installation and performance tuning, and can help you with your Confluence configuration.

Confluence Installation Guide

Before you start

Before installing Confluence, please check that you meet the minimum system requirementsandSupported Platforms.

If you're planning to run Confluence in a virtualized environment seeRunning Confluence in a Virtualized Environment.

Choose your installation method

There are a number of ways to install Confluence. Choose the method that is best for your environment.

Install method	Is this right for you?
Install a Confluence trial • Windows, Linux or OS X	This is the fastest way to get a Confluence site up and running. If you want to see what Confluence can do, use this option or try Confluence Cloud free.
Install Confluence using an installer • Windows • Linux	This option uses an installer, and is the most straightforward way to get your production site up and running on a Windows or Linux server.
Install Confluence from a zip or archive file • Windows • Linux	This option requires you to manually install files and configure some system properties. It gives you the most control over the install process. Use this option if there isn't an installer for your operating system.
Run Confluence Server in a Docker container • Docker	This option gets Confluence Server up and running using a pre-configured Docker image. Head to https://docs.docker.com/ to find out more about Docker. Atlassian supports running Confluence in a Docker container, but we cannot offer support for problems which are related to the environment itself.
Install Confluence Data Center in a cluster • Windows and Linux • Kubernetes • AWS Quick Start • Azure	You can deploy a Confluence Data Center cluster on your own infrastructure or a public cloud platform like AWS or Azure. Read the Confluence Data Center Technical Overview for more details on clustering

Note:we do not support installing Confluence as a production system on OS X. An OS X download is available for the purposes of evaluating Confluence only. There are no limitations to using Confluence on a mac with any one of thesupported browsers.

The EAR/WAR distribution is no longer available, you'll need to install Confluence from a zip or archive file if you previously deployed Confluence into an existing application server.

Installing Confluence

There are a number of ways to install Confluence. Choose the method that is best for your environment.

Install method	Is this right for you?
Install a Confluence trial • Windows, Linux or OS X	This is the fastest way to get a Confluence site up and running. If you want to see what Confluence can do, use this option or try Confluence Cloud free.
Install Confluence using an installer • Windows • Linux	This option uses an installer, and is the most straightforward way to get your production site up and running on a Windows or Linux server.
Install Confluence from a zip or archive file • Windows • Linux	This option requires you to manually install files and configure some system properties. It gives you the most control over the install process. Use this option if there isn't an installer for your operating system.
Run Confluence Server in a Docker container • Docker	This option gets Confluence Server up and running using a pre-configured Docker image. Head to https://docs.docker.com/ to find out more about Docker. Atlassian supports running Confluence in a Docker container, but we cannot offer support for problems which are related to the environment itself.
Install Confluence Data Center in a cluster • Windows and Linux • Kubernetes • AWS Quick Start • Azure	You can deploy a Confluence Data Center cluster on your own infrastructure or a public cloud platform like AWS or Azure. Read the Confluence Data Center Technical Overview for more details on clustering

Get a Confluence Data Center trial license

A trial license gives you access to a full instance of Confluence Data Center for 30 days. At the end of the trial period your Confluence Data Center site will become read-only and youll have the option to buy a full license to continue using it, so you wont lose any of your projects or data.

We support single-node Confluence Data Center both for trial and full license instances, so you dont have to modify your current number of application nodes if you dont want to scale up to a cluster yet.

To create a Confluence Data Center trial license:

1. Head tomy.atlassian.com and log in with your Atlassian ID.
2. From the list of Atlassian products, select **Confluence**, then select the **Data Center** option and fill out the form with your organizations information.
3. Select **Generate license**.

If you're ready to scale up your instance, check out how to Migrate from Server to Data Center.

If youre a new customer, the next step is to download and set up your new Confluence Data Center trial instance

Install a Confluence Data Center trial

Want to quickly get up and running with Confluence Data Center? This page will guide you through a few simple steps to install and set up a trial Confluence Data Center site.

A trial license gives you access to a full instance of Confluence Data Center for 30 days. At the end of the trial period your Confluence Data Center site will become read-only and youll have the option to buy a full license to continue using it. Either way, you wont lose any of your projects or data.

On this page:

1. Download the installer
2. Install Confluence
3. Set up Confluence

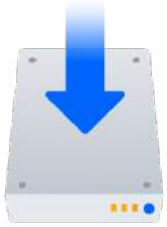

Before you begin

Confluence installers come with all the bits and pieces you need to run the application, but there's a few things you'll need to get up and running:

For a list of supported platforms, see Supported Platforms

- A computer or laptop with a supported operating system - you'll be installing Confluence so you'll need admin rights.

  You can install Confluence on a Windows or Linux operating system.

  Apple Mac isn't supported for production sites, but if you're comfortable setting up applications on your Mac from scratch, you can download the `tar.gz` file and follow the instructions forInstalling Confluence on Linux from Archive File as the process is similar.

- A valid email addressyou'll need it to generate your 30-day trial license and create an account.
- An external database.To run Confluence you'll need an external database. Check the Supported Platforms page for the version you're installing for the list of databases we currently support. If you don't already have a database, PostgreSQL is free and easy to set up.

  Good to know:

  - Set up your database before you begin. Step-by-step guides are available for PostgreSQL, Oracle, MySQL, and SQL Server.
  - If you're using Oracle or MySQL you'll need to download the driver for your database.

Ready to get going? Let's start with downloading the installer.

1. Download the installer

Head towww.atlassian.com/software/confluence/downloadand download the installer for your operating system.

2. Install Confluence

The installer allows you to select **Express** or **Custom** installations.

The **Custom** installation allows you to pick some specific options for Confluence, but for this guide we'll use the **Express** installation.

1. Run the installer. We recommend running with a Windows administrator account.
   If prompted, make sure you allow the installer to make changes to your computer. This way youll be able to install Confluence as a service.
2. Select **Express install**, then select **Next**.
3. Once installation is complete, it will ask you if you want to open Confluence in your browser. Make sure this option is selected, then select **Done**.
4. Confluence will open in your default browser, and you're ready to start the setup wizard.

1. Change to the directory where you downloaded Confluence, then execute this command to make it executable:

   ```
   $ chmod a+x atlassian-confluence-X.X.X-x64.bin
   ```

   Where `X.X.X` is is the Confluence version you downloaded.
2. Run the installer. We recommend using `sudo`to run the installer as this will create a dedicated account to run Confluence and allow you to run Confluence as a service.

   ```
   $ sudo ./atlassian-confluence-X.X.X-x64.bin
   ```

3. When prompted, choose **Express Install** (option 1).
4. Once installation is complete head tohttp://localhost:8090/in your browser to begin the setup process.

3. Set up Confluence

The set up wizard is the last step in getting Confluence up and running. You'll need your email address to generate your evaluation license.

1. Select**Trial**, then select**Next**.
2. Select**Get an evaluation license** andfollow the prompts to generate your trial Confluence Data Center license.
3. Select whether you want to try a **Standalone** (single node) or **Clustered** installation. **Standalone** is the fastest way to get started. If you select **Clustered**, you'll need to configure your cluster before continuing.
4. Enter the details for your database. See the Before you begin section of this page for details and connection options.
5. Select**Manage users with Confluence**, then select**Next**.
6. Enter and confirm the details you want to use for your administrator account, then select**Done**.

That's it! You're ready to team up with some colleagues and start using Confluence.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Installing Confluence on Windows

In this guide we'll run you through installing Confluence in a production environment, with an external database, using the Windows installer.

This is the most straightforward way to get your production site up and running on a Windows server.

On this page:

- Before you begin
- Install Confluence
  - 1. Download Confluence
  - 2. Run the installer
- Set up Confluence
  - 3. Choose installation type
  - 4. Enter your license
  - 5. Connect to your database
  - 6. Populate your new site with content
  - 7. Choose where to manage users
  - 8. Create your administrator account
  - 9. Start using Confluence
- Troubleshooting

Other ways to install Confluence:

- Evaluation -get your free trial up and running in no time.
- Zip install Confluence manually from a zip file.
- Linux install Confluence on a Linux operating system

Before you begin

Before you install Confluence, there's a few questions you need to answer.

Are you using a supported operating system?	Check the Supported Platforms page for the version of Confluence you are installing. This will give you info on supported operating systems, databases and browsers. **Good to know:** • We don't support installing Confluence on OSX. • The Confluence installer includes Java (JRE) and Tomcat, so you don't need to install these separately.

Do you want to run Confluence as a Windows Service?	Running Confluence as a service in Windows means that Confluence will automatically start up when Windows is started. **If you choose to run Confluence as a service:** • You must run the installer as administrator to be able to install Confluence as a service. • The Confluence service will be run as the Windows 'SYSTEM' user account. To change this user account see Changing the Windows user that the Confluence service uses . • We strongly recommend creating a dedicated user account (e.g. with username 'confluence') for running Confluence. See Creating a Dedicated User Account on the Operating System to Run Confluence to find out what directories this user will need to be able to read and write to. **If you choose not to run Confluence as a service:** • You will start and stop Confluence using the Windows Start menu, or by running a file in your Confluence installation directory. • Confluence will be run as the Windows user account that was used to install Confluence, or you can choose to run as a dedicated user. • Confluence will need to be restarted manually if your server is restarted.
Are ports 8090 and 8091 available?	Confluence runs on port 8090 by default. If this port is already in use, the installer will prompt you to choose a different port. Synchrony, which is required for collaborative editing, runs on port 8091 by default. If this port is already in use, you will need to change the port that Synchrony runs on after your Confluence installation is complete. See Administering Collaborative Editing to find out how to change the port Synchrony runs on.You won't be able to edit pages until Synchrony has an available port. SeePorts used by Atlassian Applicationsfor a summary of all the ports used.
Is your database set up and ready to use?	To run Confluence you'll need an external database. Check the Supported Platforms page for the version you're installing for the list of databases we currently support. If you don't already have a database, PostgreSQL is free and easy to set up. **Good to know:** • Set up your database before you begin. Step-by-step guides are available for PostgreSQL, Oracle, MySQL, and SQL Server. • If you're using Oracle or MySQL you'll need to download the driver for your database.
Do you have a Confluence license?	You'll need a valid license to use Confluence. **Good to know:** • If you have not yet purchased a Confluence license you'll be able to create an evaluation license during setup. • If you already have a license key you'll be prompted to log in to my.atlassian.com to retrieve it, or you can enter the key manually during setup. • If you're migrating from Confluence Cloud, you'll need a new license. • Weve ended sales for new server licenses and will end support for server on February 2, 2024. Were continuing our investment in Data Center.Learn more

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Do you want to store your attachment data on object storage?	By default, Confluence stores attachments in the home directory (e.g. in a file system). If your team has large or increasing data sets and requires the ability to scale efficiently, we recommend you use S3 object storage. Amazon S3 is currently the only Confluence-supported object storage solution. **Good to know:** • Amazon S3 object storage is an optional attachment storage method available to anyone on a Data Center license and running Confluence in AWS. • If you're a new customer, seeS3 object storagefor setup instructions. • If you're an existing customer, you'll need to migrate your attachment data to S3 object storage from the file system or another storage method. SeeAttachment storage configurationfor steps to do this. • Even if you use S3 object storage, other non-attachment data will still be stored in your home directory.

Theres a known issue during setup where a load balancer (or proxy) pings the server and breaks Confluence installation or migration to Data Center. See

CONFSERVER-61189 - Opening the base URL multiple times during Data Center migration will break the migration process. GATHERING IMPACT

During installation, you need to disable load balancer health checks and make sure you dont open multiple tabs that point to the same Confluence URL.

Install Confluence

1.Download Confluence

Download the installer for your operating system -https://www.atlassian.com/software/confluence/download

2. Run the installer

1. Run the installer. We recommend using a dedicated Windows administrator account.
2. Follow the prompts to install Confluence. You'll be asked for the following info:
    a. **Destination directory**this is whereConfluencewill be installed.
    b. **Home directory**this is where Confluence data like logs, search indexes and files will be stored.
    c. **TCP ports**these arethe HTTP connector port and control port Confluence will run on. Stick with the default unless you're running another application on the same port.
    d. **Install as service**this option is only available if you ran the installer as administrator.
3. Confluence will start up in your browser once installation is complete.

Set up Confluence

3. Choose installation type

1. Choose**Production installation**.
2. Choose any**apps**you'd also like to install.

4. Enter your license

Follow the prompts to log in tomy.atlassian.comto retrieve your license, or enter a license key.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

5. Connect to your database

1. If you've not already done so, it's time to create your database. See the 'Before you begin' section of this page for details and connection options.
2. For MySQL and Oracle, follow the prompts to download and install the required driver.
3. Enter your database details. Use**test connection** to check your database is set up correctly.

   If you want to specify particular parameters, you can choose to connect **By connection string**. You'll be prompted to enter:

   - **Database URL** the JDBC URL for your database. If you're not sure, check the documentation for your database.
   - **Username and Password** A valid username and password that Confluence can use to access your database.

6. Populate your new site with content

Choose whether you'd like Confluence to populate your site with content:

This option will create a space that you and your users can use to get to know Confluence. You can delete this space at any time.

Use this option if you have a full site export of an existing Confluence site. This is useful when youre migrating to another database or setting up a test site.

Good to know:

- You can only import sites from the **same** or **earlier** Confluence version.
- The system administrator account and all other user data and content will be imported from your previous installation.

In the setup wizard:

- **Upload a backup file** use this option if your site export file is small (25mb or less).
- **Restore a backup file from the file system** use this option if your backup file is large. Drop the file into your `<confluence-home>/restore`directory then follow the prompts to restore the backup.
- **Build Index** well need to build an index before your imported content is searchable. This can take a long time for large sites, so deselect this option if you would rather build the index later.Your content won't be searchable until the index is built.

7. Choose where to manage users

Choose to manage Confluence's users and groups inside Confluence or in a Jira application, such as Jira Software or Jira Service Management:

Choose this option if you're happy to manage users in Confluence, or don't have a Jira application installed.

Good to know:

- If you do plan to manage users in a Jira application, but have not yet installed it, we recommend installing Jira first, and then returning to the Confluence setup.
- You can add external user management (for example LDAP, Crowd or Jira) later if you choose.

Choose this option if you have a Jira application installed and want to manage users across both applications.

Good to know:

- This is a quick way of setting up your Jira integration with the most common options.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

- It will configure a Jira user directory for Confluence, and set up application links between Jira and Confluence for easy sharing of data.
- You'll be able to specify exactly which groups in your Jira app should also be allowed to log in to Confluence. Your license tiers do not need to be the same for each application.
- You'll need either Jira 4.3 or later, Jira Core 7.0 or later, Jira Software 7.0 or later, or Jira Service Management 3.0 or later.

In the setup wizard:

- **Jira Base URL** the address of your Jira server, such as `http://www.example.com:8080/jira/` or `http://jira.example.com/`
- **Jira Administrator Login** this is the username and password of a user account that has the Jira System Administrator global permission in your Jira application. Confluence will also use this username and password to create a local administrator account which will let you access Confluence if Jira is unavailable. Note that this single account is stored in Confluence's internal user directory, so if you change the password in Jira, it will not automatically update in Confluence.
- **Confluence Base URL** this is the URL Jira will use to access your Confluence server. The URL you give here overrides the base URL specified in Confluence, for the purposes of connecting to the Jira application.
- **User Groups** these are the Jira groups whose members should be allowed to use Confluence. Members of these groups will get the 'Can use' permission for Confluence, and will be counted in your Confluence license. The default user group name differs depending on your Jira version:
  - Jira 6.4 and earlier: `jira-users`.
  - Jira Software 7.x and later: `jira-software-users`
  - Jira Core 7.x and later: `jira-core-users`
  - Jira Service Management (formerly Jira Service Desk) 3.x and later: `jira-servicedesk-users`
- **Admin Groups** provide one or more Jira groups whose members should have administrative access to Confluence. The default group is `jira-administrators`. These groups will get the system administrator and Confluence administrator global permissions in Confluence.

8. Create your administrator account

Enter details for the administrator account.

Skip this step if you chose to manage users in a Jira application or you imported data from an existing site.

9. Start using Confluence

That's it!Your Confluence site is accessible from a URL like this: http://<computer_name_or_IP_address>:<port>

If you plan to run Confluence behind a reverse proxy, check outProxy and SSL considerationsbefore you go any further.

Here's a few things that will help you get your team up and running:

- Set the server base URL this is the URL people will use to access Confluence.
- Set up a mail server this allows Confluence to send people notification about content.
- Add and invite users get your team on board!
- Start and stop Confluence find out how to start and stop Confluence.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

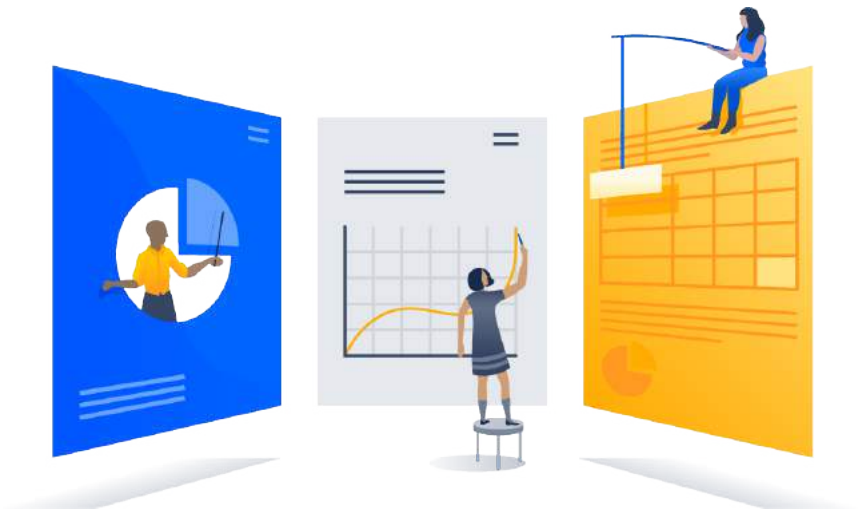

Troubleshooting

- Some anti-virus or other Internet security tools may interfere with the Confluence installation process and prevent the process from completing successfully. If you experience or anticipate experiencing such an issue with your anti-virus/Internet security tool, disable this tool first before proceeding with theConfluence installation.
- Can't start Confluence? See Confluence does not start due to Spring Application context has not been set .
- Collaborative editing errors? See Troubleshooting Collaborative Editing.

Head to Installation Troubleshooting in our Knowledge Base for more help.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Installing Confluence on Windows from Zip File

In this guide we'll run you through installing Confluence in a production environment, with an external database, manually using a zip file.

This method gives you the most control of the installation process.

Other ways to install Confluence:

- Evaluation-get your free trial up and running in no time.
- Installer install Confluence using the Windows installer.
- Linux install Confluence on a Linux operating system.

On this page:

- Before you begin
- Install Confluence
  - 1. Download Confluence
  - 2. Create the installation directory
  - 3. Create the home directory
  - 4. Check the ports
  - 5. Start Confluence
- Set up Confluence
  - 6. Choose installation type
  - 7. Enter your license
  - 8. Connect to your database
  - 9. Populate your new site with content
  - 10. Choose where to manage users
  - 11. Create your administrator account
  - 12. Start using Confluence
- Troubleshooting

Before you begin

Before you install Confluence, there's a few questions you need to answer.

Are you using a supported operating system and Java version?	Check the Supported Platforms page for the version of Confluence you are installing. This will give you info on supported operating systems, databases and browsers. **Good to know:** • We don't support installing Confluence on OS X or mac OS for production environments. • You'll need to install either Adoptium OpenJDK (formerly AdoptOpenJDK)or Oracle JDK. We don't support other OpenJDK binaries. • You can use either the JDK (Java Development Kit) or JRE (Java Runtime Environment). • We only support the version of Apache Tomcat that is bundled with Confluence.

Do you want to run Confluence as a Windows Service?	Running Confluence as a service in Windows means that Confluence will automatically start up when Windows is started. You should use the Windows installer if you want to run Confluence as a Service. **If you choose not to run Confluence as a service:** • You will start and stop Confluence by running the `start-confluence.bat` file in your Confluence installation directory. • Confluence will be run as the Windows user account that was used to install Confluence, or you can choose to run as a dedicated user (this user must have full read and write access to the installation directory and home directory). • Confluence will need to be restarted manually if your server is restarted.
Are ports 8090 and 8091 available?	Confluence runs on port 8090 by default. If this port is already in use, the installer will prompt you to choose a different port. Synchrony, which is required for collaborative editing, runs on port 8091 by default. If this port is already in use, you will need to change the port that Synchrony runs on after your Confluence installation is complete. See Administering Collaborative Editing to find out how to change the port Synchrony runs on.You won't be able to edit pages until Synchrony has an available port. SeePorts used by Atlassian Applicationsfor a summary of all the ports used.
What database do you plan to use?	To run Confluence you'll need an external database. Check the Supported Platforms page for the version you're installing for the list of databases we currently support. If you don't already have a database, PostgreSQL is free and easy to set up. **Good to know:** • Set up your database before you begin. Step-by-step guides are available for PostgreSQL, Oracle, MySQL, and SQL Server. • If you're using Oracle or MySQL you'll need to download the driver for your database.
Do you have a Confluence license?	You'll need a valid license to use Confluence. **Good to know:** • If you have not yet purchased a Confluence license you'll be able to create an evaluation license during setup. • If you already have a license key you'll be prompted to log in to my.atlassian.com to retrieve it, or you can enter the key manually during setup. • If you're migrating from Confluence Cloud, you'll need a new license. • Weve ended sales for new server licenses and will end support for server on February 2, 2024. Were continuing our investment in Data Center.Learn more

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Is your JRE_HOME variable set correctly?	Before you install Confluence, check that you're running a supported Java version and that the `JRE_HOME` (or `JAVA_HOME`) environment variable is set correctly. **To check the JRE_HOME variable**: Open a command prompt and type `echo %JRE_HOME%` and hit Enter. • If you see a path to your Java installation directory, the `JRE_Home` environment variable has been set correctly. • If nothing is displayed, or only `%JRE_HOME%` is returned, you'll need to set the `JRE_HOME` environment variable manually. See Setting the JAVA_HOME Variable in Windows for a step by step guide.

⚠ Theres a known issue during setup where a load balancer (or proxy) pings the server and breaks Confluence installation or migration to Data Center. See

CONFSERVER-61189 - Opening the base URL multiple times during Data Center migration will break the migration process. GATHERING IMPACT

During installation, you need to disable load balancer health checks and make sure you dont open multiple tabs that point to the same Confluence URL.

Install Confluence

1.Download Confluence

Download the zip file for your operating system https://www.atlassian.com/software/confluence/download.

2. Create the installation directory

1. Create your installation directory (with full control permissionfor thededicated Windows administrator accountyou'll use to run Confluence) this is where Confluence will be installed. Avoid using spaces or special characters in the path. We'll refer to this directory as your`<installation-directory>`.
2. Extract the Confluence zip file to your`<installation-directory>`.We recommend using7zipor Winzip.

3. Create the home directory

1. Create your home directory (with full control permission) this is where Confluence data like logs, search indexes and files will be stored. This should be seperate to your installation directory. We'll refer to this directory as your`<home-directory>`.
2. Edit`<installation-directory>\confluence\WEB-INF\classes\confluence-init.properties`.
3. At the bottom of the file, enter the path to your`<home directory>`.

   You can edit the confluence-init.properties file in Notepad or any other text editor.

   a. Scroll to the bottom of the text and find this line:

   ```
   # confluence.home=c:/confluence/data
   ```

   b. Remove the '#' and the space at the beginning of this line (so Confluence doesn't regard the line as a comment)

   ```
   confluence.home=c:/data/confluence-home
   ```

   c. If you decide to use a different directory as the home directory you should:
      • Avoid spaces in the directory path or file name.
      • Use forward slashes '/' to define the path in this file.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

4. Check the ports

By default Confluence listens on port`8090`. If you have another application running on your server that uses the same ports, you'll need to tell Confluence to use a different port.

To change the ports:

1. Edit `<installation-directory>\conf\server.xml`
2. Change the **Server** port (8000) and the **Connector** port (8090) to free ports on your server.

   In the example below we've changed the **Server** port to 5000 and the **Connector** port to 5050.

```
Server port="5000" shutdown="SHUTDOWN" debug="0">
  <Service name="Tomcat-Standalone">
    <Connector port="5050" connectionTimeout="20000" redirectPort="8443"
        maxThreads="48" minSpareThreads="10"
        enableLookups="false" acceptCount="10" debug="0" URIEncoding="UTF-8"
        protocol="org.apache.coyote.http11.Http11NioProtocol" />
```

5. Start Confluence

1. Run`<installation-directory>/bin/start-confluence.bat`to start the install process. We recommend using a dedicated user account.

   A command prompt will open. Closing this window will stop Confluence.

2. Go to `http://localhost:8090/` to launch Confluence in your browser (change the port if you've updated the Connector port).

- If the command prompt window closes immediately, your JAVA_HOME variable may not be set correctly. See Setting the JAVA_HOME Variable in Windows.
- If you see an error, see Confluence does not start due to Spring Application context has not been set for troubleshooting options.

Set up Confluence

6. Choose installation type

1. Choose**Production installation**.

2. Choose any**apps**you'd also like to install.

7. Enter your license

Follow the prompts to log in tomy.atlassian.comto retrieve your license, or enter a license key.

8. Connect to your database

1. If you've not already done so, it's time to create your database. See the 'Before you begin' section of this page for details and connection options.

2. For MySQL and Oracle, follow the prompts to download and install the required driver.

3. Enter your database details. Use**test connection** to check your database is set up correctly.

   If you want to specify particular parameters, you can choose to connect **By connection string**. You'll be prompted to enter:

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

- **Database URL** the JDBC URL for your database. If you're not sure, check the documentation for your database.
- **Username and Password** A valid username and password that Confluence can use to access your database.

9. Populate your new site with content

Choose whether you'd like Confluence to populate your site with content:

This option will create a space that you and your users can use to get to know Confluence. You can delete this space at any time.

Use this option if you have a full site export of an existing Confluence site. This is useful when youre migrating to another database or setting up a test site.

Good to know:

- You can only import sites from the **same** or **earlier** Confluence version.
- The system administrator account and all other user data and content will be imported from your previous installation.

In the setup wizard:

- **Upload a backup file** use this option if your site export file is small (25mb or less).
- **Restore a backup file from the file system** use this option if your backup file is large. Drop the file into your `<confluence-home>/restore`directory then follow the prompts to restore the backup.
- **Build Index** well need to build an index before your imported content is searchable. This can take a long time for large sites, so deselect this option if you would rather build the index later.Your content won't be searchable until the index is built.

10. Choose where to manage users

Choose to manage Confluence's users and groups inside Confluence or in a Jira application, such as Jira Software or Jira Service Management:

Choose this option if you're happy to manage users in Confluence, or don't have a Jira application installed.

Good to know:

- If you do plan to manage users in a Jira application, but have not yet installed it, we recommend installing Jira first, and then returning to the Confluence setup.
- You can add external user management (for example LDAP, Crowd or Jira) later if you choose.

Choose this option if you have a Jira application installed and want to manage users across both applications.

Good to know:

- This is a quick way of setting up your Jira integration with the most common options.
- It will configure a Jira user directory for Confluence, and set up application links between Jira and Confluence for easy sharing of data.
- You'll be able to specify exactly which groups in your Jira app should also be allowed to log in to Confluence. Your license tiers do not need to be the same for each application.
- You'll need either Jira 4.3 or later, Jira Core 7.0 or later, Jira Software 7.0 or later, or Jira Service Management 3.0 or later.

In the setup wizard:

- **Jira Base URL** the address of your Jira server, such as `http://www.example.com:8080/jira/` or `http://jira.example.com/`

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

- **Jira Administrator Login** this is the username and password of a user account that has the Jira System Administrator global permission in your Jira application. Confluence will also use this username and password to create a local administrator account which will let you access Confluence if Jira is unavailable. Note that this single account is stored in Confluence's internal user directory, so if you change the password in Jira, it will not automatically update in Confluence.
- **Confluence Base URL** this is the URL Jira will use to access your Confluence server. The URL you give here overrides the base URL specified in Confluence, for the purposes of connecting to the Jira application.
- **User Groups** these are the Jira groups whose members should be allowed to use Confluence. Members of these groups will get the 'Can use' permission for Confluence, and will be counted in your Confluence license. The default user group name differs depending on your Jira version:
  - Jira 6.4 and earlier: `jira-users`.
  - Jira Software 7.x and later: `jira-software-users`
  - Jira Core 7.x and later: `jira-core-users`
  - Jira Service Management (formerly Jira Service Desk) 3.x and later: `jira-servicedesk-users`
- **Admin Groups** provide one or more Jira groups whose members should have administrative access to Confluence. The default group is `jira-administrators`. These groups will get the system administrator and Confluence administrator global permissions in Confluence.

11. Create your administrator account

Enter details for the administrator account.

Skip this step if you chose to manage users in a Jira application or you imported data from an existing site.

12. Start using Confluence

That's it!Your Confluence site is accessible from a URL like this: http://<computer_name_or_IP_address>:<port>

If you plan to run Confluence behind a reverse proxy, check outProxy and SSL considerationsbefore you go any further.

Here's a few things that will help you get your team up and running:

- Set the server base URL this is the URL people will use to access Confluence.
- Set up a mail server this allows Confluence to send people notification about content.
- Add and invite users get your team on board!
- Start and stop Confluence find out how to start and stop Confluence.

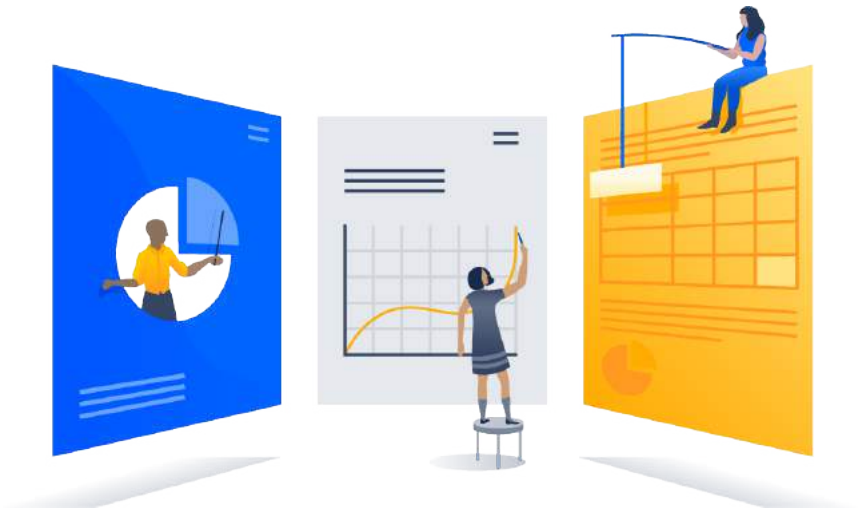

Troubleshooting

- If your web browser window shows an error the first time you try to access Confluence, wait for 30 seconds or so and then refresh the page.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

- If the command prompt window closes immediately, your JAVA_HOME variable may not be set correctly. See Setting the JAVA_HOME Variable in Windows.

- If you see an error, see Confluence does not start due to Spring Application context has not been set for troubleshooting options.
- Collaborative editing errors? See Troubleshooting Collaborative Editing.

Head to Installation Troubleshooting in our Knowledge Base for more help.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Uninstalling Confluence from Windows

This page describes the procedure for uninstalling an instance of Confluence which has been installed using the Windows Installer.

To uninstall Confluence from Windows:

1. Log in to Windows as the same user that was used to install Confluence with the Windows Installer.
2. Start the uninstaller by doing either of the following:
    - Click the Windows **Start Menu** > **All Programs** > **Confluence** > **Uninstall Confluence**
      OR
    - Open the Windows Control Panel, choose **Add or Remove Programs** (on Windows XP) or **Programs and Features** on (Windows 7, Vista) and then select **Confluence X.Y** from the list of applications and click **Uninstall/Change**.
      OR
    - Open the Windows command prompt and do the following:
        a. Change directory to your Confluence installation directory
        b. Run the `uninstall.exe` file
3. Follow the prompts to uninstall Confluence from your computer.

Please note:

- The uninstaller will not delete the Confluence Home Directory.
- All log files that were generated while Confluence was running will not be deleted.
- All files within the Confluence Installation Directory will be deleted (with the exception of the Tomcat `log` folder located in the Confluence Installation Directory).
- The uninstaller can be made to operate in unattended mode by specifying the `-q` option at the Windows command prompt i.e. `uninstall -q`
- If you wish to re-install Confluence in 'unattended mode', do not uninstall your previous installation of Confluence just yet. See Using the Silent Installation Feature for more information.

Installing Confluence on Linux

In this guide we'll run you through installing Confluence in a production environment, with an external database, using the Linux installer.

This is the most straightforward way to get your production site up and running on a Linux server.

On this page:

Before you begin
Install Confluence
1. Download Confluence
2. Run the installer
Set up Confluence
3. Choose installation type
4. Enter your license
5. Connect to your database
6. Populate your new site with content
7. Choose where to manage users
8. Create your administrator account
9. Start using Confluence
Troubleshooting

Other ways to install Confluence:

- Evaluation-get your free trial up and running in no time.
- TAR.GZ install Confluence manually from an archive file.
- Windows install Confluence on a Windows server.

Before you begin

Before you install Confluence, there are a few questions you need to answer.

Are you using a supported operating system?	Check the Supported Platforms page for the version of Confluence you are installing. This will give you info on supported operating systems, databases and browsers. **Good to know:** • We don't support installing Confluence on OSX for production sites. • The Confluence installer includes Java (JRE) and Tomcat, so you don't need to install these separately. • Confluence can only run on Oracle JDK or AdoptOpenJDK.
Does your Linux server have a font config package installed?	Many Linux distributions don't include a suitable font config package by default, so you will need to install one before you can run the Confluence installer. See Confluence Server 6.13 or later fails with FontConfiguration error when installing on Linux operating systems for commands to install a suitable package on several popular Linux distributions.

Do you want to run Confluence as a service?	Running Confluence as a service means that Confluence will automatically start up when Linux is started. **If you choose to run Confluence as a service:** • You must use `sudo` to run the installer to be able to install Confluence as a service. • The installer will create a dedicated user account, `confluence`, that will run the service. **If you choose not to run Confluence as a service:** • You will start and stop Confluence by running the `start-confluence.sh` file in your Confluence installation directory. • Confluence will be run as the user account that was used to install Confluence, or you can choose to run as a dedicated user. • Confluence will need to be restarted manually if your server is restarted.
Are ports 8090 and 8091 available?	Confluence runs on port 8090 by default. If this port is already in use, the installer will prompt you to choose a different port. Synchrony, which is required for collaborative editing, runs on port 8091 by default. If this port is already in use, you will need to change the port that Synchrony runs on after your Confluence installation is complete. See Administering Collaborative Editing to find out how to change the port Synchrony runs on.You won't be able to edit pages until Synchrony has an available port. SeePorts used by Atlassian Applicationsfor a summary of all the ports used.
Is your database set up and ready to use?	To run Confluence you'll need an external database. Check the Supported Platforms page for the version you're installing for the list of databases we currently support. If you don't already have a database, PostgreSQL is free and easy to set up. **Good to know:** • Set up your database before you begin. Step-by-step guides are available for PostgreSQL, Oracle, MySQL, and SQL Server. • If you're using Oracle or MySQL you'll need to download the driver for your database.
Do you have a Confluence license?	You'll need a valid license to use Confluence. **Good to know:** • If you have not yet purchased a Confluence license you'll be able to create an evaluation license during setup. • If you already have a license key you'll be prompted to log in to my.atlassian.com to retrieve it, or you can enter the key manually during setup. • If you're migrating from Confluence Cloud, you'll need a new license. • Weve ended sales for new server licenses and will end support for server on February 2, 2024. Were continuing our investment in Data Center.Learn more

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Do you want to store your attachment data on object storage?	By default, Confluence stores attachments in the home directory (e.g. in a file system). If your team has large or increasing data sets and requires the ability to scale efficiently, we recommend you use S3 object storage. Amazon S3 is currently the only Confluence-supported object storage solution. **Good to know:** • Amazon S3 object storage is an optional attachment storage method available to anyone on a Data Center license and running Confluence in AWS. • If you're a new customer, seeS3 object storagefor setup instructions. • If you're an existing customer, you'll need to migrate your attachment data to S3 object storage from the file system or another storage method. SeeAttachment storage configurationfor steps to do this. • Even if you use S3 object storage, other non-attachment data will still be stored in your home directory.

Theres a known issue during setup where a load balancer (or proxy) pings the server and breaks Confluence installation or migration to Data Center. See

CONFSERVER-61189 - Opening the base URL multiple times during Data Center migration will break the migration process. GATHERING IMPACT

During installation, you need to disable load balancer health checks and make sure you dont open multiple tabs that point to the same Confluence URL.

Install Confluence

1.Download Confluence

Download the installer for your operating system https://www.atlassian.com/software/confluence/download

2. Run the installer

1. Make the installer executable.

   Change to the directory where you downloaded Confluence then execute this command:

   ```
   $ chmod a+x atlassian-confluence-X.X.X-x64.bin
   ```

   Where `X.X.X` is is the Confluence version you downloaded.

2. Run the installer we recommend using`sudo`to run the installer as this will create a dedicated account to run Confluence and allow you to run Confluence as a service.

   To use `sudo` to run the installer execute this command:

   ```
   $ sudo ./atlassian-confluence-X.X.X-x64.bin
   ```

   Where `X.X.X` is is the Confluence version you downloaded.

   You can also choose to run the installer as with root user privileges.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

3. Follow the prompts to install Confluence. You'll be asked for the following info:

    - **Install type** choose option 2 (custom) for the most control.
    - **Destination directory** this is whereConfluencewill be installed.
    - **Home directory** this is where Confluence data like logs, search indexes and files will be stored.
    - **TCP ports** these arethe HTTP connector port and control port Confluence will run on. Stick with the default unless you're running another application on the same port.
    - **Install as service** this option is only available if you ran the installer as`sudo`.
4. Once installation is complete head tohttp://localhost:8090/in your browser to begin the setup process. (Replace`8090`if you chose a different port during installation).

If you're installing Confluence on a fresh Linux installation seeConfluence throws a Confluence is vacant error on install for troubleshooting options.

FontConfiguration error? See Confluence Server 6.13 or later fails with FontConfiguration error when installing on Linux operating systems to find out how to install a suitable font configuration package.

Set up Confluence

3. Choose installation type

1. Choose**Production installation**.
2. Choose any**apps**you'd also like to install.

4. Enter your license

Follow the prompts to log in tomy.atlassian.comto retrieve your license, or enter a license key.

5. Connect to your database

1. If you've not already done so, it's time to create your database. See the 'Before you begin' section of this page for details and connection options.
2. For MySQL and Oracle, follow the prompts to download and install the required driver.
3. Enter your database details. Use**test connection** to check your database is set up correctly.

    If you want to specify particular parameters, you can choose to connect **By connection string**. You'll be prompted to enter:

    - **Database URL** the JDBC URL for your database. If you're not sure, check the documentation for your database.
    - **Username and Password** A valid username and password that Confluence can use to access your database.

6. Populate your new site with content

Choose whether you'd like Confluence to populate your site with content:

This option will create a space that you and your users can use to get to know Confluence. You can delete this space at any time.

Use this option if you have a full site export of an existing Confluence site. This is useful when youre migrating to another database or setting up a test site.

Good to know:

- You can only import sites from the **same** or **earlier** Confluence version.
- The system administrator account and all other user data and content will be imported from your previous installation.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

In the setup wizard:

- **Upload a backup file** use this option if your site export file is small (25mb or less).
- **Restore a backup file from the file system** use this option if your backup file is large. Drop the file into your `<confluence-home>/restore`directory then follow the prompts to restore the backup.
- **Build Index** well need to build an index before your imported content is searchable. This can take a long time for large sites, so deselect this option if you would rather build the index later.Your content won't be searchable until the index is built.

7. Choose where to manage users

Choose to manage Confluence's users and groups inside Confluence or in a Jira application, such as Jira Software or Jira Service Management:

Choose this option if you're happy to manage users in Confluence, or don't have a Jira application installed.

Good to know:

- If you do plan to manage users in a Jira application, but have not yet installed it, we recommend installing Jira first, and then returning to the Confluence setup.
- You can add external user management (for example LDAP, Crowd or Jira) later if you choose.

Choose this option if you have a Jira application installed and want to manage users across both applications.

Good to know:

- This is a quick way of setting up your Jira integration with the most common options.
- It will configure a Jira user directory for Confluence, and set up application links between Jira and Confluence for easy sharing of data.
- You'll be able to specify exactly which groups in your Jira app should also be allowed to log in to Confluence. Your license tiers do not need to be the same for each application.
- You'll need either Jira 4.3 or later, Jira Core 7.0 or later, Jira Software 7.0 or later, or Jira Service Management 3.0 or later.

In the setup wizard:

- **Jira Base URL** the address of your Jira server, such as `http://www.example.com:8080/jira/` or `http://jira.example.com/`
- **Jira Administrator Login** this is the username and password of a user account that has the Jira System Administrator global permission in your Jira application. Confluence will also use this username and password to create a local administrator account which will let you access Confluence if Jira is unavailable. Note that this single account is stored in Confluence's internal user directory, so if you change the password in Jira, it will not automatically update in Confluence.
- **Confluence Base URL** this is the URL Jira will use to access your Confluence server. The URL you give here overrides the base URL specified in Confluence, for the purposes of connecting to the Jira application.
- **User Groups** these are the Jira groups whose members should be allowed to use Confluence. Members of these groups will get the 'Can use' permission for Confluence, and will be counted in your Confluence license. The default user group name differs depending on your Jira version:
  - Jira 6.4 and earlier: `jira-users`.
  - Jira Software 7.x and later: `jira-software-users`
  - Jira Core 7.x and later: `jira-core-users`
  - Jira Service Management (formerly Jira Service Desk) 3.x and later: `jira-servicedesk-users`
- **Admin Groups** provide one or more Jira groups whose members should have administrative access to Confluence. The default group is `jira-administrators`. These groups will get the system administrator and Confluence administrator global permissions in Confluence.

8. Create your administrator account

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Enter details for the administrator account.

Skip this step if you chose to manage users in a Jira application or you imported data from an existing site.

9. Start using Confluence

That's it!Your Confluence site is accessible from a URL like this: http://<computer_name_or_IP_address>:<port>

If you plan to run Confluence behind a reverse proxy, check outProxy and SSL considerationsbefore you go any further.

Here's a few things that will help you get your team up and running:

- Set the server base URL this is the URL people will use to access Confluence.
- Set up a mail server this allows Confluence to send people notification about content.
- Add and invite users get your team on board!
- Start and stop Confluence find out how to start and stop Confluence.

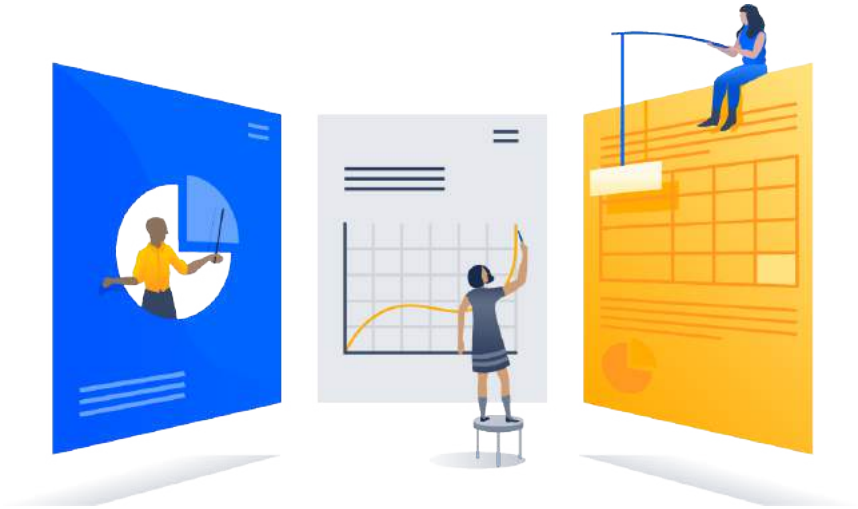

Troubleshooting

- If the installer fails with a FontConfiguration error, you'll need to install a font package. See Confluence Server 6.13 or later fails with FontConfiguration error when installing on Linux operating systems for info on how to do this.
- Some anti-virus or other Internet security tools may interfere with the Confluence installation process and prevent the process from completing successfully. If you experience or anticipate experiencing such an issue with your anti-virus/Internet security tool, disable this tool first before proceeding with theConfluence installation.
- The Linux OOM Killer can sometimes kill Confluence processes when memory on the server becomes too low. See How to Configure the Linux Out-of-Memory Killer.
- Collaborative editing errors? See Troubleshooting Collaborative Editing.

Head to Installation Troubleshooting in our Knowledge Base for more help.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Installing Confluence on Linux from Archive File

In this guide we'll run you through installing Confluence in a production environment, with an external database, manually using a zip file.

This method gives you the most control over the installation process.

Other ways to install Confluence:

- Evaluation-get your free trial up and running in no time.
- Installer install Confluence using the Linux installer.
- Windows install Confluence on a Windows server.

On this page:

- Before you begin
- Install Confluence
  - 1. Download Confluence
  - 2. Create the installation directory
  - 3. Create the home directory
  - 4. Check the ports
  - 5. Start Confluence
- Set up Confluence
  - 6. Choose installation type
  - 7. Enter your license
  - 8. Connect to your database
  - 9. Populate your new site with content
  - 10. Choose where to manage users
  - 11. Create your administrator account
  - 12. Start using Confluence
- Troubleshooting

Before you begin

Before you install Confluence, there are a few questions you need to answer.

Are you using a supported operating system and Java version?	Check the Supported Platforms page for the version of Confluence you are installing. This will give you info on supported operating systems, databases and browsers. **Good to know:** • We don't support installing Confluence on OS X or mac OS for production environments. • You'll need to install either Adoptium OpenJDK (formerly AdoptOpenJDK)or Oracle JDK. We don't support other OpenJDK binaries. • You can use either the JDK (Java Development Kit) or JRE (Java Runtime Environment). • We only support the version of Apache Tomcat that is bundled with Confluence.
Do you want to run Confluence as a service?	Running Confluence as a service means that Confluence will automatically start up when your Linux server is started. You should use the Linux installer if you want to run Confluence as a service. **If you choose not to run Confluence as a service:** • You will start Confluence by running the `start-confluence.sh` file in your Confluence installation directory. • We recommend creating a dedicated user to run Confluence. This user must have full read, write and execute access to the installation directory and home directory. • Confluence will need to be restarted manually if your server is restarted.

Are ports 8090 and 8091 available?	Confluence runs on port 8090 by default. If this port is already in use, the installer will prompt you to choose a different port. Synchrony, which is required for collaborative editing, runs on port 8091 by default. If this port is already in use, you will need to change the port that Synchrony runs on after your Confluence installation is complete. See Administering Collaborative Editing to find out how to change the port Synchrony runs on.You won't be able to edit pages until Synchrony has an available port. SeePorts used by Atlassian Applicationsfor a summary of all the ports used.
What database do you plan to use?	To run Confluence you'll need an external database. Check the Supported Platforms page for the version you're installing for the list of databases we currently support. If you don't already have a database, PostgreSQL is free and easy to set up. **Good to know:** • Set up your database before you begin. Step-by-step guides are available for PostgreSQL, Oracle, MySQL, and SQL Server. • If you're using Oracle or MySQL you'll need to download the driver for your database.
Do you have a Confluence license?	You'll need a valid license to use Confluence. **Good to know:** • If you have not yet purchased a Confluence license you'll be able to create an evaluation license during setup. • If you already have a license key you'll be prompted to log in to my.atlassian.com to retrieve it, or you can enter the key manually during setup. • If you're migrating from Confluence Cloud, you'll need a new license. • Weve ended sales for new server licenses and will end support for server on February 2, 2024. Were continuing our investment in Data Center.Learn more

Is your JAVA_HOME variable set correctly?	Before you install Confluence, check that you're running a supported Java version and that the `JAVA_HOME` environment variable is set correctly.
	
	Confluence can only run with Oracle JDK or JRE.
	
	To check your Java version:
	
	`$ java -version`
	
	To check your JAVA_HOME variable is set correctly:
	
	`$ echo $JAVA_HOME`
	
	If you see a path to your Java installation directory, the `JAVA_Home` environment variable has been set correctly. If a path is not returned you'll need to set your `JAVA_HOME` environment variable manually before installing Confluence.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Have you created a dedicated user to run Confluence?	We strongly recommend running Confluence as a dedicated user. You should create this user before you begin, so that when creating the installation and home directories, you can give this user appropriate read and write permissions. In this example, we'll create a user called `confluence`: `$ sudo /usr/sbin/useradd --create-home --comment "Account for running Confluence" --shell /bin/bash confluence` See Creating a Dedicated User Account on the Operating System to Run Confluence for more information.

⚠ Theres a known issue during setup where a load balancer (or proxy) pings the server and breaks Confluence installation or migration to Data Center. See

CONFSERVER-61189 - Opening the base URL multiple times during Data Center migration will break the migration process. GATHERING IMPACT

During installation, you need to disable load balancer health checks and make sure you dont open multiple tabs that point to the same Confluence URL.

Install Confluence

1.Download Confluence

Download the`tar.gz`file for your operating system -https://www.atlassian.com/software/confluence/download.

2. Create the installation directory

1. Create your installation directory this is where Confluence will be installed. Avoid using spaces or special characters in the path. We'll refer to this directory as your`<installation-directory>`.

   In this example we'll call our installation directory `confluence`:

   ```
   $ mkdir confluence
   ```

2. Extract the Confluence`tar.gz`file to your`<installation-directory>`. We recommend using aGNU version of the archive utility, especially on Solaris.

   Change to the directory where you downloaded Confluence then execute these commands:

   ```
   $ tar -xzf atlassian-confluence-X.X.X.tar.gz -C <installation-directory>
   $ cd <installation-directory>
   $ tar -xf atlassian-confluence-X.X.X.tar
   ```

   Replace `x.x.x` with your Confluence version and `<installation-directory>` with the full path to the directory you created in the last step.

3. Give your dedicated Confluence user read, write and execute permission toyour`<installation-directory>`.

   In this example we're changing ownership of the installation directory and giving the user `confluence` read, write and execute permissions.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

```
$ chown -R confluence <installation-directory>
$ chmod -R u=rwx,go-rwx <installation-directory>
```

3. Create the home directory

1. Create your home directory this is where Confluence application data like logs, search indexes and files will be stored. This should be separate to your installation directory, with nospaces or special characters in the path. We'll refer to this directory as your`<home-directory>`.

   In this example we'll call our home directory `confluence-home`:

   ```
   $ mkdir confluence-home
   ```

2. Give your dedicated Confluence user read, write and execute permissions to the`<home-directory>`.

   In this example we're changing ownership of the home directory and giving the user `confluence` read, write and execute permissions.

   ```
   $ chown -R confluence <home-directory>
   $ chmod -R u=rwx,go-rwx <home-directory>
   ```

3. Edit`<installation-directory>\confluence\WEB-INF\classes\confluence-init.properties`.
4. At the bottom of the file, enter the absolute path to your`<home-directory>`. This tells Confluence where to find your`<home-directory>`when it starts up.

   You can edit the `confluence.init.properties` file any text editor.

   a. Scroll to the bottom of the text and find this line:

      ```
      # confluence.home=c:/confluence/data
      ```

   b. Remove the # and the space at the beginning of this line (so Confluence doesn't read the line as a comment) and add the absolute path to your home directory (not a symlink). For example:

      ```
      confluence.home=/var/confluence-home
      ```

4. Check the ports

By default Confluence listens on port`8090`. If you have another application running on your server that uses the same ports, you'll need to tell Confluence to use a different port.

To change the ports:

1. Edit `<installation-directory>\conf\server.xml`
2. Change the **Server** port (8000) and the **Connector** port (8090) to free ports on your server.

   In the example below we've changed the **Server** port to 5000 and the **Connector** port to 5050.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

```
Server port="5000" shutdown="SHUTDOWN" debug="0">
  <Service name="Tomcat-Standalone">
    <Connector port="5050" connectionTimeout="20000" redirectPort="8443"
        maxThreads="48" minSpareThreads="10"
        enableLookups="false" acceptCount="10" debug="0" URIEncoding="UTF-8"
        protocol="org.apache.coyote.http11.Http11NioProtocol" />
```

Linux won't allow you to bind to ports less than 1024. If you want to run Confluence on port 80, for example, you could use a reverse proxy to redirect traffic from port 80. See Using Apache with mod_proxy.

5. Start Confluence

1. Run`<installation-directory>/bin/start-confluence.sh`to start the setup process.

   We recommend running Confluence as your dedicated user.

   ```
   $ su -u <user>
   $ ./start-confluence.sh
   ```

   If you're using Ubuntu the command is a little different:

   ```
   $ sudo su <user>
   $ ./start-confluence.sh
   ```

2. Go to`http://localhost:8090/`to launch Confluence in your browser (change the port if you've updated the Connector port).

   - Check your JAVA_HOME variable is set correctly.
   - If you see an error, see Confluence does not start due to Spring Application context has not been set for troubleshooting options.

Set up Confluence

6. Choose installation type

1. Choose**Production installation**.
2. Choose any**apps**you'd also like to install.

7. Enter your license

Follow the prompts to log in tomy.atlassian.comto retrieve your license, or enter a license key.

8. Connect to your database

1. If you've not already done so, it's time to create your database. See the 'Before you begin' section of this page for details and connection options.
2. For MySQL and Oracle, follow the prompts to download and install the required driver.
3. Enter your database details. Use**test connection** to check your database is set up correctly.

   If you want to specify particular parameters, you can choose to connect **By connection string**. You'll be prompted to enter:

   - **Database URL** the JDBC URL for your database. If you're not sure, check the documentation for your database.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

- **Username and Password** A valid username and password that Confluence can use to access your database.

9. Populate your new site with content

Choose whether you'd like Confluence to populate your site with content:

This option will create a space that you and your users can use to get to know Confluence. You can delete this space at any time.

Use this option if you have a full site export of an existing Confluence site. This is useful when youre migrating to another database or setting up a test site.

Good to know:

- You can only import sites from the **same** or **earlier** Confluence version.
- The system administrator account and all other user data and content will be imported from your previous installation.

In the setup wizard:

- **Upload a backup file** use this option if your site export file is small (25mb or less).
- **Restore a backup file from the file system** use this option if your backup file is large. Drop the file into your `<confluence-home>/restore`directory then follow the prompts to restore the backup.
- **Build Index** well need to build an index before your imported content is searchable. This can take a long time for large sites, so deselect this option if you would rather build the index later.Your content won't be searchable until the index is built.

10. Choose where to manage users

Choose to manage Confluence's users and groups inside Confluence or in a Jira application, such as Jira Software or Jira Service Management:

Choose this option if you're happy to manage users in Confluence, or don't have a Jira application installed.

Good to know:

- If you do plan to manage users in a Jira application, but have not yet installed it, we recommend installing Jira first, and then returning to the Confluence setup.
- You can add external user management (for example LDAP, Crowd or Jira) later if you choose.

Choose this option if you have a Jira application installed and want to manage users across both applications.

Good to know:

- This is a quick way of setting up your Jira integration with the most common options.
- It will configure a Jira user directory for Confluence, and set up application links between Jira and Confluence for easy sharing of data.
- You'll be able to specify exactly which groups in your Jira app should also be allowed to log in to Confluence. Your license tiers do not need to be the same for each application.
- You'll need either Jira 4.3 or later, Jira Core 7.0 or later, Jira Software 7.0 or later, or Jira Service Management 3.0 or later.

In the setup wizard:

- **Jira Base URL** the address of your Jira server, such as `http://www.example.com:8080/jira/` 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

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Confluence if Jira is unavailable. Note that this single account is stored in Confluence's internal user directory, so if you change the password in Jira, it will not automatically update in Confluence.

- **Confluence Base URL** this is the URL Jira will use to access your Confluence server. The URL you give here overrides the base URL specified in Confluence, for the purposes of connecting to the Jira application.
- **User Groups** these are the Jira groups whose members should be allowed to use Confluence. Members of these groups will get the 'Can use' permission for Confluence, and will be counted in your Confluence license. The default user group name differs depending on your Jira version:
  - Jira 6.4 and earlier: `jira-users`.
  - Jira Software 7.x and later: `jira-software-users`
  - Jira Core 7.x and later: `jira-core-users`
  - Jira Service Management (formerly Jira Service Desk) 3.x and later: `jira-servicedesk-users`
- **Admin Groups** provide one or more Jira groups whose members should have administrative access to Confluence. The default group is `jira-administrators`. These groups will get the system administrator and Confluence administrator global permissions in Confluence.

11. Create your administrator account

Enter details for the administrator account.

Skip this step if you chose to manage users in a Jira application or you imported data from an existing site.

12. Start using Confluence

That's it!Your Confluence site is accessible from a URL like this: http://<computer_name_or_IP_address>:<port>

If you plan to run Confluence behind a reverse proxy, check outProxy and SSL considerationsbefore you go any further.

Here's a few things that will help you get your team up and running:

- Set the server base URL this is the URL people will use to access Confluence.
- Set up a mail server this allows Confluence to send people notification about content.
- Add and invite users get your team on board!
- Start and stop Confluence find out how to start and stop Confluence.

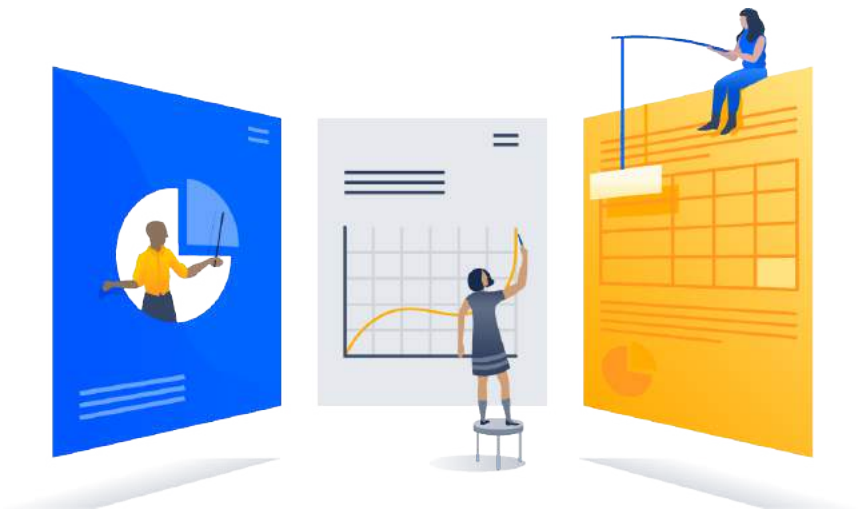

Troubleshooting

- Check your JAVA_HOME is set correctly.
- If you see an error, see Confluence does not start due to Spring Application context has not been set for troubleshooting options.
- Use a GNU version of the unzip utility. There are known issues extracting the `tar.gz` file on Solaris and AIX. See 'extractBundledPlugins Couldn't find atlassian-bundled-plugins.zip on classpath' Due to Solaris TAR Utility.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

- Collaborative editing errors? See Troubleshooting Collaborative Editing.

Head to Installation Troubleshooting in our Knowledge Base for more help.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Uninstalling Confluence from Linux

This page describes the procedure for uninstalling Confluence, which had been installed using the Linux Installer.

To uninstall Confluence from Linux:

1. Open a Linux console.
2. Change directory (`cd`) to your Confluence installation directory.
3. Execute the command `uninstall`. This command must be executed as the same user account that was used to install Confluence with the Linux Installer.
4. Follow the prompts to uninstall Confluence from your computer.

Please note:

- The uninstaller will not delete the Confluence Home Directory.
- All log files that were generated while Confluence was running will not be deleted.
- All files within the Confluence Installation Directory will be deleted (with the exception of the Tomcat `log` folder located in the Confluence Installation Directory).
- The uninstaller can be made to operate in unattended mode by specifying the `-q` option i.e. `uninstall -q`
- If you wish to re-install Confluence in 'unattended mode', do not uninstall your previous installation of Confluence just yet. See Using the Silent Installation Feature for more information.

Unattended installation

If you've previously installed Confluence using the Windows or Linux installer, you can use a configuration file from your existing Confluence installation (`response.varfile`) to re-install Confluence in unattended mode, no user input required.

This can be useful when you have installed Confluence on a test server and are ready to install on your production server with the same configuration.

Good to know

- The`response.varfile`file contains the options specified during the installation wizard steps of your previous Confluence installation. Don't uninstall your previous Confluence installation until after you've copied this file to your new install location.
- If you decide to modify the`response.varfile`file, make sure all directory paths specified are absolute, for example,`sys.installationDir=C\:\\Program Files\\Atlassian\\Confluence` (Windows) or`sys.installationDir=/opt/atlassian/confluence`(Linux).

  Unattended installations will fail the file contains relative directory paths.
- It's not possible to automate the database configuration step. This must be done via the setup wizard in your browser.

Install Confluence in unattended mode

These steps cover where you have an existing Confluence installation.

1. Downloadthe appropriate installer for your operating system.
2. Copy `<installation-directory>/.install4j/response.varfile`from your existing Confluence installation to where you downloaded the installer.
3. In command prompt or terminal change directory (cd) to where you downloaded the installer.
4. Run the following command to install Confluence:

   Windows

   ```
   > atlassian-confluence-X.X.X-x64.exe -q -varfile response.varfile
   ```

   Linux

   ```
   $ atlassian-confluence-X.X.X-x64.bin -q -varfile response.varfile
   ```

   Where`X.X.X`is the Confluence version you downloaded.

   `-q`instructs the installer to run in unattended mode (quietly). `-varfile`specifies the location and name of the configuration file containing the options used by the installer.
5. Confluence will start automatically once the silent installation finishes.

Once Confluence is installed, you will still need to head to `http://localhost:<port>`to finish setting up Confluence.

See the **Set up Confluence** section onInstalling Confluence on Windows orInstalling Confluence on Linux for more info.

Create your own response.varfile

It is also possible to create your own `response.varfile`, rather than one generated by an existing installation, if you are installing Confluence for the first time.

Example response.varfile

```
app.confHome=/var/atlassian/application-data/confluence6_15_5
app.install.service$Boolean=false
portChoice=custom
httpPort$Long=26112
rmiPort$Long=8001
launch.application$Boolean=false
sys.adminRights$Boolean=true
sys.confirmedUpdateInstallationString=false
sys.installationDir=/opt/atlassian/confluence6_15_5
sys.languageId=en
```

The following parameters can be included in the file.

Parameters	Accepted values	Description
app.confHome		This is the path to your target local home directory.
app.install.service$Boolean	• true • false	Determines whether Confluence should be installed as a service.
portChoice	• custom • default	Determines whether Confluence should be installed with default ports.
httpPort$Long		If portChoice is custom, this sets the HTTP connector port in Tomcat.
rmiPort$Long		If portChoice is custom, this sets the Tomcat server port.
launch.application$Boolean	• true • false	Determines whether the installer should start Confluence once installation is complete
sys.adminRights$Boolean=true	• true • false	Indicates whether the user running the installer has admin privileges on the machine.
sys.confirmedUpdateInstallationString	• true • false	Set this to `false` for a fresh unattended installation. Set to `true` to perform an unattended upgrade. Always back up your existing site before attempting to upgrade.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

sys.installationDir	path to install directory	This is the path to your target installation directory for a new install, or existing installation directory to be upgraded.
sys.languageId		Default application language.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

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, seeChanging JIRA application TCP ports.

[3] You will need to restart Confluence after editing`server.xml`for the changes to take effect.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Start and Stop Confluence

How you start and stop Confluence depends on whether you are running Confluence as a Service.

To check whether Confluence is already running you can go to`http://<base-url>/status`.

Windows

If you installed Confluence as a service, you can **Start Confluence Server** and **Stop Confluence Server** from the Windows Start menu.

You can't start or stop Confluence manually using the `start-confluence.bat` and `stop-confluence.bat` file.

If you didn't install Confluence as a service you'll need to start and stop Confluence manually. The way you do this depends on how Confluence was originally installed.

If you **installed Confluence manually**, and have Java installed on your server:

- To start Confluence run`<installation-directory>\bin\start-confluence.bat`
- To stop Confluence run`<installation-directory>\bin\stop-confluence.bat`

We recommend running Confluence with a dedicated user account. To do this, use use the`runas`command to execute`start-confluence.bat`.

```
> runas /env /user:<DOMAIN>\<confluence> start-confluence.bat
```

Where`<DOMAIN>`is your Windows domain or computer name and `<confluence>` is the name of your dedicated user.

If you **installed Confluence using the installer**, and don't have Java installed, use the Start and Stop Confluence options in the Start menu, or:

- To start Confluence run `<installation-directory>\startup-bundled-jre.bat`
- To stop Confluence run `<installation-directory>\shutdown-bundled-jre.bat`

It is possible to start Confluence Server with user installed apps temporarily disabled. This is useful if you need to troubleshoot problems with your site, particularly if an app may be preventing Confluence from starting up successfully.

To start Confluence with all user installed apps temporarily disabled:

```
> cd <installation-directory>/bin
> start-confluence.bat /disablealladdons
```

To start Confluence with a particular app temporarily disabled:

```
> cd <installation-directory>/bin
> start-confluence.bat /disableaddon=com.atlassian.test.plugin
```

where `com.atlassian.test.plugin` is the app key. To disable multiple apps, use a colon separated list. Regex/wildcards are not permitted, the full key of the plugin must be provided.

These parameters are applied at startup only, they do not persist. If you want to permanently disable an app, go to

> **Manage apps**

to do this via UPM.

Notes

- If the app key contains a space, disabling the app using this method will not work, you need to manually deal with that app.
- This feature does **not** work for Confluence Data Center.
- replace `/bin/start-confluence.bat` with `startup-bundled-jre.bat` if you installed Confluence using the installer, and are using the bundled JRE (Java Runtime Engine).

Linux

If you installed Confluence as a service, use one of the following commands to **start**, **stop** or **restart** Confluence.

```
$ sudo /etc/init.d/confluence start
$ sudo /etc/init.d/confluence stop
$ sudo /etc/init.d/confluence restart
```

You can't start or stop Confluence manually using the `start-confluence.sh` and `stop-confluence.sh` files.

If you didn't install Confluence as a service you'll need to start and stop Confluence manually.

- To start Confluence run`<installation-directory>\bin\start-confluence.sh`
- To stop Confluence run`<installation-directory>\bin\stop-confluence.sh`

We recommend running Confluence with a dedicated user account:

```
$ su -u <user>
$ ./start-confluence.sh
```

Where`<user>`is the name of your dedicated user.

If you're using Ubuntu the command is a little different:

```
$ sudo su <user>
$ ./start-confluence.sh
```

It is possible to start Confluence with user installed apps temporarily disabled. This is useful if you need to troubleshoot problems with your site, particularly if an app may be preventing Confluence from starting up successfully.

To start Confluence with **all** user installed apps temporarily disabled:

```
$ cd <installation-directory>/bin
$ ./start-confluence.sh --disable-all-addons
```

To start Confluence with a particular app temporarily disabled:

```
$ cd <installation-directory>/bin
$ ./start-confluence.sh --disable-addons=com.atlassian.test.plugin
```

where `com.atlassian.test.plugin` is the app key.

To disable multiple apps, use a colon separated list, for example, `com.atlassian.test.plugin:com.atlassian.another.plugin`. Regex/wildcards are not permitted, the full key of the plugin must be provided.

These parameters are applied at startup only, they do not persist. If you want to permanently disable an app, go to

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

> **Manage apps**

to do this via UPM.

Notes

- If the app key contains a space, disabling the app using this method will not work, you need to manually deal with that app.
- This feature does **not** work for Confluence Data Center.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Installing Confluence Data Center

In this guide we'll run you through installing Confluence Data Center in a Windows or Linux Environment. You can run Data Center as a standalone installation, or in a cluster, depending on your organisation's needs.

This guide covers installing for the first time, with no existing data.If you already have a Confluence Server instance, seeUpgrade from Confluence Server to Data Center.

On this page:

- Before you begin
  - Supported platforms
  - Requirements
- Install Confluence Data Center
- Install Confluence Data Center in a cluster
  - Terminology
  - Install and set up Confluence
  - Add more Confluence nodes
  - Security
  - Troubleshooting
  - Upgrading a cluster

Other ways to install Confluence Data Center:

- Kubernetes - install on a Kubernetes cluster using our Helm charts
- AWS- hassle free deployment in AWS using our Quick Start
- Azure - reference templates for Microsoft Azure deployment
- Move to Data Center- for existing Confluence Server sites

Interested in learning more about Data Center? Find out more about the benefits of Confluence Data Center.

Before you begin

Supported platforms

See ourSupported Platformspage for information on the database, Java, and operating systems you'll be able to use. These requirements are the same for Server and Data Center deployments.

Requirements

To use Confluence Data Center you must:

- Have a Data Center license (you canpurchase a Data Center licenseor create an evaluation license at my.atlassian.com)
- Use asupportedexternal database, operating system and Java version
- Use OAuth authentication if you haveapplication linksto other Atlassian products (such as Jira)

To run Confluence in a cluster you must also:

- Use a load balancer with session affinity in front of the Confluence cluster. WebSockets support is also recommended for collaborative editing.
- Have a shared directory accessible to all cluster nodes in the same path (this will be your shared home directory). This must be a separate directory, and not located within the local home or install directory.

Theres a known issue during setup where a load balancer (or proxy) pings the server and breaks Confluence installation or migration to Data Center. See

CONFSERVER-61189 - Opening the base URL multiple times during Data Center migration will break the migration process. GATHERING IMPACT

During installation, you need to disable load balancer health checks and make sure you dont open multiple tabs that point to the same Confluence URL.

Install Confluence Data Center

If your organization doesn't need high availability or disaster recovery capabilities right now, you can install Confluence Data Center without setting up a cluster.

To install Confluence Data Center, without setting up a cluster, follow the instructions for Confluence Server:

- Installing Confluence on Windows
- Installing Confluence on Linux

The process is almost identical to an ordinary Confluence Server installation, just be sure to choose **Standal one**after you've entered your Data Center license.

Install Confluence Data Center in a cluster

If your organization requirescontinuous uptime, scalability, and performance under heavy load, you'll want to run Confluence Data Center in a cluster.

SeeClustering with Confluence Data Centerfor a complete overview of hardware and infrastructure considerations.

Terminology

In this guide we'll use the following terminology:

- Installation directory The directory where you installed Confluence.
- Local home directory The home or data directory stored locally on each cluster node (if Confluence is not running in a cluster, this is simply known as the home directory).
- Shared home directory The directory you created that is accessible to all nodes in the cluster via the same path.

At the end of the installation process, you'll have an installation and*local*home directory on each node, and a single*shared*home directory (a total of 5 directories in a two node cluster) for Confluence plus directories for Synchrony.

Install and set up Confluence

1. Install Confluence on the first node

1. Install Confluence on node 1
   See Installing Confluence on Windows from Zip Fileor Installing Confluence on Linux from Archive File for more information.
2. Start Confluence on Node 1
3. Follow the prompts to enter your Data Center license then choose **Clustered** as the deployment type.
4. The setup wizard will prompt you to configure the cluster, by entering:
   - A name for your cluster
   - The path to the shared home directory you created earlier
   - The network interface Confluence will use to communicate between nodes
   - How you want Confluence to discover cluster nodes:
     - Multicast - enter your own multicast address or automatically generate one.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

- TCP/IP - enter the IP address of each cluster node
- AWS - enter your IAM Role or secret key, and region.

We recommend using our Quick Start or Cloud Formation Template to deploy Confluence Data Center in AWS, as it will automatically provision, configure and connect everything you need.

If you do decide to do your own custom deployment, you can provide the following information to allow Confluence to auto-discover cluster nodes:

Field	Description
IAM Role or **Secret Key**	This is your authentication method. You can choose to authenticate by IAM Role or Secret Key.
Region	This is the region your cluster nodes (EC2 instances) will be running in.
Host header	Optional. This is the AWS endpoint for Confluence to use (the address where the EC2 API can be found, for example 'ec2.amazonaws.com'). Leave blank to use the default endpoint.
Security group name	Optional. Use to narrow the members of your cluster to only resources in a particular security group (specified in the EC2 console).
Tag key and **Tag value**	Optional. Use to narrow the members of your cluster to only resources with particular tags (specified in the EC2 console).

5. Follow the prompts to set up your database and administrator account.
6. Confirm that you can log in to Confluence and everything is working as expected, then stop Confluence on Node 1.

Add more Confluence nodes

2.Copy Confluence to second node

To copy Confluence to the second node:

1. Shut down Confluence on node 1.
2. Copy the installation directory from node 1 to node 2.
3. Copy the local home directory from node 1 to node 2.

Copying the local home directory ensures the Confluence search index, the database and cluster configuration, and any other settings are copied to node 2.

3. Configure load balancer

Configure your load balancer for Confluence. You can use the load balancer of your choice, but it needs to support session affinity and WebSockets.

You can verify that your load balancer is sending requests correctly to your existing Confluence server by accessing Confluence through the load balancer and creating a page, then checking that this page can be viewed/edited by another machine through the load balancer.

4. Start Confluence one node at a time

You must only start Confluence **one node at a time**. The first node must be up and available before starting the next one.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

1. Start Confluence on node 1
2. Wait for Confluence to become available on node 1
3. Start Confluence on node 2
4. Wait for Confluence to become available on node 2.

TheCluster monitoring console (> **General Configuration** > **Clustering**) shows information about the active cluster.

When the cluster is running properly, this page displays the details of each node, including system usage and uptime. Use the ••• menu to see more information about each node in the cluster.

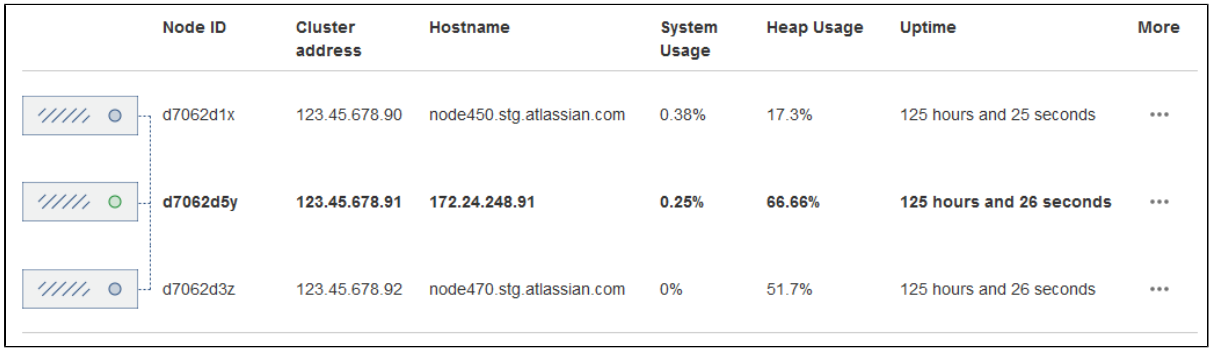

	Node ID	Cluster address	Hostname	System Usage	Heap Usage	Uptime	More
	d7062d1x	123.45.678.90	node450.stg.atlassian.com	0.38%	17.3%	125 hours and 25 seconds	•••
	d7062d5y	**123.45.678.91**	**172.24.248.91**	**0.25%**	**66.66%**	**125 hours and 26 seconds**	•••
	d7062d3z	123.45.678.92	node470.stg.atlassian.com	0%	51.7%	125 hours and 26 seconds	•••

5. Test your Confluence cluster

To test creating content you'll need to access Confluence via your load balancer URL. You can't create or edit pages when accessing a node directly.

A simple process to ensure your cluster is working correctly is:

1. Access a node via your load balancer URL, and create a new document on this node.
2. Ensure the new document is visible by accessing it directly on a different node.
3. Search for the new document on the original node, and ensure it appears.
4. Search for the new document on another node, and ensure it appears.

> If Confluence detects more than one instance accessing the database, but not in a working cluster, it will shut itself down in a *cluster panic*. This can be fixed by troubleshooting the network connectivity of the cluster.

6. Set up your Synchrony cluster (optional)

Synchrony is required for collaborative editing. You have two options for running Synchrony with a Data Center license:

- **managed by Confluence**(recommended)
  This is the default setup. Confluence will automatically launch a Synchrony process on the same node, and manage it for you. No manual steps are required.
- **Standalone Synchrony cluster** (managed by you)
  You deploy and manage Synchrony standalone in its own cluster with as many nodes as you need. Significant setup is required.See Set up a Synchrony cluster for Confluence Data Center for a step-by-step guide.

Head to Administering Collaborative Editing to find out more about collaborative editing.

Security

Ensure that only permittedcluster nodes are allowed to connect to the following portsthrough the use of a firewall and / or network segregation:

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

- 5801 -Hazelcastport for Confluence
- 5701 -Hazelcastport for Synchrony
- 25500 - Cluster base port for Synchrony

If you use multicast for cluster discovery:

- 54327- Multicast port for Synchrony (only required if running Synchrony standalone cluster)

Troubleshooting

If you have problems with the above procedure, please see ourCluster Troubleshooting guide.

If you're testing Confluence Data Center by running the cluster on a single machine, please refer to our developer instructions onStarting a Confluence cluster on a single machine.

Upgrading a cluster

It's important that upgrades follow the procedure for Upgrading Confluence Data Center.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Upgrading Confluence Data Center

This page contains instructions for **upgrading an existing Confluence cluster**.

If you are not running Confluence in a cluster, follow the instructions in Upgrading Confluence.

If you're running Confluence in a cluster in AWS, follow the instructions inRunning Confluence Data Center in AWS.

If you are upgrading to the next bug fix update (for example, 7.9.0 to 7.9.3), you can do so with no downtime. Follow the instructions inUpgrade Confluence without downtime.

In this guide we'll use the following terminology:

- Installation directory The directory where you installed Confluence.
- Local home directory The home or data directory stored locally on each cluster node (if Confluence is not running in a cluster, this is simply known as the home directory).
- Shared home directory The directory you created that is accessible to all nodes in the cluster via the same path.

Currently using Confluence Server? Learn more about thebenefits of Confluence Data Center.

On this page:

- 1. Back up
- 2. Download Confluence
- 3. Stop the cluster
- 4. Upgrade the first node
- 5. Upgrade Synchrony (optional)
- 6. Copy Confluence to remaining nodes
- 7. Start Confluence and check cluster connectivity

1. Back up

We strongly recommend that you backup your Confluence home and install directories and your database before proceeding.

More information on specific files and directories to backup can be found inUpgrading Confluence.

2. Download Confluence

Download the appropriate file for your operating system fromhttps://www.atlassian.com/software/confluence/download

3. Stop the cluster

You must stop all the nodes in the cluster before upgrading.

We recommend configuring your load balancer to redirect traffic away from Confluence until the upgrade is complete on all nodes.

4. Upgrade the first node

To upgrade the first node:

1. Extract (unzip) the files to adirectory (this will be your new installation directory, and must be different to your existing installation directory)
2. Update the following line in the`<Installation-Directory>\confluence\WEB-INF\classes\confluence-init.properties`file to point to the**existinglocal home**directory on that node.

3. If your deployment uses a MySQL database, copy the jdbc driver jar file from your existing Confluence installation directory to `confluence/WEB-INF/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. See Database Setup For MySQL for more details.
4. If you run Confluence as a service:
    - On Windows, delete the existing service then re-install the service by running `<install-directory>/bin/service.bat`.
    - On Linux, update the service to point to the new installation directory (or use symbolic links to do this).
5. Copy any other immediately required customizations from the old version to the new one (for example if you are not running Confluence on the default ports or if you manage users externally, you'll need to update / copy the relevant files - find out more in Upgrading Confluence Manually).

    > If you configured Confluence to run as a Windows or Linux service, don't forget to update its service configuration as well. For related information, see Start Confluence Automatically on Windows as a Service or Run Confluence as a systemd service on linux.

6. Start Confluence, and confirm that you can log in and view pages before continuing to the next step.

You should now stop Confluence, and reapply any additional customizations from the old version to the new version, before upgrading the remaining nodes.

5. Upgrade Synchrony (optional)

If you've chosen to let Confluence manage Synchrony for you (recommended), you don't need to do anything. Synchrony was automatically upgraded with Confluence.

If you're running your own Synchrony cluster, you should:

1. Grab the new `synchrony-standalone.jar` from the `<local-home>` directory on your upgraded Confluence node.
2. Copy the new `synchrony-standalone.jar` to each of your Synchrony nodes, and start Synchrony as normal.

6. Copy Confluence to remaining nodes

The next step is to replicate your upgraded Confluence directories to other nodes in the cluster.

1. Copy the installation directory and local home directory from the first node to the next node.
2. If the path to the local home directory is different on this node, edit the `confluence-init.properties` to point to the correct location.
3. Start Confluence, and and confirm that you can log in and view pages on this node.

Stop Confluence on this node, then repeat this process for each remaining node.

7. Start Confluence and check cluster connectivity

Once all nodes have been upgraded you can start Confluence Data Center on each node, **one at a time** (starting up multiple nodes simultaneously can lead to serious failures).

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.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Adding and Removing Data Center Nodes

Your Data Center license is based on the number of users in your cluster, rather than the number of nodes. This meansyou can add and remove nodes from your Data Center cluster at any time.

If you deployed Confluence Data Center on AWS using theQuick Start, your Confluence and Synchrony nodes will be inauto-scaling groups. You will add and remove nodes in the AWS console either by changing the minimum and maximum size of each group or using a scaling plan.

On this page:

- Adding a node
- Removing a node
- Changing the node identifier
- Moving to a non-clustered installation

Adding a node

To add a node:

1. Copy the installation directory and local home directory from the stopped node to your new node.
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.

If the discovery mode is set to TCP/IP, you'll need to update the `confluence.cluster.peers` property in the `confluence.cfg.xml` file for each node so the file lists all nodes in your cluster:

```
<property name="confluence.cluster.peers">[node 1 IP],[node 2 IP],[node 3 IP]</property> <!-- A comma-separated list of node IP addresses, without spaces -->
```

Removing a node

To remove a node, stop Confluence on that node. You can then remove the installation and local home directory as required.

To see the number of nodes remaining go to > **General Configuration**>**Clustering**.

Changing the node identifier

Confluence generates an identifier for each node in your cluster. You can use the`confluence.cluster.node.name`system propertyto set the node identifier on each node so that it's easier for your users and administrators to read.

SeeConfiguring System Properties for more information on how to set the system property.

Moving to a non-clustered installation

If you no longer need clustering, and want to avoid the overhead that comes from running a cluster with just one node, you can go back to a non-clustered Data Center installation. You'll need to make some infrastructure changes as part of the switch.

SeeMove to a non-clustered installationto find out how to do this.

Change Node Discovery from Multicast to TCP/IP or AWS

On this page:

- To change from multicast to TCP/IP
- To change from multicast to AWS
- To change from TCP/IP to AWS
- To change from TCP/IP to multicast
- Reference of properties in the confluence.cfg.xml file

If you're setting up Confluence Data Center for the first time, it'll step you through the process of choosing your discovery mode and adding cluster nodes.If you decide to change the node discovery for the cluster, you'll need to edit the`confluence.cfg.xml`file in thelocalhome 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` filein each local home directory)
- Always perform a safety backup before making manual edits to these files

The changes you need to make may differ slightly, depending on whether you've upgraded from an older version of Confluence Data Center or if you've started with version 5.9. We've detailed both methods, below.

To change from multicast to TCP/IP

Look for the following two lines in 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`propert y exists, butthere'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. Please, make sure to remove the brackets from around the IP addresses.

You can now restart your cluster nodes.

To change from multicast to AWS

Look for the following two lines in the**`confluence.cfg.xml`**file and remove them:

```
<property name="confluence.cluster.address">[multicast IP]</property>
<property name="confluence.cluster.join.type">multicast</property>
```

Depending on which type of credentials you are passing to Confluence, you will add*one*of the following two blocks with your AWS configuration.

Option 1: For Access Key/Secret Key based credentials:

```
<property name="confluence.cluster.join.type">aws</property>
<property name="confluence.cluster.aws.host.header">[---VALUE---]</property>
<property name="confluence.cluster.aws.region">[---VALUE---]</property>
<property name="confluence.cluster.aws.tag.key">[---VALUE---]</property>
<property name="confluence.cluster.aws.tag.value">[---VALUE---]</property>
<property name="confluence.cluster.aws.access.key">[---VALUE---]</property>
<property name="confluence.cluster.aws.secret.key">[---VALUE---]</property>
```

Option 2: For IAM role based credentials:

```
<property name="confluence.cluster.join.type">aws</property>
<property name="confluence.cluster.aws.host.header">[---VALUE---]</property>
<property name="confluence.cluster.aws.region">[---VALUE---]</property>
<property name="confluence.cluster.aws.tag.key">[---VALUE---]</property>
<property name="confluence.cluster.aws.tag.value">[---VALUE---]</property>
<property name="confluence.cluster.aws.iam.role">[---VALUE---]</property>
```

To change from TCP/IP to AWS

Look for the following two lines in the**`confluence.cfg.xml`**file and remove them:

```
<property name="confluence.cluster.join.type">tcp_ip</property>
<property name="confluence.cluster.peers">[node 1 IP],[node 2 IP],[node 3 IP]</property>
```

Depending on which type of credentials you are passing to Confluence, you will add*one*of the following two blocks with your AWS configuration.

Option 1: For Access Key/Secret Key based credentials:

```
<property name="confluence.cluster.join.type">aws</property>
<property name="confluence.cluster.aws.host.header">[---VALUE---]</property>
<property name="confluence.cluster.aws.region">[---VALUE---]</property>
<property name="confluence.cluster.aws.tag.key">[---VALUE---]</property>
<property name="confluence.cluster.aws.tag.value">[---VALUE---]</property>
<property name="confluence.cluster.aws.access.key">[---VALUE---]</property>
<property name="confluence.cluster.aws.secret.key">[---VALUE---]</property>
```

Option 2: For IAM role based credentials:

```
<property name="confluence.cluster.join.type">aws</property>
<property name="confluence.cluster.aws.host.header">[---VALUE---]</property>
<property name="confluence.cluster.aws.region">[---VALUE---]</property>
<property name="confluence.cluster.aws.tag.key">[---VALUE---]</property>
<property name="confluence.cluster.aws.tag.value">[---VALUE---]</property>
<property name="confluence.cluster.aws.iam.role">[---VALUE---]</property>
```

You can now restart your cluster nodes.

Note that if you're using a CloudFormation YAML template you need to make sure you have these appropriate values as a minimum and they should be reflected on the AWS side as well. If you switch to AWS mode cluster type, please also reviewRunning Confluence Data Center in AWSand make sure you have the following set up in your YAML:

```
Key: Cluster
Value: !Ref AWS::StackName
PropagateAtLaunch: true
```

To change from TCP/IP to multicast

To switch from TCP/IP to multicast, just perform the reverse of the changes outlined above.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Reference of properties in the confluence.cfg.xml file

key	valid values	notes
`confluence.cluster.join.type`	`'multicast'` or `'tcp_ip'`or `'aws'`	Pre-5.9 Data Center installations won't have this key. By default, if the key is missing, Confluence will choose `multicast`
`confluence.cluster.address`	a single multicast IP address	This key is only used by confluence if `confluence.cluster.join.type` is set to `multicast`
`confluence.cluster.peers`	a comma-separated string of IP addresses (no spaces)	There must be at least one address here. The addresses are the IP address of each node in the cluster, for example `<property name="confluence.cluster.peers">[node 1 IP],[node 2 IP],[node 3 IP]</property>` This key is only used by confluence if `confluence.cluster.join.type` is set to `tcp_ip`
`confluence.cluster.authentication.enabled`	true, false	Set this property to false if you don't want to authenticate Confluence nodes as they join the cluster. This is not recommended.
`confluence.cluster.authentication.secret`	(automatically generated)	Set this property to change the shared secret used to authenticate nodes as they join the cluster. The secret must be a string of maximum 40 characters.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Running Confluence Data Center in AWS

The AWS Quick Start template as a method of deployment is nearing its end-of-support date on Februar y 29, 2024. You can still use the template after this date but we won't maintain or update it. Learn more about the AWS deployment template

We recommend deploying your Data Center products on a Kubernetes cluster using our Helm charts for a more efficient and robust infrastructure and operational setup. Learn more about deploying on Kubernetes

AWS now recommends switching launch configurations, which our AWS Quick Start template uses, to la unch templates. We wont do this switch, since were ending our support for the AWS Quick Start template. Youll still be able to create launch configurations until December 31, 2023.

If you decide to deploy your Data Center instance in a clustered environment, consider usingAmazon Web Services (AWS). AWS allows you to scale your deployment elastically by resizing and quickly launching additional nodes, and provides a number of managed services that work with Data Center products. These services make it easier to configure, manage, and maintain your deployment's clustered infrastructure.

We recommend deploying your Data Center instance on a Kubernetes cluster using our Helm charts. This allows you to stay in control of your data and meet your compliance needs while still using a modern infrastructure. Learn more about running Data Center products on Kubernetes

Interested in learning more about what Data Center provides? Check out the Data Center overview

Non-clustered VS clustered environment

A single node is adequate for most Small or Medium size deployments, unless you need high availability or zero-downtime upgrades.

If you have an existing Server installation, you can still use its infrastructure when you upgrade to Data Center. Many features exclusive to Data Center (likeSAML single sign-on,self-protection via rate limiting,andCDN support) don't require clustered infrastructure.You can start using these Data Center features by simply upgrading your Server installations license.

For more information on whether clustering is right for you, check outAtlassian Data Center architecture and infrastructure options

Deploying Data Center products in a cluster using the AWS EKS

You can deploy your Data Center instance using a managed Kubernetes cluster service. Learn how to prepare a Kubernetes cluster using Amazon EKS

Here's an overview of the architecture for a Data Center instance running in Kubernetes:

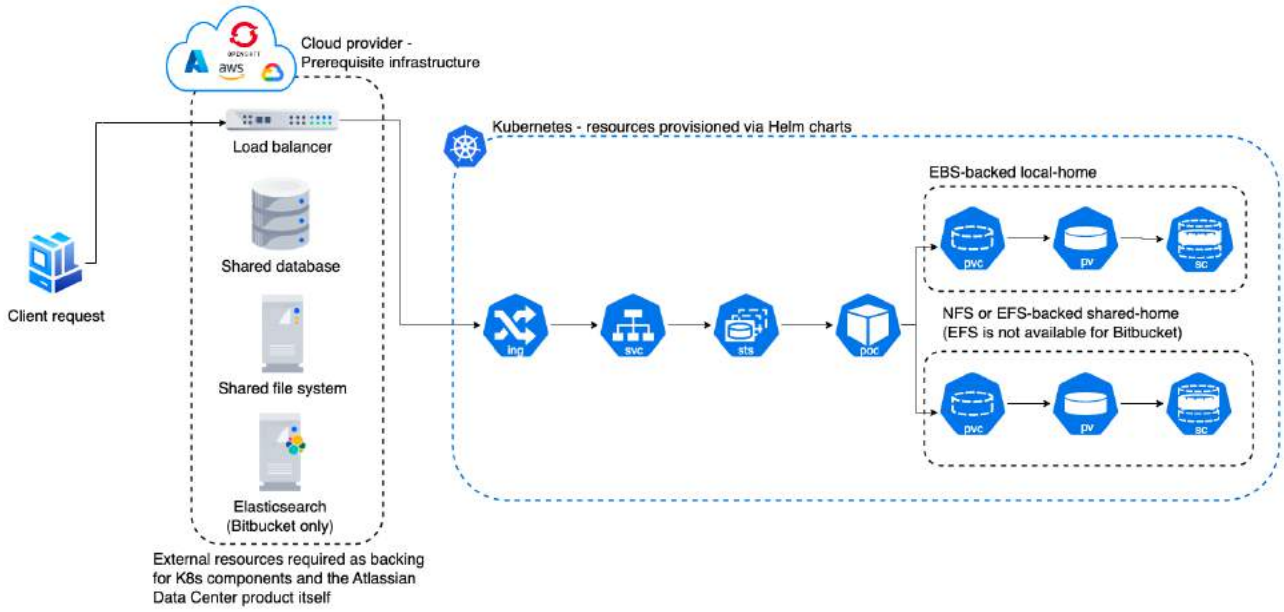

For more information, seeAtlassian products on AWS.

> ⚠ Even though you can deploy our Data Center products on AWS GovCloud, we dont test or verify our Helm charts on the AWS GovCloud environment and cant provide any support.

Deploy your instance with AWS

Create components

Before you deploy your Data Center product with AWS, you need to create the required infrastructure components. These include a database, a Kubernetes cluster, and shared storage. Learn more about the prerequisites

Take advantage of Helm charts

If you decide to deploy your Data Center instance on AWS with Kubernetes, make sure to use our Helm charts. Learn how to install your Data Center product with Helm charts

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Getting started with Confluence Data Center on Azure

> ⚠ The Azure Resource Manager template as a method of deployment is no longer supported or maintained by Atlassian. You can still customize it for your own usage to deploy Data Center products on Azure though.
>
> We recommend deploying your Data Center products on a Kubernetes cluster using our Helm charts for a more efficient and robust infrastructure and operational setup. Learn more about deploying on Kubernetes

If you decide to deploy your Data Center instance in a clustered environment, consider using Microsoft Azure. This platform allows you to scale your deployment elastically by resizing and quickly launching additional nodes and provides a number of managed services that work out of the box with Data Center products. These services make it easier to configure, manage, and maintain your deployment's clustered infrastructure.

We recommend deploying your Data Center instance on a Kubernetes cluster using our Helm charts. This allows you to stay in control of your data and meet your compliance needs while still using a modern infrastructure. Learn more about running Data Center products on Kubernetes

> ⓘ Interested in learning more about what Data Center provides? Check out the Data Center overview

Non-clustered VS clustered environment

A single node is adequate for most small or medium size deployments, unless you need high availability orzero-downtime upgrades.

If you have an existing Server installation, you can still use its infrastructure when you upgrade to Data Center. Many features exclusive to Data Center (likeSAML single sign-on,self-protection via rate limiting,andCDN support) don't require clustered infrastructure.You can start using these Data Center features by simply upgrading your Server installations license.

For more information on whether clustering is right for you, check outAtlassian Data Center architecture and infrastructure options.

How it works

Here's an architectural overview of what you'll get when deployingData Center products with Azure:

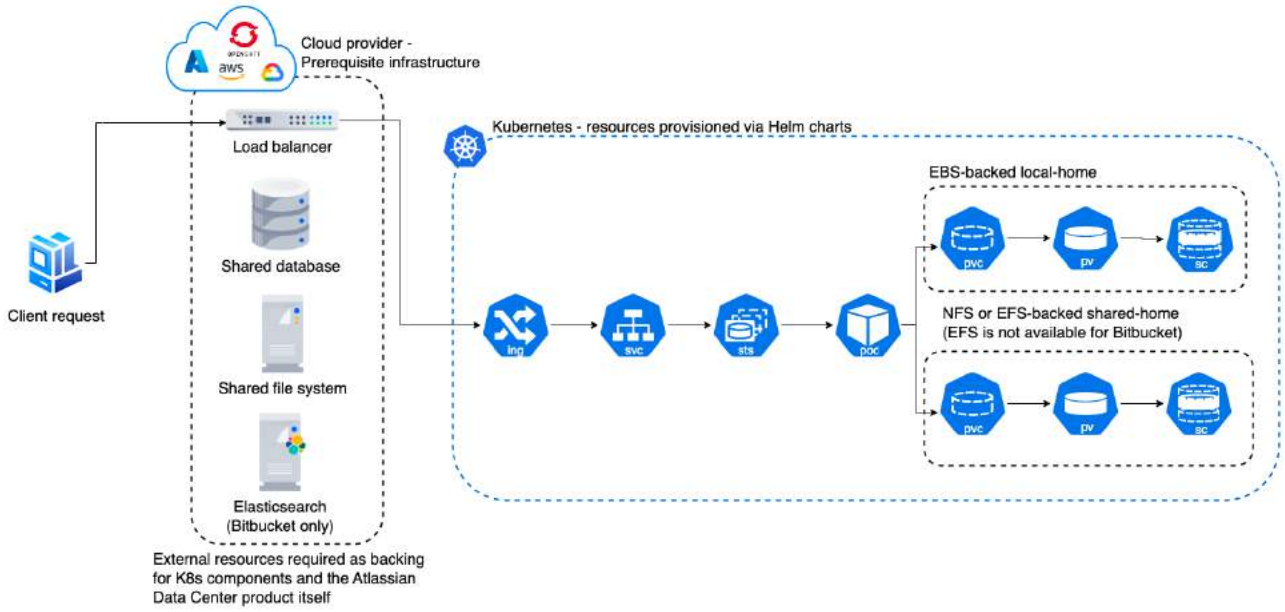

Deploy your instance with Azure

Create components

Before you deploy your Data Center product with Azure, you need to create the required infrastructure components. These include a database, a Kubernetes cluster, and shared storage. Learn more about the prerequisites

Take advantage of Helm charts

If you decide to deploy your Data Center instance on Azure with Kubernetes, make sure to use our Helm charts. Learn how to install your Data Center product with Helm charts

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Administering Confluence Data Center on Azure

The Azure Resource Manager template as a method of deployment is no longer supported or maintained by Atlassian. You can still customize it for your own usage to deploy Data Center products on Azure though.

We recommend deploying your Data Center products on a Kubernetes cluster using our Helm charts for a more efficient and robust infrastructure and operational setup. Learn more about deploying on Kubernetes

Once you've deployed Confluence Data Center to Azure using the deployment template, administering the application is similar to managing an application on your own hardware, with the exception that you'll need to go via the bastion host (jumpbox) to access your nodes.

To access your jumpbox and nodes you'll need:

- the SSH credentials you provided during setup,
- the Confluence node credentials you provided during setup
- the public DNS name or IP address of your jumpbox (you can obtain this through the Azure portal via**Menu**>**Resource groups**><**your resource group**>>**confluencenat**), and
- the node IP addresses, listed against the`confluencecluster (instance n)`row in**Connected devices**.(You can obtain this through the Azure portal via**Menu**>**Resource groups**><**your resource group**>>**confluencevnet**).

Connecting to your Azure jumpbox over SSH

You can SSH into yourConfluencecluster nodes, Synchrony nodes and shared home directory to perform configuration or maintenance tasks. Note that you must keep your SSH public key filein a safe place. This is the key to your jumpbox, and therefore all the nodes in your instance.

Access the jumpbox via a terminal or command line using:

```
$ ssh JUMPBOX_USERNAME@DNS_NAME_OR_IP_ADDRESS
```

You can find the SSH URL in the outputs section of your deployment.

Once you've accessed the jumpbox,we can jump to any of the nodes in the cluster, using:

```
$ ssh NODE_USERNAME@NODE_IP_ADDRESS
```

You'll then be asked for your node password - after providing this, you should be connected to the node.

Accessing your configuration files

For your Azure deployment, you may need to make changes to some configuration files, just as you would for a deployment on your own hardware:

- your`server.xml`lives in`/opt/atlassian/confluence/conf`
- your`setenv.sh`lives in`/opt/atlassian/confluence/bin`
- your local home`confluence.cfg.xml`lives in`/var/atlassian/application-data/confluence`
- your shared home`confluence.cfg.xml`lives in`/media/atl/confluence/shared`

These files are only accessible from the existing nodes. The*shared home*is mounted (think of it as a network hard disk) on each node under`/media/atl/confluence/shared`. So from an existing node (when you're logged in through SSH), you can go to`/media/atl/confluence/shared`.

If modifications to these files are made manually, new nodes will not pick up those modifications. You can either repeat the modifications on each node, or change the templates in the `/media/atl/confluence/shared` directory from which those files are derived. The mappings are:

- the `server.xml` file is derived from `/media/atl/confluence/shared/server.xml`
- the `setenv.sh` file is derived from `/media/atl/confluence/shared/setenv.sh`
- the local home `confluence.cfg.xml` is derived from `/media/atl/confluence/shared/home-confluence.cfg.xml`
- the shared home `confluence.cfg.xml` is derived from `/media/atl/confluence/shared/shared-confluence.cfg.xml`

These template files contain placeholders for values that are injected via the deployment script. Removing or changing them may cause breakages with the deployment. In most cases, these files should not be modified, as a lot of these settings are produced from the Azure Resource Manager templates automatically.

Upgrading

Consider upgrading to a Long Term Support release (if you're not on one already). Enterprise releases get fixes for critical bugs and security issues throughout its two-year support window. This gives you the option to keep a slower upgrade cadence without sacrificing security or stability. Long Term Support releases are suitable for companies who can't keep up with the frequency at which we ship feature releases.

Here's some useful advice for upgrading your deployment:

1. Before upgrading to a later version of Confluence Data Center, check if your apps are compatible with that version. Update your apps if needed. For more information about managing apps, see Using the Universal Plugin Manager.
2. If you need to keep Confluence Data Center running during your upgrade, we recommend using read-only mode for site maintenance. Your users will be able to view pages, but not create or change them.
3. We strongly recommend that you perform the upgrade first in a *staging* environment before upgrading your production instance. Create a staging environment for upgrading Confluence provides helpful tips on doing so.

Rolling upgrades

As of Confluence Data Center 7.9, you can now upgrade to the next bug fix version (for example, 7.9.0 to 7.9.3) with no downtime. Follow the instructions in Upgrade Confluence without downtime.

Upgrading Confluence in Azure

The process of upgrading Confluence is the same as if you were running the cluster on your own hardware. You will stop Confluence on all nodes, upgrade one node, stop that node then copy the installation directory across to each remaining node in the cluster, before restarting each node, one at a time.

See Upgrading Confluence Data Center for more details.

You can't use the `confluenceVersion` parameter in the deployment template to upgrade an existing Confluence deployment, or to provision new nodes running a different version to the rest of your cluster.

You also can't do a rolling upgrade. You will need to bring all nodes down before upgrading.

Upgrading your operating system

If you need to upgrade the operating system running on your Confluence nodes, you will need to SSH into each node, perform `sudo apt dist-upgrade` (Ubuntu) and reboot each node.

As Confluence is running as a service it will be automatically restarted on reboot.

You can't simply reimage an instance, as you might do in Jira, due to the way Hazelcast discovers cluster nodes.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Backing up and recovering from failures

We recommend you use the Azure native backup facilities where possible to make sure your data is backed up, and you can easily recover in the case of a failure.

Database backups

We use Azure-managed database instances with high availability. Azure provides several excellent options for backing up your database, so you should take some time to work out which will be the best, and most cost effective option for your needs. See the following Azure documentation for your chosen database:

- SQL Database: Automated backups
- SQL Database: Backup retention
- PostGreSQL: Backup concepts

Shared home backups

The shared home stores your your attachments, profile pictures, and export files. We create a general purpose Azure storage account, configured withlocal redundant storage(LRS), which means there are multiple copies of the data at any one time.

LRS provides a basic redundancy strategy for your shared home. As such, it shouldn't be necessary to take regular backups yourself. If you need to take point-in-time backups, usesnapshots.

Application nodes

The application nodes are VMs in an Azure Virtual Machine Scale Set. Each application node has a Confluence installation directory and a local home directory containing things like logs and search indexes.

Like the shared home, application nodes are configured withlocal redundant storage.This means there are multiple copies of the data at any one time.

If you've manually customised any configuration files in the installation directory (for example velocity templates), you may also want to manually back these up as a reference.

Bastion host

As this VM acts as a jumpbox, and doesn't store any data it doesn't need to be backed up. If the VM becomes unresponsive it can be restarted from the Azure Portal.

Application gateway

The application gateway is highly available. We deploy 2 instances by default. As with the bastion host, it doesn't need to be backed up.

Disaster recovery

SeeConfluence Data Center disaster recovery to learn about how you can develop a disaster recovery strategy. See also information in the Azure documentation about recovering from a region-wide failureAzure resiliency technical guidance: recovery from a region-wide service disruption.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Running Confluence Data Center on a Kubernetes cluster

If you're running self-managed environments and looking to adopt modern infrastructure, you can deploy your Atlassian Data Center products on Kubernetes clusters. By leveraging Kubernetes, you can drive greater agility amongst your teams and enjoy a simplified administrative experience at scale without compromising your organizations regulatory requirements.

What is Kubernetes?

Kubernetes (K8s) is a high-availability rapid deployment and container orchestration framework that allows you to easily manage and automate your deployments in one place. Because it makes heavy use of containers, your applications stay up-to-date with no downtime and always have safe and reliable access to the dependencies they require. Learn more on the official Kubernetes website

How does it work?

Kubernetes automates the management of containerized applications. It provides a centralized control plane to manage containers and the underlying infrastructure, automate scaling, rollouts and rollbacks, and more. The platform abstracts away the underlying infrastructure and provides a unified way of managing containers and applications, making it easier for developers to build, deploy, and run applications at scale.

Why Kubernetes?

Kubernetes is a powerful platform that comes with a number of benefits, including:

- Improved agility
- Simplified administration
- Deployment automation
- Automated operations for containers
- Security enhancements
- Accelerated upgrades and rollbacks
- Better scalability and resiliency

On top of that, the ability to manage your infrastructure as code by using simple YAML files helps you reduce unnecessary resource consumption.

How does Atlassian integrate with Kubernetes?

Manage Kubernetes with Helm charts

To help you deploy our products, weve created Data Center Helm chartscustomizable templates that can be configured to meet the unique needs of your business. You can even choose how to run them: either on your own hardware or on a cloud providers infrastructure. This allows you to stay in control of your data and meet your compliance needs while still using a more modern infrastructure. Helm charts have their own lifecycle, so updates contain certain features and are upgraded automatically.

Helm charts provide the essential building blocks needed to deploy Atlassian Data Center products (Jira, Confluence, Bitbucket, Bamboo, and Crowd) in Kubernetes clusters and give you the capability to integrate with your operation and automation tools. Learn more about Helm charts

Use Docker images for improved agility

To speed up development, you can take advantage of Data Centers hardened Docker container images. Using our Docker container images as part of your Data Center deployment allows you to cut significant time by streamlining and automating workflows.

After defining your required configuration once, you can instantly deploy exact replicas of your environment from the command line at every stage of your deployment lifecycle, giving you the agility needed to keep valuable work moving forward, and the flexibility to accommodate your organizations evolving development strategy over time.

Learn Kubernetes deployment architecture

The Kubernetes cluster can be a managed environment, such asAmazon EKS,Azure Kubernetes Service,Google Kubernetes Engine, or a custom on-premise system. We strongly recommend you set up user management, central logging storage, a backup strategy, and monitoring just as you would for a Data Center installation running on your own hardware.

Here's an architectural overview of what you'll get when deployingyour Data Center application on a Kubernetes cluster using the Helm charts:

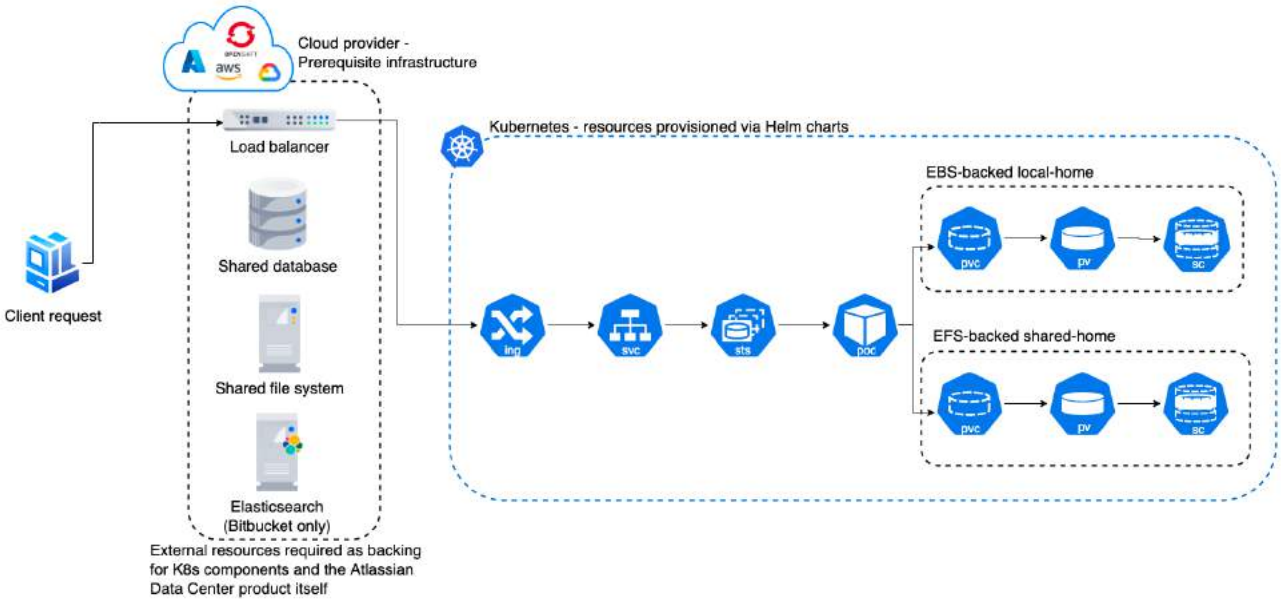

The following Kubernetes entities are required for product deployment:

- **Ingress and Ingress controller** (ing)the Ingress defines the rules for traffic routing, which indicate where a request will go in the Kubernetes cluster. The Ingress controller is the component responsible for fulfilling those rules.
- **Service** (svc)provides a single address for a set of pods to enable load-balancing between application nodes.
- **Pod**a group of one or morecontainers, with shared storage and network resources, and a specification for how to run the containers. Podsare the smallest deployable units of computing that you can create and manage in Kubernetes.
- **StatefulSets** (sts)manages the deployment and scaling of a set ofpods requiring persistent state.
- **PersistentVolume** (pv)a "physical"volumeon the host machine that stores your persistent data.
- **PersistentVolumeClaim** (pvc)reserves the Persistent Volume (PV) to be used by a pod or potentially multiple pods.
- **StorageClass** (sc)provides a way for administrators to describe the "classes" of storage they offer.

Installyour Data Center application on a Kubernetes cluster

To install and operate your Data Center application on a Kubernetes cluster using our Helm charts:

1. Follow the requirements and set up your environment according to the Prerequisites guide.
2. Perform the installation steps described in the Installation guide.
3. Learn how to upgrade applications, scale your cluster, and update resources using the Operation guide.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Installing Java for Confluence

This page contains instructions for installing the Java Development Kit (JDK). This is a manual step that's only required if you're installing Confluence from a zip or archive file.

If you're using the Confluenceinstaller, you don't need to install Java manually, but you can choose to use a different Java vendor.

Check theSupported Platforms page to find out which Java versions and vendors can be used with Confluence.

Installing Java

The JDK (Java Development Kit) needs to be installed on the same server that will have Confluence installed. We support running Confluence with the JDK or JRE (Java Runtime Environment). These instructions will just cover installing the JDK.

Before you start, go to **Control Panel** > **Programs and Features** to check whether a JDK is already installed.

To install the JDK on Windows:

1. Download the appropriate Eclipse Temurin OpenJDK or Oracle JDK version.
   Check theSupported Platforms page to find out which JDK / JRE versions and vendors are supported for your version of Confluence. Be sure to download the right one for your operating system.
2. Run the Java installer. Make a note of the installation directory, as you'll need this later.
3. Once the Java installation is complete, check that the `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.

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/JDK11`or `/user/lib/jvm/java11`, then your JDK is installed and properly configured.
- If nothing is displayed, you'll need to install the JDK or set the `$JAVA_HOME` environment variable. You can set this environment variable in your user account's 'profile' file. Alternatively, you can set this after installing Confluence, by defining this path in your Confluence installation's `setenv.sh` file, usually located in the Confluence `bin` directory.

To install the JDK on Linux:

1. Download the appropriate Eclipse Temurin OpenJDK or Oracle JDK version.
   Check theSupported Platforms page to find out which JDK / JRE versions are supported for your version of Confluence. Be sure to download the right one for your operating system.
2. Run the Java installer.
3. Open a shell console and 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

To install Confluence manually on Windows, you will need to set an environment variable to point Confluence to the your Java installation directory.

> This information is only relevant if you're installing Confluence manually on a Windows server. If you're using the installer, you don't need to do this.

Related pages

- Starting Tomcat as a Windows Service
- Installing Confluence in Linux

In most cases you should set the JRE_HOME environment variable, but if it is not set, Confluence will use JAVA_HOME.

Set the JAVA_HOME Variable

To set the JRE_HOME or JAVA_HOME variable:

1. Locate your Java installation directory

   > If you didn't change the path during installation, it'll be something like
   >
   > `C:\Program Files\Eclipse Adoptium\jdk-11.0.17.8-hotspot\` or
   >
   > `C:\Program Files\Java\jdk-17.0.4.1\`
   >
   > You can also type `where java` at the command prompt.

2. Do one of the following:
   Windows 7 Right click **My Computer** and select **Properties** > **Advanced**
   Windows 8 Go to **Control Panel** > **System** > **Advanced System Settings**
   Windows 10 Search for **Environment Variables** then select **Edit the system environment variables**
3. Click the **Environment Variables** button.
4. Under **System Variables**, click **New**.
5. In the **Variable Name** field, enter either:
   - `JAVA_HOME` if you installed the JDK (Java Development Kit)
     or
   - `JRE_HOME` if you installed the JRE (Java Runtime Environment)
6. In the **Variable Value** field, enter your JDK or JRE installation path.

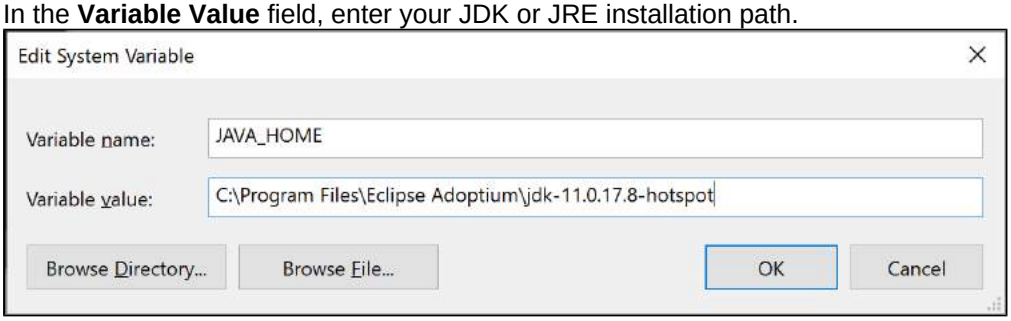

7. Click **OK** and **Apply Changes** as prompted

You'll need to close and re-open any command windows that were open before you made these changes, as there's no way to reload environment variables from an active command prompt. If the changes don't take effect after reopening the command window, restart Windows.

Set the JAVA_HOME variable via the command line

If you would prefer to set the JAVA_HOME (or JRE_HOME) variable via the command line:

1. Open Command Prompt (make sure you Run as administrator so you're able to add a system environment variable).
2. Set the value of the environment variable to your JDK (or JRE) installation path as follows:

```
setx /m JAVA_HOME "C:\Program Files\Java\jdk11.0.17.8"
```

Restart Command Prompt to reload the environment variables then use the following command to check the it's been added correctly.

3.
```
echo %JAVA_HOME%
```

You should see the path to your JDK (or JRE) installation.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Change the Java vendor or version Confluence uses

When you install Confluence Server using the installer, it will run Confluence with the Java Runtime Environment (JRE) that was bundled with that Confluence release.

If you want to use a different Java vendor, version, or you want to install the full JDK, you can tell Confluence to use the version of Java installed on your server.

Not all vendors and versions are supported, and some versions have known issues, so always check the Supported Platforms page, as using an unsupported version can cause problems in Confluence.

On this page:

- Check your current setup
- Installer method - Windows
- Installer method - Linux
- Environment variable method - Windows and Linux
- How Confluence determines which Java to use
- Which Java vendor can I use with my Confluence version?
- Known issues
- Upgrading Java

Check your current setup

How you change Confluence's Java path depends on whether you originally installed Confluence using the installer, or manually from a .zip or .tar.gz file.

The easiest way to check how Confluence is currently finding your Java is to:

1. Go to`<install-directory>/bin/setjre.sh`file (Linux) or `setjre.bat`(Windows) file.
2. Scroll to the bottom of the file and look for a line similar to the following. The file path may be different in your file.
   In Linux:

   ```
   JRE_HOME="/opt/atlassian/confluence/jre/"; export JRE_HOME
   ```

   In Windows:

   ```
   SET "JRE_HOME=C:\Program Files\Atlassian\Confluence\jre"
   ```

If a line similar to the one above is present, then JRE_HOME **is set** in this file by the installer, and you should use the **installer method**for Windows or Linux below.

If this line isn't present, JRE_HOME **is not set** in this file (because Confluence was installed manually), and you should use the **environment variable** method below.

Installer method - Windows

The way you do this depends on whether you run Confluence manually using the`start-confluence.bat`file, or as a Windows service.

In these examples we're going to point Confluence to the Eclipse Temurin OpenJDK JRE, which is installed on our Windows server at `C:\Program Files\Eclipse Adoptium\jdk-11.0.16.101-hotspot`. The location of your JRE will be different, but the steps are the same for any supported Java vendor and version.

If you start Confluence manually

To change the Java that Confluence uses if you start Confluence manually in Windows:

1. In Command Prompt, use the following command tocheck that Java is installed and has been added to your path correctly.

   ```
   > java -version
   ```

   This will return your Java version. If nothing is returned, or it returns the wrong version, check the installation instructions for your Java vendor.
2. Stop Confluence.
3. In the Confluence installation directory edit the`<install-directory>/bin/setjre.bat`file andchange the last lineto point to your local Java installation, as in the example below.

   ```
   SET "JRE_HOME=C:\Program Files\Eclipse Adoptium\jdk-11.0.16.101-hotspot\bin"
   ```

   If this line isn't present, exit this file and use the environment variable method below.
4. Start Confluence.
5. Go to > **General Configuration**>**System Information**and check that Confluence is using the expected Java version.

Remember, when you next upgrade Confluence this file will be overwritten, so you will need to re-apply this change to the new`setjre.bat`file.

If you run Confluence as a Windows service

To change the Java that Confluence uses if you run Confluence as a Windows service:

1. Open the Tomcat properties dialog. SeeHow to set system properties for Confluence running as a service on Windowsfor a step-by-step guide to locating your service and launching the Tomcat dialog.
2. Choose the**Java**tab.
3. Update the**Java Virtual Machine**line to point to the**AdoptOpenJDK jvm.dll**, as in the example below. The path to your Java installation will be different to our example.

   ```
   C:\Program Files\Eclipse Adoptium\jdk-11.0.16.101-hotspot\bin\server\jvm.dll
   ```

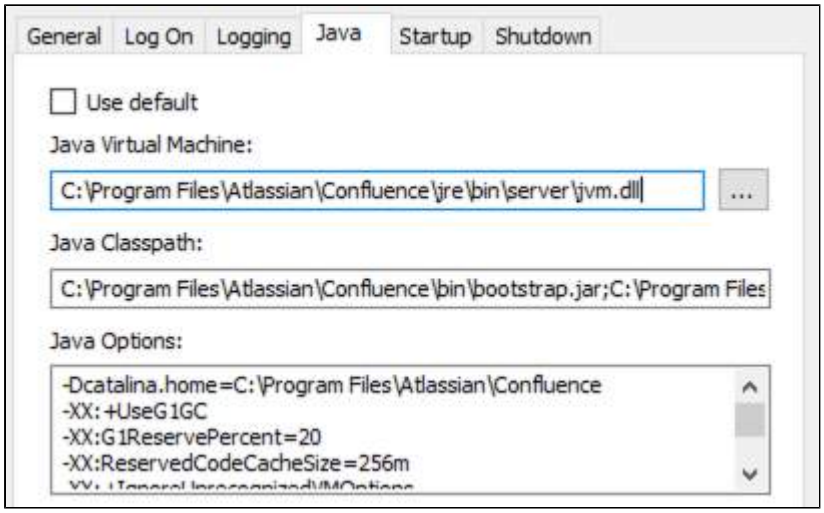

4. Restart the Confluence Windows Service.
5. Go to > **General Configuration**>**System Information**and check that Confluence is using the expected Java version.

Remember, when you next upgrade Confluence this file will be overwritten, so you will need to re-apply this change to the service.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Installer method - Linux

In this example we're going to point Confluence to the AdoptOpenJDK JRE, which is installed on our Linux server at /opt/java/adoptopenjdk/jdk-11.0.4.11-hotspot/. The location of your JRE will be different, but the steps are the same for any supported Java vendor and version.

To change the Java that Confluence uses in Linux:

1. In Terminal, use the following command to check that Java is installed and added to your path correctly.

   ```
   $ java -version
   ```

   This will return your Java version. If nothing is returned, or it returns the wrong version, seeInstalling Java for Confluenceor check the installation instructions for your Java vendor.
2. Stop Confluence.
3. In the Confluence installation directory edit the`<install-directory>/bin/setjre.sh`file andchange the last lineto point to your local Java installation, as in the example below.

   The path to your Java installation will be different to our example.

   ```
   JRE_HOME="/opt/java/adoptopenjdk/jdk-11.0.4.11-hotspot/"; export JRE_HOME
   ```

   If this line isn't present, exit this file and use the environment variable methodbelow.
4. Start Confluence.
5. Go to > **General Configuration**>**System Information**and check that Confluence is using the expected Java version.

Remember, when you next upgrade Confluence this file will be overwritten, so you will need to re-apply this change to the new `setjre.sh`file.

Environment variable method - Windows and Linux

If you installed Confluence manually (the path to the bundled JRE was not automatically set in the `setjre`file), Confluence will use the path set in the JRE_HOME environment variable. If JRE_HOME is not set, it will use the path set in JAVA_HOME.

See Setting JAVA_HOME variable for Confluence to find out how to set this environment variable in Windows.

Refer to the documentation for your Linux distribution to find out how to set an environment variable in Linux.

You won't need to update the JRE_HOME environment variable when you upgrade Confluence, but you will need to update the path if you upgrade Java.

How Confluence determines which Java to use

The JRE_HOME set in the`setjre`file takes precedence. If you installed Confluence using the installer, this will be automatically set to the Java version bundled with Confluence.

If JRE_HOME is not set in the`setjre.bat`or`setjre.sh`file, Confluence will use the JRE_HOME defined in your environment or service. If it can't find JRE_HOME, it will use the JAVA_HOME environment variable.

Which Java vendor can I use with my Confluence version?

The following table lists the supported Java vendors, and whether Oracle or AdoptOpenJDK is bundled with Confluence.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Confluence version	Supported Java vendors	Bundled Java vendor
6.6.12 and earlier	Oracle JRE	Oracle JRE
6.7.0 to 6.13.1, and 6.14.0	Oracle JRE	Oracle JRE
6.13.2 to 6.13.x, and 6.14.1 to 7.13.1, and 7.14.0	Oracle JDK/JRE AdoptOpenJDK	AdoptOpenJDK
7.13.2 to 7.13.x, and 7.14.1 to latest	Oracle JDK/JRE Adopt OpenJDK Eclipse Temurin	Eclipse Temurin

Known issues

- You may find that Oracle is still listed as the vendor in System Information. This is a known issue in Confluence which we hope to have resolved soon. The Java version will be reported correctly, so you can use that to make sure Confluence is pointing to the right version.
- AdoptOpenJDK does not include a required font configuration package, which may cause issues when installing in Linux. See Confluence Server 6.13 or later fails with FontConfiguration error when installing on Linux operating systemsfor information on how to install the required package manually.
- AdoptOpenJDK is now known as Temurin.

Upgrading Java

If you choose not to use the bundled Java version, you will need to manually update Java from time to time, to get access to new security fixes and other improvements.

Always check the Supported Platforms page before upgrading, for any known issues affecting particular Java versions.

If upgrading to a major version, for example from Java 11 or Java 17, be aware that some Java arguments will not be recognised in later versions. When you upgrade, make sure you apply your customisations manually, don't simply copy over your old `setenv.sh`/ `setenv.bat`file, or existing Java options if you run Confluence as a service.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Creating a Dedicated User Account on the Operating System to Run Confluence

A dedicated user should be created to run Confluence, because Confluence runs as the user it is invoked under and therefore can potentially be abused.

This is optional if you're evaluating Confluence, but is required for production installations.If you used the Confluence installer on Linux, the installer created this user automatically.

Create a dedicated user account

Linux

If your operating system is *nix-based (for example, Linux or Solaris), type the following in a console:

```
$ sudo /usr/sbin/useradd --create-home --comment "Account for running Confluence" --shell /bin/bash
confluence
```

Windows

If your operating system is Windows create the dedicated user account bytyping the following at the Windows command line:

```
> net user confluence mypassword /add /comment:"Account for running Confluence"
```

(This creates a user account with user name 'confluence' and password 'mypassword'. You should choose your own password.)

Alternatively, openthe Windows 'Computer Management' console to add your 'confluence' user with its own password.

Next,Use the Windows 'Computer Management' console to remove the 'confluence' user's membership of all unnecessary Windows groups, such as the default 'Users' group.

If Windows is operating under Microsoft Active Directory, ask your Active Directory administrator to create your 'confluence' account (with no prior privileges).

Allow the account to write to specific Confluence directories

Ensure that the following directories can be read and written to by this dedicated user account (e.g. 'confluence'):

- The sub-directories of theConfluence Installation Directory:
    - `logs`
    - `temp`
    - `work`
- The entireConfluence Home directory.

Set who can access Confluence directories in Linux

To achieve this in Linux run the following commands:

```
sudo chown -R confluence <confluence-home-folder>/
sudo chown -R confluence <confluence-install-folder>/logs
sudo chown -R confluence <confluence-install-folder>/work
sudo chown -R confluence <confluence-install-folder>/temp
```

The other install directories should be left as root as those are controlled by the installer and allow for future upgrades:

```
sudo chmod -R u=rwx,g=rx,o=rx <confluence-install-folder>
sudo chmod -R u=rwx,g=rx,o=rx <confluence-home-folder>
```

See alsoBest Practices for Configuring Confluence Security.

Confirm who can access Confluence directories in Windows

After installing Confluence youshould check the permissions assigned to the installation directory, and make sure there are no unnecessary permissions being inherited. You can also repeat this process for the home directory.

To check the permissions for the install directory:

1. Right click your installation directory and select **Properties**.
2. In the **Security** tab, select **Advanced**.
3. Select **Disable inheritance**, and when prompted choose **Convert inherited permissions into explicit permissions on this object**.
4. Select **OK**.
5. Select any group or user account that should not have access and choose **Remove**.
   We recommend limiting access to only the dedicated 'confluence' user and system administrator groups.
6. Select **OK** to apply changes to your install directory (and all sub-directories).

To confirm your changes, log in to Windows with a normal user account, and check that you can't access the contents of the install directory.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

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
- 3. Enter your license key
- 4. Production installation: database configuration
- 5. Production installation: load content
- 6. Set up user management
- 7. Connect to your Jira application
- 8. Set up system administrator account
- 9. Setup is Complete

1. Start the setup wizard

1. Start Confluence (if it is not already running)
   For Windows, go to**Start** > **Programs** > **Confluence** > **Start Confluence Server**.
   Or, run the start-up script found in the `bin` folder of your installation directory:
   - `start-confluence.bat` for Windows.
   - `start-confluence.sh` for Linux-based systems.
2. Go 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

In this step, you'll choose whether you want a trial or a production installation.

- **Trial installation**
  Choose this option if you don't have a license, and want to try Confluence for the first time. You'll need an external database.

- **Production installation**
  Set up Confluence with your own external database. This option is recommended for setting up Confluence in a production environment.

3.Enter your license key

Follow the prompts to generate an evaluation license, or enter an existing license key.To retrieve an existing license key head tomy.atlassian.com, or to purchase a new commercial license go towww.atlassian.com/buy.

If you selected a **Trial installation** in the previous step,Confluence will generate your license. This will take a few minutes. Once complete,go tostep 8 below.

If you selected a **Production installation**, go to the next step to set up your external database.

4. Production installation: database configuration

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 configuring character encoding.
- **Database name:** When creating a new external database, give it the name '**confluence**'.

Next it's time to set up your database. Some things to consider:

- Check thesupported platformslist to confirm that your chosen database and version is supported.
- Seedatabase configurationfor 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.

Screenshot: Database configuration

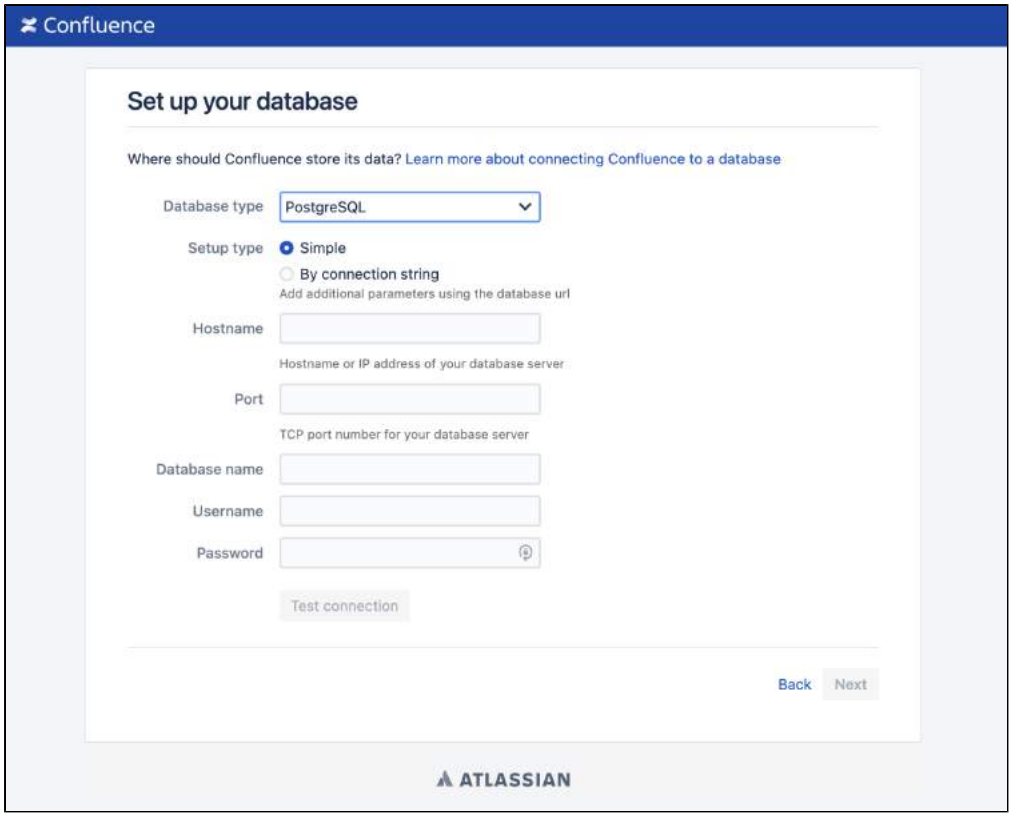

Database connection

Confluence will connect to your database with a direct JDBC connection. Connection pooling is handled within Confluence.

- **Driver Class Name** The Java class name for the appropriate database driver. This will depend on the JDBC driver, and will be found in the documentation for your database. Note that Confluence bundles some database drivers, but you'll need to install the driver yourself if it is not bundled. SeeDatabase JDBC Driversfor details.
- **Database URL** The JDBC URL for the database you will be connecting to. This will depend on the JDBC driver, and will be found in the documentation for your database.
- **User Name** and **Password** A valid username and password that Confluence can use to access your database.

You will also need to know:

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

- 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.

5. 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, you will be able to Restore from backup by import your existing Confluence data. This can only be done after the setup wizard is complete by following the instructions on Restore a Site.

Screenshot: Load content screen

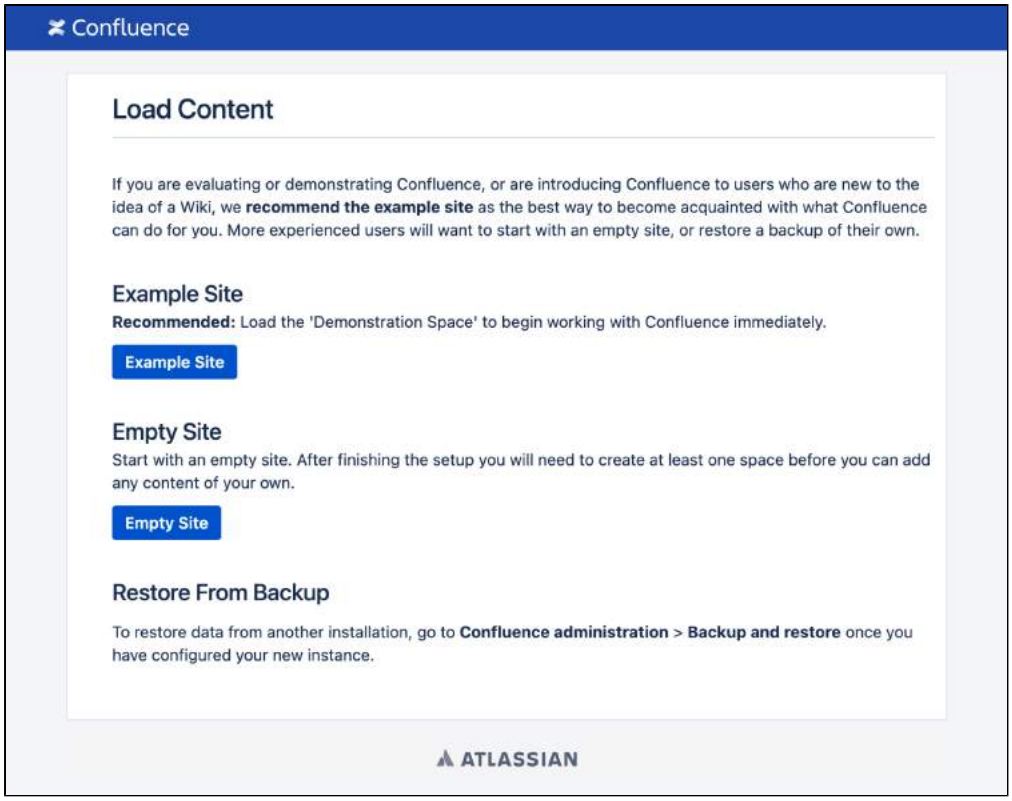

6. Set up user management

You can choose to manage Confluence's users and groups inside Confluence or in a Jira application, such as Jira Software or Jira Service Management.

- If you do not have a Jira application installed, or if you would prefer to set up external user management later, 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**.

7. Connect to your Jira application

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

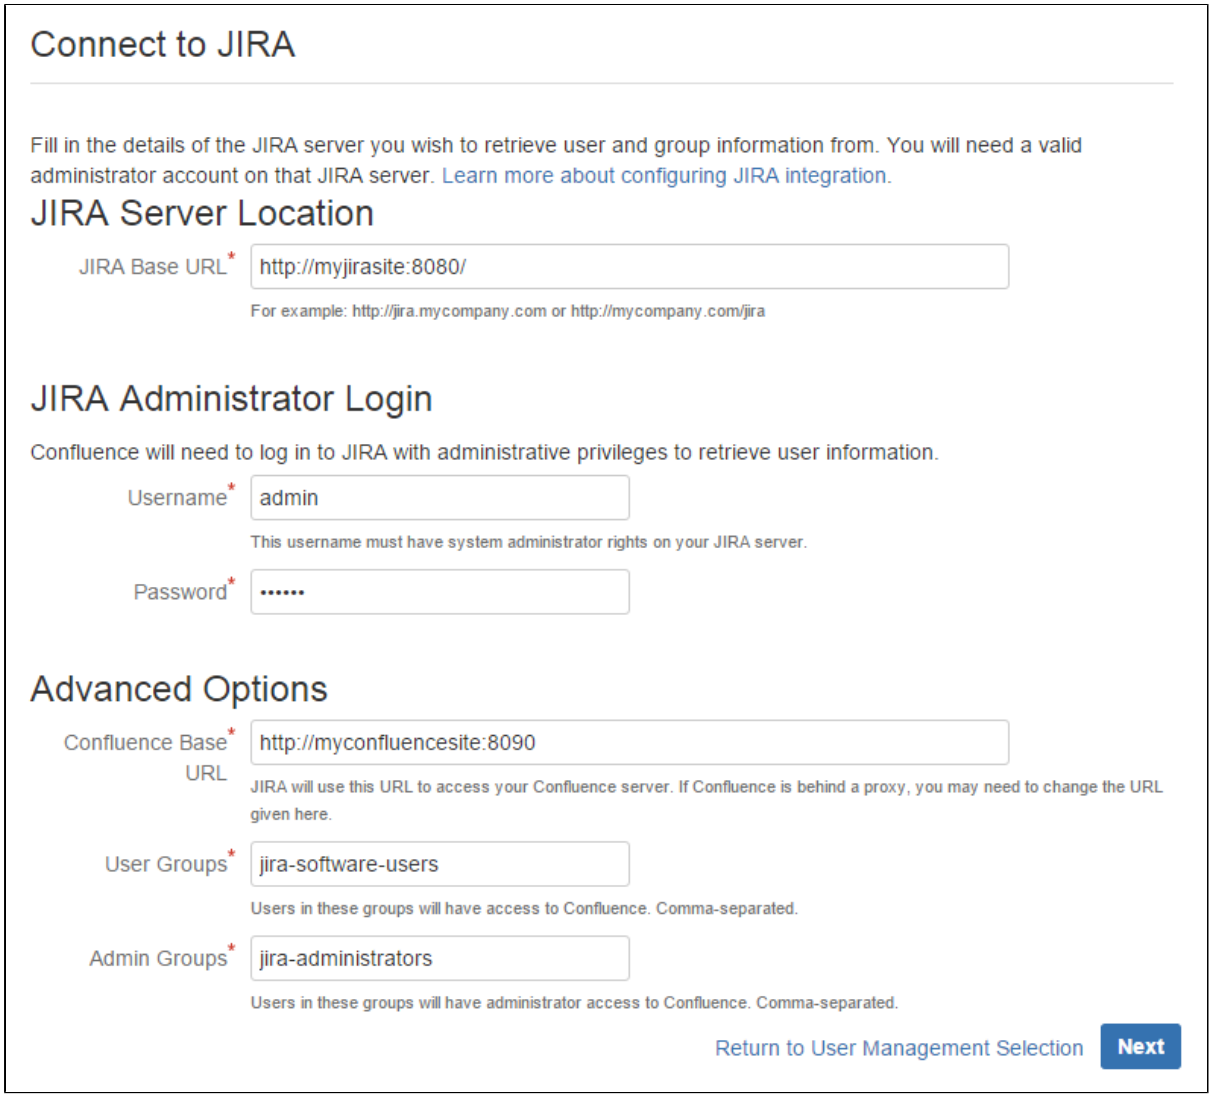

Enter the following information:

- **Jira Base URL** - the address of your Jira server, such as `http://www.example.com:8080/jira/` or `http://jira.example.com`
- **Jira Administrator Login**- this is the username and password of a user account that has the Jira System Administrator global permission in your Jira application.

  Confluence will also use this username and password to create a local administrator account which will let you access Confluence if Jira is unavailable. Note that this single account is stored in Confluence's internal user directory, so if you change the password in Jira, it will not automatically update in Confluence.

- **Confluence Base URL** - this is the URL Jira will use to access your Confluence server. The URL you give here overrides the base URL specified in Confluence, for the purposes of connecting to the Jira application.
- **User Groups** - these are the Jira groups whose members should be allowed to use Confluence. Members of these groups will get the 'Can use' permission for Confluence, and will be counted in your Confluence license. The default user group name differs depending on your Jira version:
  - Jira 6.4 and earlier: `jira-users`.
  - Jira Software 7.x and later: `jira-software-users`
  - Jira Core 7.x and later: `jira-core-users`
  - Jira Service Management (formerly Jira Service Desk) 3.x and later: `jira-servicedesk-users`
- **Admin Groups** Specify one or more Jira groups whose members should have administrative access to Confluence. The default group is `jira-administrators`. These groups will get the system administrator and Confluence administrator global permissions in Confluence.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

For full details and a troubleshooting guide, see Configuring Jira Integration in the Setup Wizard.

8. 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.

9. 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.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Configuring Jira Integration in the Setup Wizard

This page describes the **Connect to Jira** step in the Confluence setup wizard.

If you are already using a Jira application, you can choose to delegate user management to Jira, instead of separately maintaining your users in Confluence.

You'll be able to specify exactly which groups in your Jira app should also be allowed to log in to Confluence. Your license tiers do not need to be the same for each application.

It's possible to connect Confluence to Jira after completing the setup process, but it's much quicker and easier to set it up at this stage.

You can delegate Confluence's user management to:

- Jira 4.3 or later
- Jira Core 7.0 or later
- Jira Software 7.0 or later
- Jira Service Management (formerly Jira Service Desk) 3.0 or later.

On this page:

- Connecting to a Jira application in the Setup Wizard
- Troubleshooting

Related pages:

- User Management Limitations and Recommendations
- Connecting to Crowd or Jira for User Management
- Confluence Setup Guide

Connecting to a Jira application in the Setup Wizard

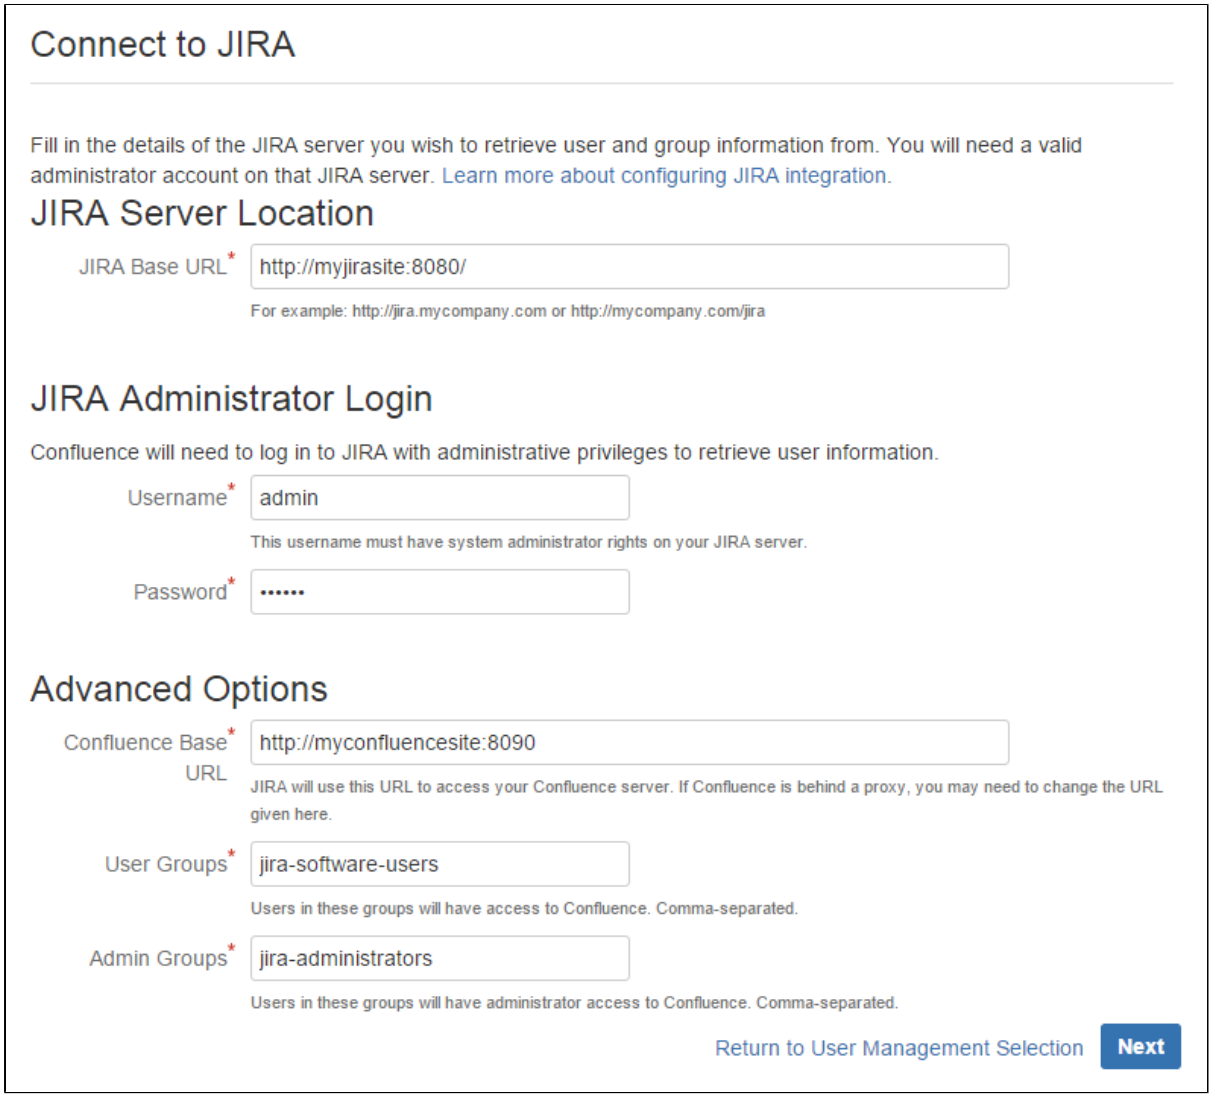

Enter the following information:

- **Jira Base URL** - the address of your Jira server, such as `http://www.example.com:8080/jira/` or `http://jira.example.com`
- **Jira Administrator Login**- this is the username and password of a user account that has the Jira System Administrator global permission in your Jira application.

  Confluence will also use this username and password to create a local administrator account which will let you access Confluence if Jira is unavailable. Note that this single account is stored in Confluence's internal user directory, so if you change the password in Jira, it will not automatically update in Confluence.

- **Confluence Base URL** - this is the URL Jira will use to access your Confluence server. The URL you give here overrides the base URL specified in Confluence, for the purposes of connecting to the Jira application.
- **User Groups** - these are the Jira groups whose members should be allowed to use Confluence. Members of these groups will get the 'Can use' permission for Confluence, and will be counted in your Confluence license. The default user group name differs depending on your Jira version:
  - Jira 6.4 and earlier: `jira-users`.
  - Jira Software 7.x and later: `jira-software-users`
  - Jira Core 7.x and later: `jira-core-users`
  - Jira Service Management (formerly Jira Service Desk) 3.x and later: `jira-servicedesk-users`
- **Admin Groups** Specify one or more Jira groups whose members should have administrative access to Confluence. The default group is `jira-administrators`. These groups will get the system administrator and Confluence administrator global permissions in Confluence.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

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 configurationand resolve any conflict with existing links then try the Connect to Jira step again.

Removing a partial configuration

If you hit a roadblock, you'll need to log in to Jira and remove the partial integration before you can try again. The specific steps will differ depending on your Jira application and version, but the essentials are the same for all versions:

- Log in to Jira as a user with system administrator permissions.
- In the Administrator screens, go to **Application Links**.
- Remove the application link that matches the base URL of your Confluence server.
- In the User Management screens, go to **Jira User Server**.
- Remove the link that matches the name and base URL of your Confluence server from the list of applications that can use Jira for user management.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

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. SeeConnecting to Crowd or Jira for User Management.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Upgrading Confluence

In this guide we'll run you through using the installer to upgrade your Confluence site to the latest Confluence version on Windows or Linux.

Upgrading to any later version is free if you have current software maintenance. See ourLicensing FAQto find out more.

Other ways to upgrade Confluence:

- Manuallyupgrade Server or single-node Data Center without using the installer.
- Cluster with downtime upgrade your Data Center cluster.
- Cluster without downtime - rolling upgrade to a compatible bug fix version, with no downtime.

XML backups should**not**be used to upgrade Confluence.

On this page:

- Before you begin
- Plan your upgrade
  - 1. Determine your upgrade path
  - 2. Complete the pre-upgrade checks
  - 3. Upgrade Confluence in a test environment
- Upgrade Confluence
  - 4. Back up
  - 5. Download Confluence
  - 6. Run the installer
- After the upgrade
  - 7. Copy your database driver
  - 8. Reinstall the service if required (Windows only)
  - 9. Re-apply any modifications
  - 10. Update your apps (add-ons)
  - 11. Update your reverse proxy and check you can access Confluence
- Troubleshooting

Before you begin

Before you upgrade Confluence, there's a few questions you need to answer.

Which upgrade method is the best option?	You can choose to upgrade using the installer, or manually using a zip or tar.gz file. In most cases the installer is the easiest way to upgrade your Confluence instance. You will need to upgrade manually if you are: • moving to another operating system or file location as part of this upgrade. • upgrading from **Confluence 3.5 or earlier** • upgrading from **Confluence 5.6** or earlier and previously used the **EAR/WAR distribution**to deploy Confluence into an existing application server. • performing a rolling upgrade, and you need to upgrade each node individually.

Are you eligible to upgrade?	To check if software maintenance is current for your license, go to > **General Configuration** and select **License Details** on the left panel to make sure the license support period has not expired. 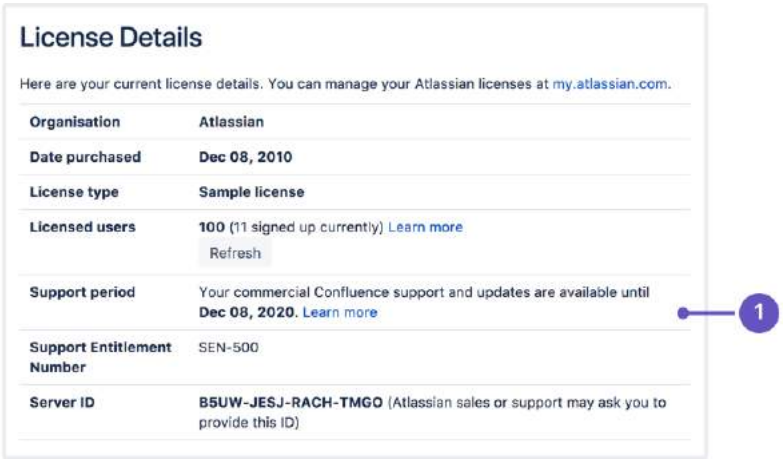 1. **Software maintenance**: upgrade at any time during this period. If your support period has expired, follow the prompts to renew your license and reapply it before upgrading.
Have our supported platforms changed?	Check the Supported Platforms page for the version of Confluence you are upgrading to. This will give you info on supported operating systems, databases and browsers. **Good to know:** • The Confluence installer includes Java (JRE) and Tomcat, so you won't need to upgrade these separately. • If you need to upgrade your database, be sure to read the upgrade notes for the Confluence version you plan to upgrade to (and any in-between) to check for any database configuration changes that you may need to make.
Do you need to make changes to your environment?	Newer Confluence versions sometimes require changes to your environment, such as providing more memory or adjusting your reverse proxy settings. **Good to know:** We useUpgrade Notesto communicate changes that will impact you, such as: • Changes to supported databases, memory requirements or other changes that will impact your environment. • Features that have significantly changed or been removed in this release. • Actions you may need to take in your instance or environment immediately after the upgrade. It's important to read the notes for the version you're upgrading to and those in-between. For example, if you are upgrading from 5.8 to 5.10 you should read the upgrade notes for 5.9 and 5.10.

Plan your upgrade

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Create a custom upgrade plan

Planning an upgrade? You can instantly generate a tailored upgrade plan from within Confluence. Head to > **General Configuration** > **Plan your upgrade**

You'll need to have a compatible version of the Troubleshooting and Support tools app installed. Learn more

1. Determine your upgrade path

Use the table below to determine the most efficient upgrade path from your current version to the latest versions of Confluence.

Your Version	Recommended upgrade path to Confluence 8
2.7 or earlier	Upgrade to **2.7.4** then upgrade to **3.5.17**, and follow paths below.
2.8 to 3.4	Upgrade to **3.5.17**, and follow paths below.
3.5	Upgrade to **5.0.3**, and follow paths below.
4.0 to 4.3	Upgrade to **5.10.x**, and follow paths below.
5.0 to 5.10	Upgrade to **7.19.x**, and follow paths below.
6.0.5 to 8.x	Upgrade directly to the latest version of Confluence 8.

Confluence 8 is a major upgrade

Be sure to check the Confluence Upgrade Matrix, take a full backup, and test your upgrade in a non-production environment before upgrading your production site.

Enterprise releases

A Long Term Support release is a feature release that gets backported critical security updates and critical bug fixes during its entire two-year support window. If you can only upgrade once a year, consider upgrading to a Long Term Support release.Learn more

Long Term Support releases were originally referred to as Enterprise Releases.

2. Complete the pre-upgrade checks

1. Check theUpgrade Notesfor the version you plan to upgrade to (and any in between).
2. Go to> **General Configuration**> **Plan your upgrade**then select the version you want to upgrade to. This will run somepre-upgrade checks.
3. Go to > **General Configuration** > **Troubleshooting and support tools**to run the health check.

   If the software maintenance period included in your license has expired you can keep using Confluence, but you'll need to renew before you can upgrade.

   Go to > **General Configuration** > **License Details** and follow the prompts to renew your license.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Database character encoding must be set to UTF8 (or UTF8MB4 for MySQL databases, or AL32UTF8 for Oracle databases). You will not be able to upgrade to the current Confluence versions unless you have the correct character encoding.

4. Go to > **Manage apps**and scroll down to the**Confluence Update Check**to check the compatibility of your Marketplace apps.
5. Choose the version you plan to upgrade to then hit**Check**.

If your users rely on particular Marketplace apps, you may want to wait until they are compatible before upgrading Confluence. Vendors generally update their apps very soon after a major release.

Good to know:

- You can disable an app temporarily while you upgrade if it is not yet compatible.
- Compatibility information for Atlassian Labs and other free apps is often not available immediatley after a new release. In many cases the app will still work, so give it a try in a test site before upgrading your production site.

3. Upgrade Confluence in a test environment

1. Create a staging copy of your current production environment.
   See Create a staging environment for upgrading Confluence for help creating an environment to test your upgrade in.
2. Follow the steps below to upgrade your test environment.
3. Test any unsupported user-installed apps, customizations (such as custom theme or layouts) and proxy configuration (if possible) before upgrading your production environment.

Upgrade Confluence

4. Back up

1. Back up your **database** and confirm the backup was created properly.
   If your database does not support online backups you'll need to stop Confluence first.
   If you have a MySQL database,make sure your back up includes stored procedures/functions.

   Once you've confirmed your database backup was successful, you can choose to disable the automatic generation of an upgrade recovery file, as this process can take a long time for sites that are medium sized or larger.
2. Back up your **installation directory**
   The installer will completely replace this directory, so any files you've added (such as a keystore or SSL certificate) won't be retained. The installation wizard will back up this directory before starting the upgrade, but you should also back it up manually first.
3. Back up your**home directory**.
   The installation wizard gives you the option to also back up your home directory as part of the installation process, but you should also back up this directory manually before starting the upgrade.

   You can find the location of your home directory in the `<installation-directory>/confluence/WEB-INF/classes/confluence-init.properties` file.

   This is where your search indexes and attachments are stored. If you store attachments outside the Confluence Home directory, you should also backup your attachments directory.

5.Download Confluence

Download the installer for your operating system.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

- Latest versionhttps://www.atlassian.com/software/confluence/download
- Older versionshttps://www.atlassian.com/software/confluence/download-archives

6. Run the installer

1. Run the installer.

   Run the .exe file. We recommend using a Windows administrator account.

   If prompted to allow the upgrade wizard to make changes to your computer, choose '**Yes**'. If you do not, the installation wizard will have restricted access to your operating system and any subsequent installation options will be limited.

   Change to the directory where you downloaded Confluence then execute this command to make the installer executable:

   ```
   $ chmod a+x atlassian-confluence-X.X.X-x64.bin
   ```

   Where `X.X.X` is is the Confluence version you downloaded.

   Next, run the installer we recommend using`sudo`to run the installer:

   ```
   $ sudo ./atlassian-confluence-X.X.X-x64.bin
   ```

   You can also choose to run the installer with root user privileges.

2. Follow the prompts to upgrade Confluence:

   a. When prompted choose**Upgrade an existingConfluence installation**(for Linux users this is option 3).

   b. Make sure the**Existing Confluence installation directory**suggested by the wizard is correct (especially important if you have multiple Confluence installations on the same machine).

   c. **Back up Confluence home**is strongly recommended. This will create a .zip backup of the Confluence home and installation directories.

   d. The installation wizard notifies you of customizations in the Confluence Installation directory. Make a note of these as you'll need to reapply them later.

      The installation wizard's ability to notify you about customizations will depend on how your existing Confluence instance was installed:

      - If your current Confluence instance was installed using the installer, the wizard will check the entire Confluence Installation directory.
      - If your current Confluence instance was installed manually it will only check the`confluence`subdirectory of the Confluence Installation directory. The installation wizard will**not**notify you of modifications in any other directory, for example modifications to start-up scripts under the`bin`directory or modifications to the`server.xml`file (such as anSSL configuration).

      You won't be notified about files you've added to the installation directory, so be sure to back them up first.

3. The wizard will shut down your Confluence instance and proceed with the upgrade. Once complete, it will restart Confluence and you can then launch Confluence in your browser to confirm the upgrade was successful.

   Depending on the size of your instance and the number of upgrade tasks to be run, this step may take a few minutes or several hours.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

After the upgrade

7. Copy your database driver

If you're using an Oracle or MySQL database, you'll need to copy the jdbc driver jar file from your existing Confluence installation directory to `confluence/WEB-INF/lib` in your new installation directory.

Microsoft SQL and Postgres users can skip this step.

8. Reinstall the service if required (Windows only)

If you run Confluence as a service on Windows you should delete the existing service then re-install the service by running `<install-directory>/bin/service.bat`.

This makes sure the service gets the most recent JVM options.

9. Re-apply any modifications

During the upgrade the wizard migrated the following from your existing Confluence installation:

- TCP port values in your `<install-directory>/conf/server.xml` file.
- Location of your Confluence home directory in `<install-directory>/confluence/WEB-INF/classes/confluence-init.properties`.

All other customizations, including `CATALINA_OPTS` parameters in your `<install-directory>/bin/setenv.sh` / `setenv.bat` files, need to be reapplied manually.

Any other configurations, customizations (including any other modifications in the `<install-directory>/conf/server.xml` file), the path to your own Java installation in `<install-directory>/bin/setjre.sh`, or `setjre.bat`, or additional files added to the installation directory are **not migrated** during the upgrade and need to be reapplied manually.

1. Stop your upgraded Confluence instance.
2. Edit each file, and reapply the customizations in your upgraded Confluence Installation directory.
3. Copy over any additional files (such as keystore or SSL certificate)
4. Restart the upgraded Confluence instance.

We **strongly recommend** you test your customizations in a test instance prior to upgrading your production instance as changes may have been made to Confluence that make your customizations unusable.

> Edit the new file manually, rather than copying over the old file, as the default configuration in these files may have changed between Confluence versions.
>
> System properties can change from time to time. Be sure to check you're using the correct Recognized System Properties.

10. Update your apps (add-ons)

You can update any apps that are compatible with the new version of Confluence.

1. Go to > **Manage apps**
2. Update your apps to the supported versions.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

At this stage, it can be useful to clear your plugin cache. Learn how to do this

This is optional, but can be useful to avoid any issues with third-party apps and plugins.

11. Update your reverse proxy and check you can access Confluence

If you are upgrading from**Confluence 5.x to Confluence 6.x**you will need to modify your reverse proxy (if used) to add Synchrony, which is required for collaborative editing. SeeProxy and SSL considerationsfor more information on the changes you'll need to make to your proxy config.

Once your upgrade is complete, you should access Confluence (via your reverse proxy, not directly) and:

- Head to > **General Configuration**>**Collaborative editing**and check the Synchrony status is**running**.
- Edit any page to check that your browser can connect to Synchrony.

SeeTroubleshooting Collaborative Editingfor suggested next steps if Synchrony is not running or you see an error in the editor, as you may have a misconfigured reverse proxy.

Troubleshooting

Did something go wrong?

If you need to retry the upgrade,**you must restore your pre-upgrade backups first.** Do not attempt to run an upgrade again, or start the older version of Confluence again after an upgrade has failed.

- **Can't proceed with upgrade because license has expired**
  If your license has expired and was not renewed and reapplied before upgrading you will receive errors during the upgrade process. Seeupgrading beyond current license periodfor information on how to resolve this problem.
- **Can't proceed with upgrade because of a conflict with anti virus**
  Some anti-virus or other Internet security tools may interfere with the Confluence upgrade process and prevent the process from completing successfully, particularly if you run Confluence as a Windows service. If you experience or anticipate experiencing such an issue with your anti-virus / Internet security tool, disable this tool first before proceeding with the Confluence upgrade.
- **Database does not support online backups**
  The upgrade wizard will prompt you to backup your database using your database's backup utilities. If your database does not support online backups, stop the upgrade process, shut down Confluence, perform your database backup and then run the installer again to continue with the upgrade.
- **Upgrade is taking a very long time**
  If you have a very large database (i.e. database backups take a very long time to complete), setting the`confluence.upgrade.recovery.file.enabled`system propertyto falsewill speed up the upgrade process. It should be used only when there is a process to back up database and verify the backup before performing an upgrade.
- **Confluence doesn't start**
  Incompatible Marketplace apps can occasionally prevent Confluence from starting successfully. You can troubleshoot the problem by starting Confluence with all user installed apps temporarily disabled. See Start and Stop Confluence for more info.
- **Collaborative editing errors**
  If Synchrony is not running or you see an error, head toTroubleshooting Collaborative Editingfor info on how to get collaborative editing up and running in your environment. The most common problems are a misconfigured reverse proxy or port 8091 not being available for Synchrony.
- **Space directory is empty after the upgrade**
  If you are upgrading from Confluence 6.3 or earlier, there's a known issue where spaces do not appear in the space directory. You'll need to reindex your site after upgrading to fix this.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

You can also refer to the Upgrade Troubleshooting guide in the Confluence Knowledge Base, or check for answers from the community at Atlassian Answers.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Upgrading Beyond Current Licensed Period

This page covers what to do ifyou have upgraded Confluence to a version beyond your current license entitlement.

Related pages:

- Upgrading Confluence
- Working with Confluence Logs

License warnings

During the upgrade you will see an error similar to the following in yourapplication logs.

```
ERROR [confluence.upgrade.impl.DefaultUpgradeManager] runUpgradePrerequisites
Current license is not valid: SUPPORT_EXPIRED
```

When you try to access Confluence in your browser, you'll see this warning:

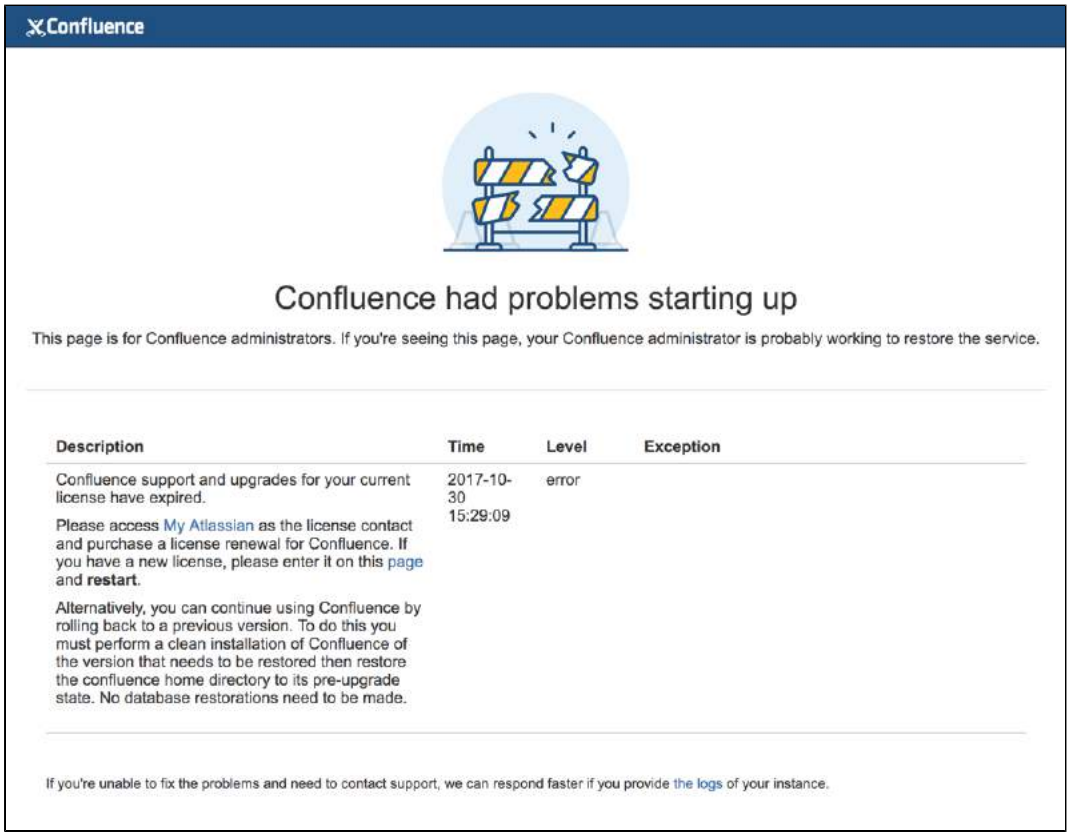

Updating the Confluence license

1. Head to my.atlassian.com to renew your license or purchase a new license.
2. Follow the prompts on the warning screen to enter your new license key.

3. Restart Confluence to pick up the license change. You should now be able to log in to Confluence as normal.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Confluence Post-Upgrade Checks

This article provides a list of items for Confluence Administrators to check after a Confluence upgrade to ensure that it has completed successfully. This list is not exhaustive, but it does cover common upgrade mistakes.

Before You Begin

After you have completed an upgrade, you should see the following message in the `atlassian-confluence.log` file:

```
2010-03-08 08:03:58,899 INFO [main] [atlassian.confluence.upgrade.
AbstractUpgradeManager] entireupgradeFinished Upgrade completed successfully
```

If you do not see the line in your log similar to the one above, this means that your upgrade may not have completed successfully. Please check our Upgrade Troubleshooting documentation to check for a suitable recommendation or fix.

Upgrade Checklist

Here's a recommended list of things to check after completing an upgrade

1.The editor

Edit a page to check your browser can connect to Synchrony, which is required for collaborative editing. SeeTroubleshooting Collaborative Editing if you are not able to edit a page.

2. Layout and Menu

Visit the Confluence dashboard and check that it is accessible and displays as expected. Test the different Internet browsers that you have in use in your environment. In addition, confirm that the layout appears as expected and that the menus are clickable and functioning.

3. Search

Try searching for content, for example pages, attachments or user names. Check that the expected results are returned.If you notice any problems, you may want to take advantage of the maintenance window to rebuild the indexes from scratch. SeeContent Index Administration.

4. Permissions

Confirm that you can visit a page that has view restrictions, but you have permission to view. Confirm that you can edit a page that has edit restrictions but you have permission to edit. Make sure that the permissions of child pages are functioning as well. Involve as many space administrators as possible to confirm they are working. Confirm that anonymous or forbidden users cannot access or modify restricted pages.

5. Attachments

Confirm that attachments are accessible and searchable.

6. Marketplace apps

Outdated third-party apps can cause upgrade failure. Quite often, they will just be incompatible and simply do not work anymore. If you discover that your app is no longer working, please check for the latest version for your app in the The Atlassian Marketplaceor check for compatibility in theUniversal Plugin Manager.

Migration from Wiki Markup to XHTML-Based Storage Format

If you are upgrading **to Confluence 4.0 or later** from an older version (From Confluence 3.5.x or earler) then as part of the upgrade an automatic migration of your content will take place. This is a non-destructive process. Your existing content is not overwritten. Instead, the migration process will create a new version of each wiki markup page. The new version will use the new XHTML-based storage format, so that you can edit the page in the Confluence rich text editor.

In addition, if you are **upgrading to Confluence 4.3 or later** from an older version then as part of the upgrade an automatic migration of your page templates will take place. See Migration of Templates from Wiki Markup to XHTML-Based Storage Format.

Note: Even though the process is non-destructive, you must be sure to perform a backup of your database and home directory prior to starting the new version of Confluence, as we recommend for any Confluence upgrade.

Migration process

Depending on the size of your Confluence installation, the migration from wiki markup to the new XHTML-based storage format could prove time consuming. The duration of the migration is difficult to estimate; this is due to a number of site specific factors. As a rough guide, a test dataset we migrated was 130,000 pages, totalling approximately 700Mb, which took six minutes.

On this page:

- Migration process
- Watching the migration logs during the upgrade
- Re-running the migration for content that completely failed the migration
- Re-attempting the migration for content in 'unmigrated-wiki-markup' macro
- Notes

Related pages:

- Migration of Templates from Wiki Markup to XHTML-Based Storage Format
- Upgrading Confluence

The following properties that can be modified to allow finer control over the migration process:

Property	Purpose	Default
confluence.wiki.migration.threads	The number of concurrent worker threads migrating content	4
confluence.wiki.migration.batch.size	The number of items migrated in each batch of work	500
confluence.wiki.migration.versioncomment	The comment associated with the newly migrated version of each piece of content	"Migrated to Confluence 4.0"

(For instructions on setting Confluence system properties see this document.)

Again, due to the large variability in Confluence installations it is hard to give specific recommendations for the above settings. One point to note though that both increasing batch size and the number of threads (or both) will increase the peak memory required for migration. If memory is an issue then as you increase one of these settings consider decreasing the other.

Another factor to be aware of if modifying these defaults is that of the cache settings employed in your site. The migration will quickly populate certain Confluence caches so be sure that if you have customized caches as described here that there is enough memory on the server for these caches should they reach maximum capacity.

Watching the migration logs during the upgrade

To monitor the progress of a site migration you should watch the output in the application log.

Typical logging progress will be shown by multiple log entries at the INFO level of the following format:

```
WikiToXhtmlMigrationThread-n - Migrated 2500 of 158432 pages, this batch migrated 500/500 without error
```

There may be a wide array of messages logged from each individual page but any errors are also collected for display in a single migration report once all content has been processed. Here is a typical example of such a report:

```
Wiki to XHTML Exception Report:
Summary:
        0 settings values failed.
        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.UnknownMacroMigrationException: The macro
link is unknown.. Message: The macro link is unknown.
        2) Type: comment, Id: 6919, Title: null, Global Scope. Cause: com.atlassian.confluence.content.
render.xhtml.migration.exceptions.UnknownMacroMigrationException: The macro mymacro is unknown.. Message:
The macro mymacro is unknown.
```

Each entry in the report will identify the content that caused migration exceptions as well as displaying the exceptions themselves.

In almost all cases any content reported as errored will have been migrated to the new XHTML-based storage format, but will actually consist of wiki markup content wrapped within an XML 'unmigrated-wiki-markup' macro. This content will still be viewable in Confluence and editable within the new Confluence Editor.

However, in some cases a batch of content may actually have completely failed to migrated. This is most typically due to an unhandled exception causing a database transaction rollback. This would be reported in the log with a message like this:

```
Unable to start up Confluence. Fatal error during startup sequence: confluence.lifecycle.core:
pluginframeworkdependentupgrades (Run all the upgrades that require the plugin framework to be available) -
com.atlassian.confluence.content.render.xhtml.migration.exceptions.MigrationException: java.util.concurrent.
ExecutionException: org.springframework.transaction.UnexpectedRollbackException: Transaction rolled back
because it has been marked as rollback-only
```

Confluence provides no further report about this scenario and will also allow Confluence to restart as normal without retrying a migration. If a user tries to view any such unmigrated content they will see an exception similar to this:

```
java.lang.UnsupportedOperationException: The body of this ContentEntityObject ('Page Title') was 'WIKI' but
was expected to be 'XHTML'
```

The solution is to ensure you manually re-run the site migration after the restart.

Re-running the migration for content that completely failed the migration

A Confluence Administrator can restart the site migration if there was any content that failed migration (see previous section). Only the content that is still formatted in wiki markup will be migrated, so typically a re-migration will take less time than the original migration.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

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**.

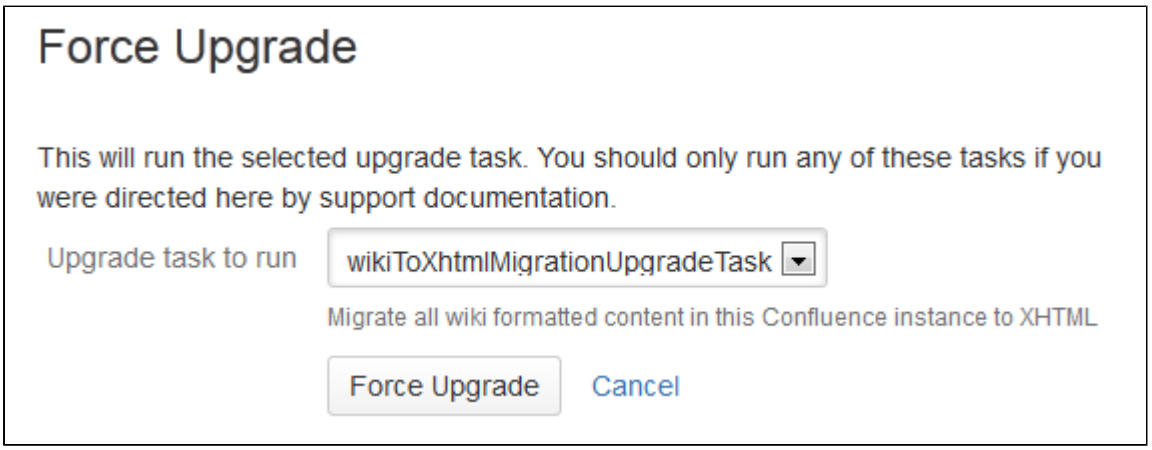

Re-attempting the migration for content in 'unmigrated-wiki-markup' macro

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`

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

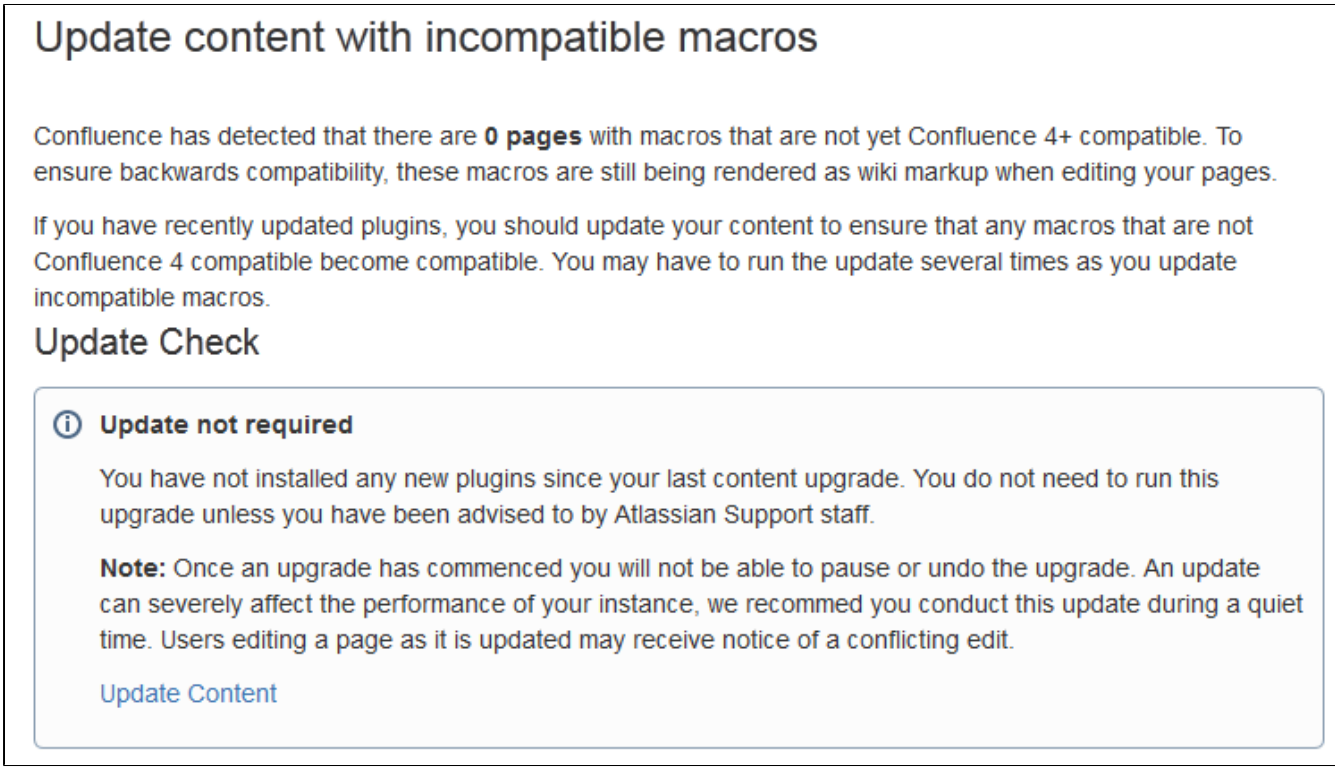

Notes

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.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Migration of Templates from Wiki Markup to XHTML-Based Storage Format

If you are **upgrading to Confluence 4.3 or later from an older version** (from Confluence 4.2.x or earlier) then as part of the upgrade an automatic migration of your page templates will take place. This is a non-destructive process. Your existing content is not overwritten. Instead, the migration process will create a new version of each space template and each global template on your Confluence site. The new version will use the new XHTML-based storage format, so that you can edit the template in the Confluence rich text editor.

Note: Nevertheless, you must be sure to perform a backup of your database and home directory prior to starting the new version of Confluence, as we recommend for any Confluence upgrade.

Watching the migration logs during the upgrade

To monitor the progress of a site migration you should watch the output in the application log.

A typical logging progress will be shown by multiple log entries at the INFO level of the following format:

```
WikiToXhtmlMigrationThread-n - Migrated 22 of 29 PageTemplates.
```

On this page:

- Watching the migration logs during the upgrade
- Re-running the migration
- Notes

Related pages:

- Migration from Wiki Markup to XHTML-Based Storage Format
- Page Templates
- Upgrading Confluence

There may be a wide array of messages logged from each individual template, but any errors are also collected for display in a single migration report once all content has been processed. Here is a typical example of such a report:

```
Wiki to XHTML Exception Report:
Summary:
        0 settings values failed.
        2 PageTemplates failed.
        0 ContentEntityObjects failed.
Content Exceptions:
        1) Type: page, Id: 332, Title: Release Notes 1.0b3, Space: DOC - Confluence 4.0 Beta. Cause: com.
atlassian.confluence.content.render.xhtml.migration.exceptions.UnknownMacroMigrationException: The macro
link is unknown.. Message: The macro link is unknown.
        2) Type: comment, Id: 6919, Title: null, Global Scope. Cause: com.atlassian.confluence.content.
render.xhtml.migration.exceptions.UnknownMacroMigrationException: The macro mymacro is unknown.. Message:
The macro mymacro is unknown.
```

Each entry in the report will identify the content that caused migration exceptions as well as displaying the exceptions themselves.

In almost all cases any content reported as errored will have been migrated to the new XHTML-based storage format, but will actually consist of wiki markup content wrapped within an XML 'unmigrated-wiki-markup' macro. This content will still be viewable in Confluence and editable within the Confluence rich text editor.

However, in some cases a batch of content may actually have completely failed to migrate. This is most typically due to an unhandled exception causing a database transaction rollback. This would be reported in the log with a message like this:

```
Unable to start up Confluence. Fatal error during startup sequence: confluence.lifecycle.core:
pluginframeworkdependentupgrades (Run all the upgrades that require the plugin framework to be available) -
com.atlassian.confluence.content.render.xhtml.migration.exceptions.MigrationException: java.util.concurrent.
ExecutionException: org.springframework.transaction.UnexpectedRollbackException: Transaction rolled back
because it has been marked as rollback-only
```

Confluence provides no further report about this scenario and will also allow Confluence to restart as normal without retrying a migration. If a user tries to view or edit an unmigrated template, the wiki template editor will be used.

The solution is to manually re-run the site migration after the restart, as described below.

Re-running the migration

A Confluence administrator can restart the template migration if any templates have failed the migration (see previous section). Only the templates that are still formatted in wiki markup will be migrated again. Typically, a re-migration will take less time than the original migration.

To manually re-run the migration:

1. Open this URL in your browser:`<Confluence Address>/admin/force-upgrade.action`
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

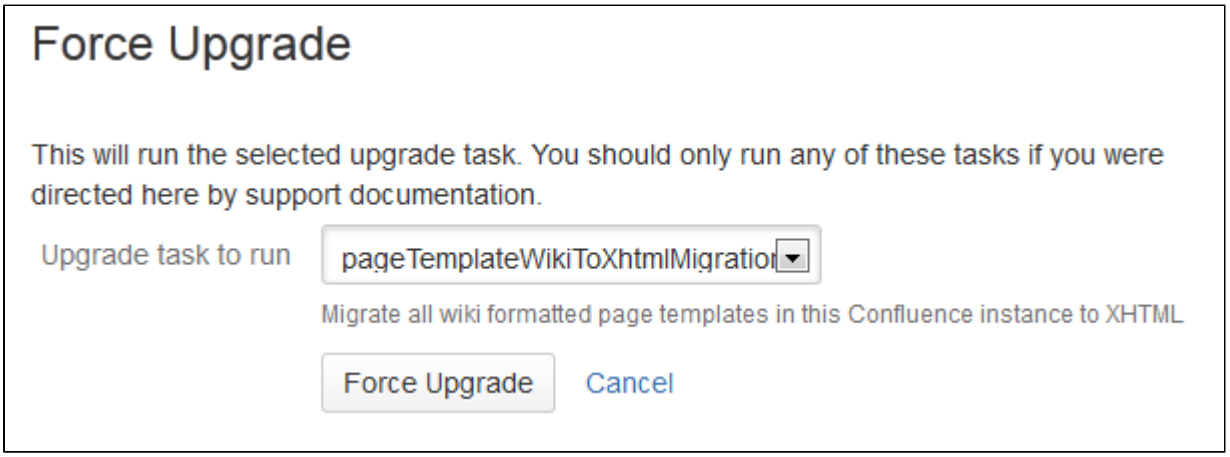

Notes

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.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Upgrading Confluence Manually

In this guide we'll run you through upgrading your Confluence site to the latest Confluence version on Windows or Linux using the zip / tar.gz file.

Upgrading to any later version is free if you have current software maintenance. See ourLicensing FAQto find out more.

Other ways to upgrade Confluence:

- Installer the simplest way to upgrade Confluence.
- Data Center upgrade your Data Center cluster.
- Rolling upgrade- upgrade your Data Center cluster to the latest available bug fix version, with no downtime.

XML backups should**not**be used to upgrade Confluence.

On this page:

- Before you begin
- Plan your upgrade
  - 1. Determine your upgrade path
  - 2. Complete the pre-upgrade checks
  - 3. Upgrade Confluence in a test environment
- Upgrade Confluence
  - 4. Back up
  - 5. Download Confluence
  - 6. Extract the file and upgrade Confluence
- After the upgrade
  - 7. Reinstall the service (Windows only)
  - 8. Re-apply any modifications
  - 9. Update your reverse proxy and check you can access Confluence
- Troubleshooting

Before you begin

Before you upgrade Confluence, there's a few questions you need to answer.

Is manual the right upgrade method for you?	You can choose to upgrade using the installer, or manually using a zip or tar.gz file. In most cases the installer is the easiest way to upgrade your Confluence instance. You will need to upgrade manually if you are: • moving to another operating system or file location as part of this upgrade. • upgrading from **Confluence 3.5 or earlier** • upgrading from **Confluence 5.6** or earlier and previously used the **EAR/WAR distribution**to deploy Confluence into an existing application server. • performing a rolling upgrade, and you need to upgrade each node individually.

Are you eligible to upgrade?	To check if software maintenance is current for your license, go to > **General Configuration** and select **License Details** on the left panel to make sure the license support period has not expired. 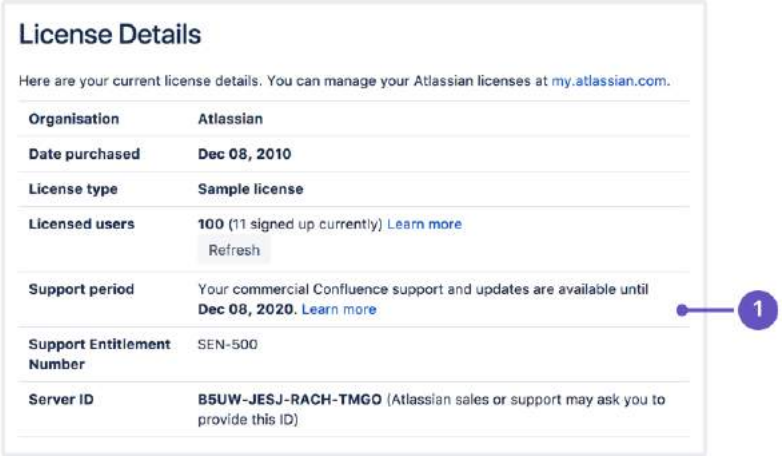 1. **Software maintenance**: upgrade at any time during this period. If your support period has expired, follow the prompts to renew your license and reapply it before upgrading.
Have our supported platforms changed?	Check the Supported Platforms page for the version of Confluence you are upgrading to. This will give you info on supported operating systems, databases and browsers. **Good to know:** • If you need to upgrade Java, remember to update your JAVA_HOME variable to the new version. • The Confluence installer includes Tomcat, so you won't need to upgrade it separately. • If you need to upgrade your database, be sure to read the upgrade notes for the Confluence version you plan to upgrade to (and any in-between) to check for any database configuration changes that you may need to make.
Do you need to make changes to your environment?	Newer Confluence versions sometimes require changes to your environment, such as providing more memory or adjusting your reverse proxy settings. **Good to know:** We useUpgrade Notesto communicate changes that will impact you, such as: • Changes to supported databases, memory requirements or other changes that will impact your environment. • Features that have significantly changed or been removed in this release. • Actions you may need to take in your instance or environment immediately after the upgrade. It's important to read the notes for the version you're upgrading to and those in-between. For example, if you are upgrading from 5.8 to 5.10 you should read the upgrade notes for 5.9 and 5.10.

Plan your upgrade

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

1. Determine your upgrade path

Use the table below to determine the most efficient upgrade path from your current version to the latest versions of Confluence.

Your Version	Recommended upgrade path to Confluence 8
2.7 or earlier	Upgrade to **2.7.4** then upgrade to **3.5.17**, and follow paths below.
2.8 to 3.4	Upgrade to **3.5.17**, and follow paths below.
3.5	Upgrade to **5.0.3**, and follow paths below.
4.0 to 4.3	Upgrade to **5.10.x**, and follow paths below.
5.0 to 5.10	Upgrade to **7.19.x**, and follow paths below.
6.0.5 to 8.x	Upgrade directly to the latest version of Confluence 8.

Confluence 8 is a major upgrade

Be sure to check the Confluence Upgrade Matrix, take a full backup, and test your upgrade in a non-production environment before upgrading your production site.

2. Complete the pre-upgrade checks

1. Check theUpgrade Notesfor the version you plan to upgrade to (and any in between).
2. Go to > **General Configuration**> **Plan your upgrade**then select the version you want to upgrade to. This will run somepre-upgrade checks.
3. Go to > **General Configuration** > **Troubleshooting and support tools**to run the health check.

   If the software maintenance period included in your license has expired you can keep using Confluence, but you'll need to renew before you can upgrade.

   Go to > **General Configuration** > **License Details** and follow the prompts to renew your license.

   Database character encoding must be set to UTF8 (or UTF8MB4 for MySQL databases, or AL32UTF8 for Oracle databases). You will not be able to upgrade to the current Confluence versions unless you have the correct character encoding.

4. Go to > **Manage apps**and scroll down to the**Confluence Update Check**to check the compatibility of your Marketplace apps.
5. Choose the version you plan to upgrade to then hit**Check**.

   If your users rely on particular Marketplace apps, you may want to wait until they are compatible before upgrading Confluence. Vendors generally update their apps very soon after a major release.

   Good to know:

   - You can disable an app temporarily while you upgrade if it is not yet compatible.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

- Compatibility information for Atlassian Labs and other free apps is often not available immediatley after a new release. In many cases the app will still work, so give it a try in a test site before upgrading your production site.

3. Upgrade Confluence in a test environment

1. Create a staging copy of your current production environment.
   See Create a staging environment for upgrading Confluence for help creating an environment to test your upgrade in.

2. Follow the steps below to upgrade your test environment.

3. Test any unsupported user-installed apps, customizations (such as custom theme or layouts) and proxy configuration (if possible) before upgrading your production environment.

Upgrade Confluence

4. Back up

1. Back up your database and confirm the backup was created properly.
   If your database does not support online backups you'll need to stop Confluence first.

   Once you've confirmed your database backup was successful, you can choose todisable the automatic generation of an upgrade recovery file, as this process can take a long time for sites that are medium sized or larger.

2. Back up your**installation directory**and**home directory**.

   You can find the location of your home directory in the `<installation-directory>/confluence/WEB-INF/classes/confluence-init.properties` file.

   This is where your search indexes and attachments are stored. If you store attachments outside the Confluence Home directory, you should also backup your attachments directory.

5.Download Confluence

Download the appropriate file for your operating system -https://www.atlassian.com/software/confluence/download

6. Extract the file and upgrade Confluence

1. Stop Confluence.
   SeeUsing read-only mode for site maintenance if you need to provide uninterrupted access.

2. Extract (unzip) the files to adirectory (this is your new installation directory, and must be different to your existing installation directory)
   Note: There are some known issues with unzipping the archive on Windows. We recommend using 7Zip or Winzip.

3. Edit `<Installation-Directory>\confluence\WEB-INF\classes\confluence-init.properties`file to point to your**existing**Confluence home directory.

4. If you're using an Oracle or MySQL database, you'll need to copyyour jdbc driver jar file from your existing Confluence installation directory to`confluence/WEB-INF/lib`in your **new installation directory**.

5. There are some additional steps you make need to take if:
   - you are running Confluence as a**Windows Service**

     If you are running Confluence as a Windows service, go to the command prompt and type:

     ```
     <Installation-Directory>\bin\service.bat remove Confluence
     ```


Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

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 theConfluence installer, you'll need to run`<confluence auto installer installation folder>\UninstallService.bat`.

- You are running Confluence on a different port (not the default 8090)

  If you are not running Confluence on port 8090 update `<Installation-Directory>\conf\server.xml`file to include your ports.

6. Start your new Confluence. You should not see the setup wizard.

After the upgrade

7. Reinstall the service (Windows only)

If you run Confluence as a service on Windows you should delete the existing service then re-install the service by running`<install-directory>/bin/service.bat`.

This makes sure the service gets the most recent JVM options.

8. Re-apply any modifications

If you have customized Confluence (such as an SSL configuration in the`server.xml`file, or `CATALINA_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 your upgraded Confluence instance.
2. Reapply the customizations to the relevant files in the newly upgraded Confluence Installation directory.
3. Restart the upgraded Confluence instance.

We**strongly recommend**you test your customizations in a test instance prior to upgrading your production instance as changes may have been made to Confluence that make your customizations unsuable.

9. Update your reverse proxy and check you can access Confluence

If you are upgrading from**Confluence 5.x to Confluence 6.x**you will need to modify your reverse proxy (if used) to add Synchrony, which is required for collaborative editing. SeeProxy and SSL considerationsfor more information on the changes you'll need to make to your proxy config.

Once your upgrade is complete, you should access Confluence (via your reverse proxy, not directly) and:

- Head to > **General Configuration**>**Collaborative editing**and check the Synchrony status is**running**.
- Edit any page to check that your browser can connect to Synchrony.

SeeTroubleshooting Collaborative Editingfor suggested next steps if Synchrony is not running or you see an error in the editor, as you may have a misconfigured reverse proxy.

Troubleshooting

Did something go wrong?

If you need to retry the upgrade,**you must restore your pre-upgrade backups first.** Do not attempt to run an upgrade again, or start the older version of Confluence again after an upgrade has failed.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

- **Can't proceed with upgrade because license has expired**
  If your license has expired and was not renewed and reapplied before upgrading you will receive errors during the upgrade process. Seeupgrading beyond current license periodfor information on how to resolve this problem.

- **Collaborative editing errors**
  If Synchrony is not running or you see an error, head toTroubleshooting Collaborative Editingfor info on how to get collaborative editing up and running in your environment. The most common problems are a misconfigured reverse proxy or port 8091 not being available for Synchrony.
- **Upgrade is taking a very long time**
  If you have a very large database (i.e. database backups take a very long time to complete), setting the`confluence.upgrade.recovery.file.enabled`system propertyto falsewill speed up the upgrade process. It should be used only when there is a process to back up database and verify the backup before performing an upgrade.

You can also refer to theUpgrade Troubleshootingguide in the Confluence Knowledge Base, or check for answers from the community atAtlassian Answers.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Create a staging environment for upgrading Confluence

When you upgrade Confluence we strongly recommend performing the upgrade in a test environment before upgrading your production site. In this guide we'll refer to this test environment as *staging*.

Most Confluence licenses include a free developer license for use in a staging environment. SeeHow to get a Confluence Developer licenseto find out how to access your license.

On this page:

- Create a staging environment
- Additional configuration options
- Upgrade your staging environment

Create a staging environment

1. Replicate your environment

Your staging environment should closely replicate your real-live environment (production), including any reverse proxies, SSL configuration, or load balancer (for Data Center).You may decide to use a different physical server or a virtualized solution. The main thing is to make sure it is an appropriate replica of your production environment.

For the purposes of these instructions, we assume your staging environment is physically separate fromyour production environment, and has the same operating system (and Java version if you've installed Confluence manually).

2. Replicate your database

To replicate your database:

1. Back up your production database. Refer to the documentation for your database for more info on the best way to do this.
2. Install your database on the staging server and restore the backup.
3. Re-create any database triggers that may still reference the original database name. If you're using SQLServer, the following query may help identifying them:

```
SELECT table_name = so.name
,trigger_name = st.name
,trigger_text = sc.text
,create_date = st.create_date
FROM sys.triggers st
JOIN sysobjects so ON st.parent_id = so.id
JOIN syscomments sc ON sc.id = st.[object_id];
```

The steps for restoring your database backup will differ depending on your chosen database and backup tool. Make sure:

- Your new staging database has a**different**name from your production database.
- Your staging database user account has the**same**username and password as your production database user account.
- Character encoding and other configurations are the same as your production database (for example character encoding should be Unicode UTF-8 (or AL32UTF8 for Oracle databases).

3. Replicate Confluence

To replicate Confluence, make a copy of your Confluence installation and point it to your staging database. These instructions are for Confluence Server (for Data Center there are some additional steps before you start Confluence).

1. Copy your entire **production installation directory** to your staging server.
2. Copy your entire **production home directory** to your staging server.
3. Edit `<installation-directory>/confluence/WEB-INF/classes/confluence-init.properties`to point to your staging home directory.

4. Edit<home-directory>/confluence.cfg.xmlor <installation-directory>/server.xmlto point to your staging database.

```
<property name="hibernate.connection.url">jdbc:postgresql://localhost:5432/confluencestaging<
/property>
```

5. Start Confluence with the followingSystem Properties to make sure your staging site does not send notifications to real users.

```
-Datlassian.notifications.disabled=true
-Datlassian.mail.senddisabled=true
```

6. Head tohttp://localhost:<port>and log in to Confluence on your staging server.
7. Go to > **General Configuration**and change the**base URL**of your staging site (for example mysite.staging.com)
8. Go to > **General Configuration**> **License Details** and apply your development license.
9. Go to > **General Configuration**> **System Information** and check that Confluence is correctly pointing to your staging database, and staging home directory.

> ⚠ It's essential to check that you are not still connected to your production database.

Additional steps for Data Center

If you have Confluence Data Center, the process is much the same as for Confluence Server described above. You will copy each local home and installation directory to each staging node, and then:

1. Copy the**production shared home directory** to the staging server.
2. Edit<local-home-directory>/confluence.cfg.xmlto point to your staging shared home directory. This change **must** be made on every staging node.

Changes to the <installation-directory>/confluence/WEB-INF/classes/confluence-init.propertiesand <home-directory>/confluence.cfg.xml must be made on **every** staging node.

When it comes time to start Confluence, start one node at a time, as usual.

4. Replicate external user management (optional)

If you're managing users in Jira, Crowd, or in an external LDAP directory you can:

- replicate Jira, Crowd, or your external directory in your staging environment and point your Confluence staging site to your staging external directory (recommended).
- provide your staging server with network or local access to the same hosts as your production server.

Additional configuration options

There are a number of additional things you may want to change in your staging environment, to make sure it does not interact with your production environment, or to clearly differentiate it for users.

Bypass single sign-on

If you've configured single sign-on, you might want to bypass this in your staging environment.

SeeEnable default login page to bypass SAML in Confluence Data Center.

Disable CDN

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

If you've configured a CDN to cache static assets, you may experience problems with broken resources as the CDN configuration is for your production environment, not the staging environment.

To find out how to disable CDN, seeConfigure your CDN for Confluence Data Center.

Modify application links (recommended)

If you have application links between Confluence and other Atlassian applications you should change the server ID on each staging application. SeeHow to change the server ID of Confluence andChange the server ID for an instance of Jira server for Jira.

If you don't change the server ID and update your application links there is a chance that when you create a new application link in production it will point to your staging server instead.

To review the Application Links manually in the database,use the following following SQL query:

```
select * from bandana where bandanakey like 'applinks%';
```

Modify external gadgets

If you have external gadgets configured, you can update these from the database, using the following SQL query:

```
select * from bandana where bandanakey = 'confluence.ExternalGadgetSpecStore.specs'
```

Change the global color scheme

If can be helpful to use a different color scheme on your staging site, to differentiate it from your production site. See Customizing Color Schemesfor how to do this.

You can also find this data in the database using the following SQL query:

```
select * from bandana where bandanakey = 'atlassian.confluence.colour.scheme';
```

Change the instance name (recommended)

It is a good idea to change the name of your staging site, to differentiate it from your production site. Head to > **General Configuration**and update the **Site Title** if Confluence is running.

If Confluence is not running, you can do this from the database.You can find the site title using the following SQL query:

```
select * from bandana where bandanakey = 'atlassian.confluence.settings';
```

The attribute you are looking for is`setTitle.`

Add a banner

It can be useful to add a banner to your staging site, to provide useful information like the date of the last refresh, or who to contact if you want to make changes.

If you have a Confluence Data Center license, you can do this by enabling the banner that is used by read-only mode (you don't need to enable read-only mode to use the banner).

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

If you have a Confluence Server license, you can manually add a banner using HTML. Head to > **General Configuration** > **Custom HTML**.Remember to close your tags properly, or Confluence may not display correctly.

If you want to add a banner before starting Confluence, you can do it in the database.You can find the custom HTML using the following SQL query:

```
select * from bandana where bandanakey = 'atlassian.confluence.settings';
```

The attribute you are looking for is`customHtmlSettingsafterBodyStart`

Disable specific plugins

You might want to disable specific plugins or check whether these plugins are already disabled or not. See theHow to reset all Confluence plugins back to their default state through the databaseknowledge base article to find how to do this.

You can alsodisable plugins in Confluence in 6.1+ using Java system properties.

Upgrade your staging environment

Once you have created your staging environment, you can upgrade it in the same way you would your production environment.

Make a note of how long the upgrade takes, as this information will help you plan your production system outage and communicate with your users.

You can also use your staging environment to test any customizations or essential Marketplace apps in your site.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Upgrade Confluence without downtime

If you run Confluence Data Center in a cluster, you may be able to upgrade Confluence without any downtime for your users. This method is known as a rolling upgrade.

In a rolling upgrade, your site is put into upgrade mode, which temporarily allows nodes running different Confluence versions to join the cluster. As you take each node offline to upgrade it, the other active nodes keep your Confluence site available to users. Once all nodes have been upgraded in turn, you finalize the upgrade and turn off upgrade mode.

On this page:

- Can I upgrade without downtime?
- Before you begin
- Prepare for the rolling upgrade
  - 1. Complete pre-upgrade checks
  - 2. Prevent the installation or upgrade of apps during the upgrade period
  - 3. Back up Confluence Data Center
  - 4. Set up a staging environment to test the rolling upgrade
- Perform the rolling upgrade

Can I upgrade without downtime?

Whether you can upgrade your Confluence Data Center cluster without downtime depends on the version you are upgrading from, and the version you are upgrading to.Learn more about thedifferent types of releases.

Upgrading from	Upgrading to		
	Bugfix	**Feature**	**Platform**
Confluence 7.8 and earlier	✗ Requires downtime	✗ Requires downtime	✗ Requires downtime
Confluence 7.8 to 7.13	✓ No downtime (for example, from 7.12.0 to 7.12.2)	✗ Requires downtime	✗ Requires downtime
Confluence 7.14 and later	✓ No downtime (for example, from 7.14.0 to 7.14.2)	✓ No downtime when upgrading to the *next* feature version (for example, from 7.14.x to 7.15.x) ✗ Requires downtime if the upgrade spans *more than one* feature version (for example, from 7.14.x to 7.17.x)	✗ Requires downtime

Before you begin

Before you start planning a rolling upgrade, theres a few questions you need to answer.

Does my Confluence deployment support rolling upgrades?	You can only perform a rolling upgrade with no downtime on a multi-node Confluence cluster. Clustering is only supported on a Confluence Data Center license. In addition, a rolling upgrade involves enabling upgrade mode, which is only available in Confluence Data Center. Learn more about multi-node clustering in Confluence
Do I have enough nodes to support user requests during the rolling upgrade?	You need to take a node offline to upgrade it. During this time, other active nodes will take over the offline nodes workload. Make sure you have enough active nodes to handle user traffic at any given time. If possible, add a node temporarily to your cluster to compensate for offline nodes.
Is the version compatible with rolling upgrades?	Whether you can upgrade without downtime depends on the version you are upgrading from, and the version you are upgrading to. The pre-upgrade check will confirm whether you can upgrade without downtime.

Prepare for the rolling upgrade

1.Complete pre-upgrade checks

1. Check theUpgrade Notesfor the version you plan to upgrade to (and any in between).
2. Go to > **General Configuration**> **Plan your upgrade**then select the version you want to upgrade to. This will run somepre-upgrade checks.
3. Go to > **General Configuration** > **Troubleshooting and support tools**to run the health check.

   If the software maintenance period included in your license has expired you can keep using Confluence, but you'll need to renew before you can upgrade.

   Go to > **General Configuration** > **License Details** and follow the prompts to renew your license.

   Database character encoding must be set to UTF8 (or UTF8MB4 for MySQL databases, or AL32UTF8 for Oracle databases). You will not be able to upgrade to the current Confluence versions unless you have the correct character encoding.

4. Go to > **Manage apps**and scroll down to the**Confluence Update Check**to check the compatibility of your Marketplace apps.
5. Choose the version you plan to upgrade to then hit**Check**.

   If your users rely on particular Marketplace apps, you may want to wait until they are compatible before upgrading Confluence. Vendors generally update their apps very soon after a major release.

   Good to know:

   - You can disable an app temporarily while you upgrade if it is not yet compatible.
   - Compatibility information for Atlassian Labs and other free apps is often not available immediatley after a new release. In many cases the app will still work, so give it a try in a test site before upgrading your production site.

2. Prevent the installation or upgrade of apps during the upgrade period

If you manage Confluence with a team of admins, schedule the rolling upgrade with them. Notify them to postpone any app installs or upgrades until after the rolling upgrade. Installing or upgrading apps during a rolling upgrade could result in unexpected errors.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

3. Back up Confluence Data Center

Backup and Restore provides an overview of manual and scheduled backup methods in Confluence. For larger sites, we recommend having a robust production backup strategy.

If your deployment is hosted on AWS, we recommend that you use the AWS native backup facility, which utilizes snapshots to back up your site.For more information, seeAWS Backup.

⚠ During a rolling upgrade, backup and restore tasks are placed in a queue unless they were started before the rolling upgrade. Canceling a restore task during a rolling upgrade isn't supported and isn't recommended.

4. Set up a staging environment to test the rolling upgrade

We strongly recommend that you perform the rolling upgrade on a staging or test environment first.

1. Create a staging copy of your current production environment.
   See Create a staging environment for upgrading Confluence for help creating an environment to test your upgrade in.
2. Follow the steps below to upgrade your test environment.
3. Test any unsupported user-installed apps, customizations (such as custom theme or layouts) and proxy configuration (if possible) before upgrading your production environment.

Perform the rolling upgrade

There are three methods for performing a rolling upgrade, depending on what orchestration tools your deployment uses.

Method	Description	Instructions
Manual upgrade	A manual upgrade is suitable for deployments that feature minimal orchestration, particularly in node upgrades. If your deployment is based on ourAzure templates, you'll also need to perform a manual upgrade.	Upgrade a Confluence cluster manually without downtime
AWS CloudFormation	If your deployment is defined by an AWS CloudFormation template (like ourAWS Quick Start), then you can use the same template to orchestrate your upgrade.	Upgrade a Confluence cluster on AWS without downtime
API-driven	You can orchestrate the entire rolling upgrade process through API calls.	Upgrade a Confluence cluster through the API without downtime

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Upgrade a Confluence cluster manually without downtime

This document provides step-by-step instructions on how to perform a rolling upgrade on deployments with little or no automation. These instructions are also suitable for deployments based on ourAzure templates.

For an overview of rolling upgrades (including planning and preparation information), seeUpgrade Confluence without downtime.

Step 1: Download upgrade files

Before you start the upgrade, you'll need to download the right Confluence version. You'll be installing this on each node. Remember, you can only upgrade to a higher bug fix version (for example, from Confluence 7.9.0 to 7.9.4) or to the next feature version (for example, from Confluence 7.14.2 to 7.15.0).

Download Confluence

Alternatively, go to> **General Configuration**> **Plan your upgrade** to run the pre-upgrade checks and download a compatible bug fix version.

Step 2: Enable upgrade mode

You need System Administrator global permissions to do this.

To enable upgrade mode:

1. Go to> **General Configuration**> **Rolling upgrades**.
2. Select the **Upgrade mode** toggle (1).

Screenshot: The Rolling upgrades screen.

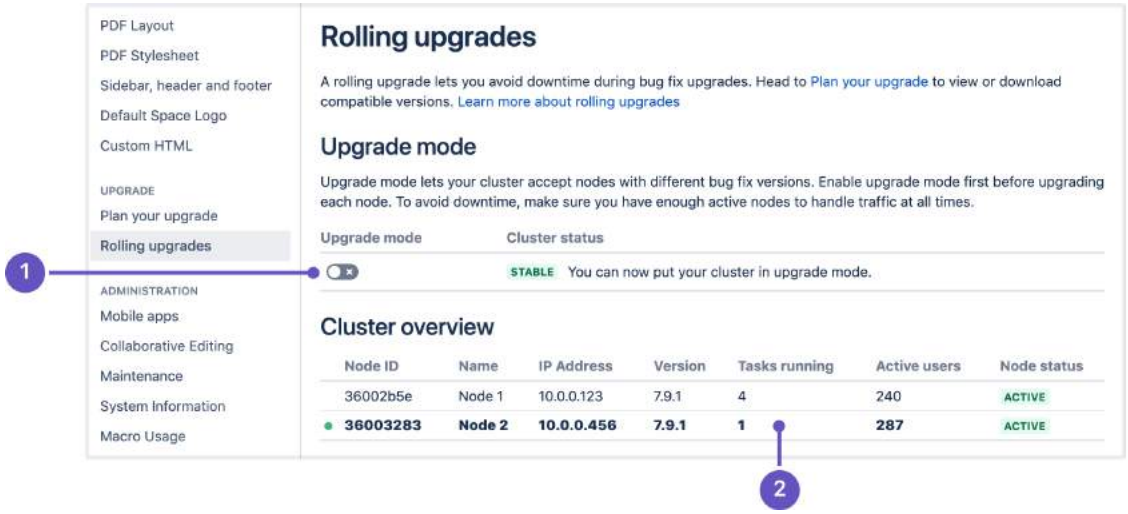

The cluster overview can help you choose which node to upgrade first. The **Tasks running** (2) column shows how many long-running tasks are running on that node, and the **Active users** shows how many users are logged in. When choosing which node to upgrade first, start with the ones with the least number of tasks running and active users.

Upgrade mode allows your cluster to temporarily accept nodes running different Confluence versions. This lets you upgrade a node and let it rejoin the cluster (along with the other non-upgraded nodes). Both upgraded and non-upgraded active nodes work together to keep Confluence available to all users.You can disable upgrade mode as long as you havent upgraded any nodes yet.

Step 3: Upgrade the first node

With upgrade mode enabled, you can now upgrade your first node.

Start with the least busy node

We recommend that you start upgrading the node with the least number of running tasks and active users. You can check this on the Rolling upgrades page.

Start by shutting down Confluence gracefully on the node:

1. Access the node through a command line or SSH.
2. Shut down Confluence gracefully on the node. To do this, run the stop script corresponding to your operating system and configuration. For example, if you installed Confluence as a service on Linux, run the following command:
   `$ sudo /etc/init.d/confluence stop`
   Learn more about graceful Confluence shutdowns
   A graceful shutdown allows the Confluence node to finish all of its tasks first before going offline.During shutdown, the node's status will be**Terminating**, and user requests sent to the node will be redirected by the load balancer to other Active nodes.

   For nodes running on Linux or Docker, you can also trigger a graceful shutdown through the `kill` command (this will send a `SIGTERM` signal directly to the Confluence process).

3. Wait for the node to go offline. You can monitor its status on the Node status column of the Rolling upgrade pages Cluster overview section.

Once the status of the node is offline, you can start upgrading the node. Copy the Confluence installation file you downloaded to the local file system for that node.

To upgrade the first node:

1. Extract (unzip) the files to adirectory (this will be your new installation directory, and must be different to your existing installation directory)
2. Update the following line in the`<Installation-Directory>\confluence\WEB-INF\classes\confluence-init.properties`file to point to the**existinglocal home**directory on that node.
3. If your deployment uses a MySQL database, copy the jdbc driver jar file from your existing Confluence installation directory to `confluence/WEB-INF/lib`in your new installation directory.
   The jdbc driver will be located in either the`<Install-Directory>/common/libor<Installation-Directory>/confluence/WEB-INF/lib`directories. SeeDatabase Setup For MySQL for more details.
4. If you run Confluence as a service:
   - On Windows,delete the existing service then re-install the serviceby running`<install-directory>/bin/service.bat`.
   - On Linux, update the service to point to the new installation directory (or use symbolic links to do this).
5. Copy any other immediately required customizations from the old version to the new one(for example if you are not running Confluence on the default ports or if you manage users externally, you'll need to update / copy the relevant files - find out more inUpgrading Confluence Manually).

   If you configured Confluence to run as a Windows or Linux service, don't forget to update its service configuration as well. For related information, see Start Confluence Automatically on Windows as a Service or Run Confluence as a systemd service on linux.

6. Start Confluence, and confirm that you can log in and view pages before continuing to the next step.

As soon as the first upgraded node joins the cluster, your cluster status will transition to Mixed. This means that you wont be able to disable Upgrade mode until all nodes are running the same version.

Upgrade Synchrony (optional)

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

If you've chosen to let Confluence manage Synchrony for you (recommended), you don't need to do anything. Synchrony was automatically upgraded with Confluence.

If you're running your own Synchrony cluster, grab the new `synchrony-standalone.jar` from the `<local-home>` directory on your upgraded Confluence node. Then, perform the following steps on each Synchrony node:

1. Stop Synchrony on the node using either the `start-synchrony.sh` (for Linux) or `start-synchrony.bat` (for Windows) file from the Synchrony home directory.
2. Copy the new `synchrony-standalone.jar` to your Synchrony home directory.
3. Start Synchrony as normal.

See Set up a Synchrony cluster for Confluence Data Center for related information.

Step 4: Upgrade all other nodes individually

After starting the upgraded node, wait for its status to change to Active in the Cluster overview. At this point you should check the application logs for that node, and log in to Confluence on that node to make sure everything is working. It's still possible to roll back the upgrade at this point, so taking some time to test is recommended.

Once you've tested the first node, you can start upgrading another node, following the same steps. Do this for each remaining node as always, we recommend that you upgrade the node with the least number of running tasks each time.

Step 5: Finalize the upgrade

The steps to finalize your upgrade will differ slightly depending on whether you are upgrading to a bugfix version, or to the next feature version which may require upgrade tasks to be run. You should do this soon as possible, as some tasks are put on hold while your cluster is in upgrade mode.

Finalize upgrade to a bugfix version

To finalize the upgrade:

1. Wait for the cluster status to change to **Ready to finalize**. This won't happen until all nodes are active, and running the same upgraded version.
2. Select the **Finalize upgrade** button.
3. Wait for confirmation that the upgrade is complete. The cluster status will change to **Stable**.

Your upgrade is now complete.

Finalize upgrade to a feature version

To finalize the upgrade:

1. Wait for the cluster status to change to **Ready to run upgrade tasks**. This won't happen until all nodes are active, and running the same upgraded version.
2. Select the **Run upgrade tasks and finalize upgrade** button.
3. One node will start running upgrade tasks. Tail the logs on this node if you want to monitor the process.
4. Wait for confirmation that the upgrade is complete. The cluster status will change to **Stable**.

Your upgrade is now complete.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

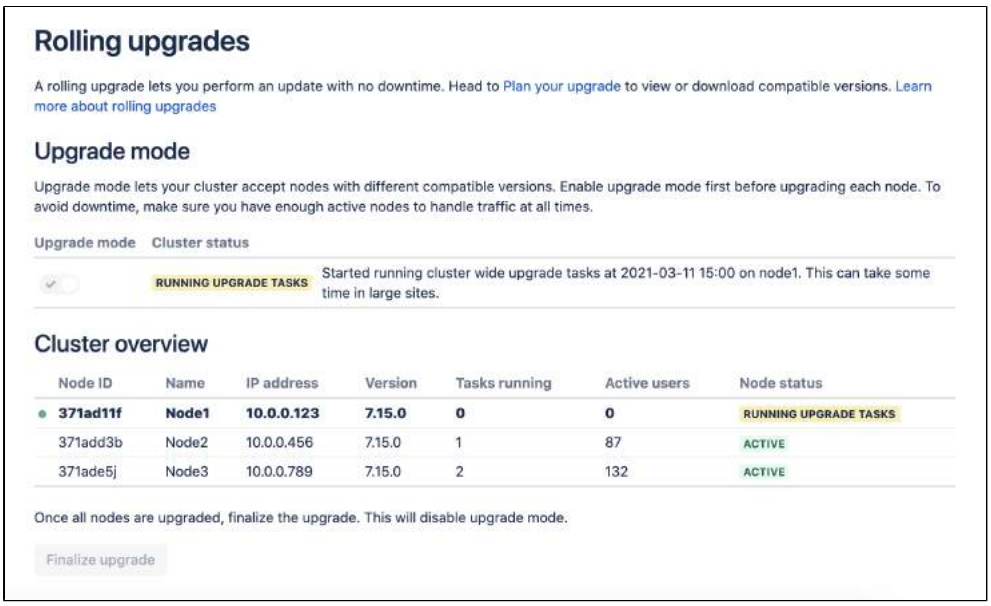

Screenshot: One cluster node running upgrade tasks for the whole cluster.

Upgrade tasks make any required changes to your database and file system, for example changing the database schema or the way index files are stored in the local home directories.

There are a few things you should know about upgrade tasks:

- One cluster node will run the upgrade tasks on the database and other nodes. If there's a problem, logs will be written to the application log on this node.
- The status of other nodes in the cluster may change to **Running upgrade tasks** momentarily to indicate that an upgrade task is making a change to the file system on that node. The node actually running the upgrade tasks does not change.
- Depending on the the size or complexity of your data, some upgrade tasks can take several hours to complete. We generally include a warning in the upgrade notes for the particular version if an upgrade task is likely to take a significant amount of time.
- It's not necessary to direct traffic away from the node running upgrade tasks, but if you know the upgrade tasks are likely to be significant, you may want to do this to avoid any performance impact.

Troubleshooting

Node errors during rolling upgrade

If a nodes status transitions to **Error**, it means something went wrong during the upgrade. You cant finish the rolling upgrade if any node has an **Error** status. However, you can still disable Upgrade mode as long as the cluster status is still **Ready to upgrade**.

There are several ways to address this:

- Shut down Confluence gracefully on the node. This should disconnect the node from the cluster, allowing the node to transition to an **Offline** status.
- If you cant shut down Confluence gracefully, shut down the node altogether.

Once all active nodes are upgraded with no nodes in Error, you can finalize the rolling upgrade. You can investigate any problems with the problematic node afterwards and re-connect it to the cluster once you address the error.

Upgrade tasks failed error

If the cluster status changes to **Upgrade tasks failed**, this means that one or more upgrade tasks did not complete successfully and the upgrade has not been finalized. You should:

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

1. Check the application log on the node running the upgrade task for errors. The node identifier is included in the cluster status message.
2. Resolve any obvious issues (such as file system permissions, or network connectivity problems)
3. Select **Re-run upgrade tasks and finalize upgrade** to try again.

If upgrade tasks are still failing, and you can't identify a cause, you should contact our Support team for assistance. You may also want to roll back the upgrade at this point. We don't recommend leaving Confluence in upgrade mode for a prolonged period of time.

Roll back a node to its original version

How you roll back depends on the upgrade stage you have reached. See Roll back a rolling upgrade for more information.

> **Mixed status with Upgrade mode disabled**
>
> If a node is in an Error state with Upgrade mode disabled, you can't enable Upgrade mode. Fix the problem or remove the node from the cluster to enable Upgrade mode.

Disconnect a node from the cluster through the load balancer

If a node error prevents you from gracefully shutting down Confluence, try disconnecting it from the cluster through the load balancer. The following table provides guidance how to do so for popular load balancers.

NGINX	NGINX defines groups of cluster nodes through the upstream directive . To prevent the load balancer from connecting to a node, delete the node's entry from its corresponding upstream group. Learn more about the upstream directive in the ngx_http_upstream_module module.
HAProxy	With HAProxy, you candisable all traffic to the node by putting it in a `maint` state: `set server <node IP or hostname> state maint` Learn more about forcing a server's administrative state.
Apache	You can disable a node (or "worker") by setting its activation member attribute to disabled. Learn more about advanced load balancer worker properties in Apache.
Azure Application Gateway	We provide adeployment template for Confluence Data Center on Azure; this template uses the Azure Application Gateway as its load balancer. The Azure Application Gateway defines each node as a target within a backend pool. Use the **Edit backend pool** interface to remove your node's corresponding entry. Learn more about adding (and removing) targets from a backend pool.

Traffic is disproportionately distributed during or after upgrade

Some load balancers might use strategies that send a disproportionate amount of active users to a newly-upgraded node. When this happens, the node might become overloaded, slowing down Confluence for all users logged in to the node.

To address this, you can also temporarily disconnect the node from the cluster. This will force the load balancer to re-distribute active users between all other available nodes. Afterwards, you can add the node again to the cluster.

Node won't start up

If a node is Offline or Starting for too long, you may have to troubleshoot Confluence on the node directly. See C onfluence Startup Problems Troubleshooting for related information.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Upgrade a Confluence cluster on AWS without downtime

This document provides step-by-step instructions on performing a rolling upgrade on an AWS deployment orchestrated through CloudFormation. In particular, these instructions are suitable for Confluence Data Center deployments based on our AWS Quick Starts.

For an overview of rolling upgrades (including planning and preparation information), seeUpgrade Confluence without downtime.

Step 1: Enable upgrade mode

You need System Administrator global permissions to do this.

To enable upgrade mode:

1. Go to > **General Configuration**> **Rolling upgrades**.
2. Select the **Upgrade mode** toggle (1).

Screenshot: The Rolling upgrades screen.

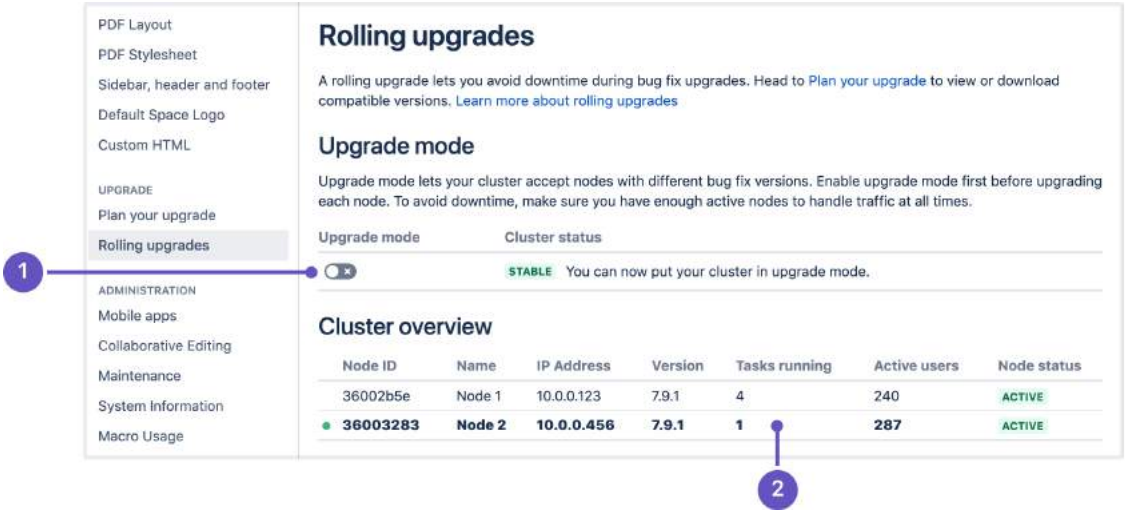

The cluster overview can help you choose which node to upgrade first. The **Tasks running** (2) column shows how many long-running tasks are running on that node, and the **Active users** shows how many users are logged in. When choosing which node to upgrade first, start with the ones with the least number of tasks running and active users.

Upgrade mode allows your cluster to temporarily accept nodes running different Confluence versions. This lets you upgrade a node and let it rejoin the cluster (along with the other non-upgraded nodes). Both upgraded and non-upgraded active nodes work together to keep Confluence available to all users.You can disable upgrade mode as long as you havent upgraded any nodes yet.

Step 2: Find all the current application nodes in your stack

In AWS, note the Instance IDs of all running application nodes in your stack. These are all the application nodes running your current version. You'll need these IDs for a later step.

1. In the AWS console, go to**Services > CloudFormation**. Select your deployments stack to view its Stack Details.
2. Expand the **Resources** drop-down. Look for the **ClusterNodeGroup** and click its Physical ID. This will take you to a page showing the Auto Scaling Group details of your application nodes.
3. In the Auto Scaling Group details, click on the **Instances** tab. Note all of the Instance IDs listed there; you'll be terminating them at a later step.

Step 3: Update your CloudFormation template

Your deployment uses a CloudFormation template that defines each component of your environment. In this case, upgrading Confluence means updating the version of Confluence used in the template. During the upgrade, we highly recommend that you add a node temporarily to your cluster as well.

1. In the AWS console, go to **Services** > **CloudFormation**. Select your deployments stack to view its Stack Details.
2. In the Stack Details screen, click **Update Stack**.
3. From the Select Template screen, select **Use current template** and click **Next**.
4. Set the **Version** parameter to the version youre updating to. Since this is a rolling upgrade, you can only set this to a later bug fix version.
5. Add an extra node to your cluster. This will help ensure that your cluster won't have a shortage of nodes for user traffic. To do this, increase the value of the following parameters by 1:
   - **Maximum number of cluster nodes**
   - **Minimum number of cluster nodes**
6. Select **Next**. Click through the next pages, and then to apply the change using the **Update** button.

After updating the stack, you will have one extra node already running the new Confluence version. With Upgrade mode enabled, that node will be allowed to join the cluster and start work. Your other nodes won't be upgraded yet.

As soon as the first upgraded node joins the cluster, your cluster status will transition to Mixed. This means that you wont be able to disable Upgrade mode until all nodes are running the same version.

Once the new upgraded node is running an in an Active state, you shouldcheck the application logs for that node, and log in to Confluence on that node to make sure everything is working. It's still possible to roll back the upgrade at this point, so taking some time to test is recommended.

Once you've tested the first node, you can start upgrading another node. To do that, shut down and terminate the node AWS will then replace the node with a new one running the updated Confluence version.

Step 4: Upgrade another node

> **Start with the least busy node**
>
> We recommend that you start upgrading the node with the least number of running tasks and active users. On the Rolling upgrades page, youll find both in the Cluster overview section.

In Step 2, you noted the instance ID of each node in your cluster. Terminate the node where you gracefully shut down Confluence. To do this:

1. In the AWS console, go to **Services** > **EC2**. From there, click **Running Instances**.
2. Check the instance of matching the node where you gracefully shut down Confluence.
3. From the **Actions** drop-down, select Instance **State** > **Terminate**.
4. Click through to terminate the instance.

Each time you terminate a node, AWS will automatically replace it. The replacement will be running the new version of Confluence. Once the new node's status is Active, you can move on to upgrading another node.

Step 5: Upgrade all other nodes individually

At this point, your cluster should have two nodes running the new version of Confluence. You can now upgrade other nodes. To do so, simply repeat the previous step on another node.As always, we recommend that you upgrade the node with the least number of running tasks each time.

> If your deployment uses standalone Synchrony, you may need to update the version used by each Synchrony node as well. To do this, terminate each Synchrony node one after the other after you upgrade all nodes to the new version.

Step 6: Finalize the upgrade

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

The steps to finalize your upgrade will differ slightly depending on whether you are upgrading to a bugfix version, or to the next feature version which may require upgrade tasks to be run. You should do this soon as possible, as some tasks are put on hold while your cluster is in upgrade mode.

Finalize upgrade to a bugfix version

To finalize the upgrade:

1. Wait for the cluster status to change to **Ready to finalize**. This won't happen until all nodes are active, and running the same upgraded version.
2. Select the **Finalize upgrade** button.
3. Wait for confirmation that the upgrade is complete. The cluster status will change to **Stable**.

Your upgrade is now complete.

Finalize upgrade to a feature version

To finalize the upgrade:

1. Wait for the cluster status to change to **Ready to run upgrade tasks**. This won't happen until all nodes are active, and running the same upgraded version.
2. Select the **Run upgrade tasks and finalize upgrade** button.
3. One node will start running upgrade tasks. Tail the logs on this node if you want to monitor the process.
4. Wait for confirmation that the upgrade is complete. The cluster status will change to **Stable**.

Your upgrade is now complete.

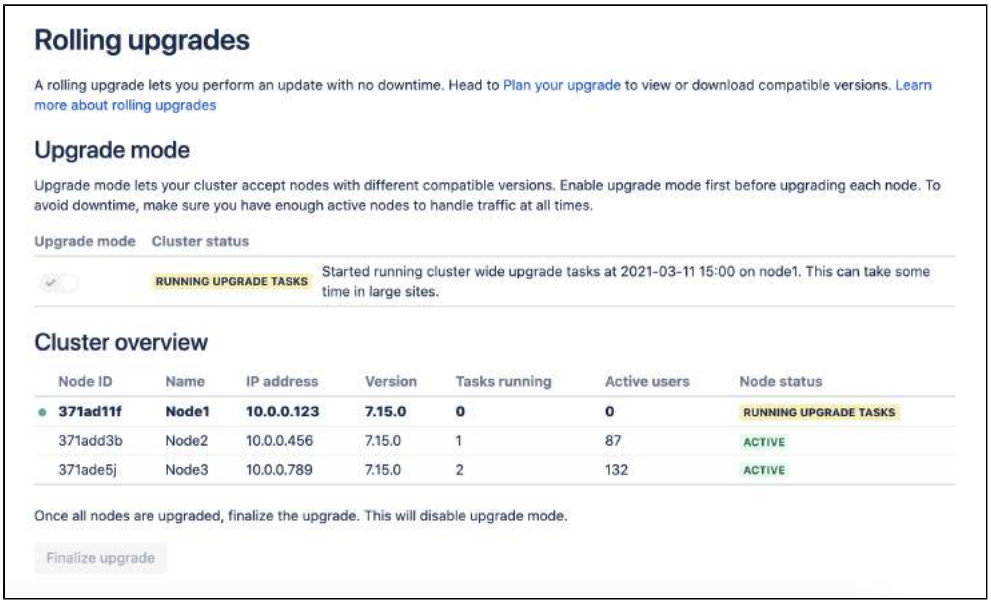

Screenshot: One cluster node running upgrade tasks for the whole cluster.

Upgrade tasks make any required changes to your database and file system, for example changing the database schema or the way index files are stored in the local home directories.

There are a few things you should know about upgrade tasks:

- One cluster node will run the upgrade tasks on the database and other nodes. If there's a problem, logs will be written to the application log on this node.
- The status of other nodes in the cluster may change to **Running upgrade tasks** momentarily to indicate that an upgrade task is making a change to the file system on that node. The node actually running the upgrade tasks does not change.
- Depending on the the size or complexity of your data, some upgrade tasks can take several hours to complete. We generally include a warning in the upgrade notes for the particular version if an upgrade task is likely to take a significant amount of time.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

- It's not necessary to direct traffic away from the node running upgrade tasks, but if you know the upgrade tasks are likely to be significant, you may want to do this to avoid any performance impact.

Step 7: Scale down your cluster

In Step 3, we added a node temporarily to the cluster as a replacement for each one we terminated. This was to help ensure we'd have enough nodes to handle normal user traffic. After finalizing the upgrade, you can remove that node:

1. In the AWS console, go to**Services** > **CloudFormation**. Select your deployments stack to view its **Stack Details**.
2. In the Stack Details screen, click**Update Stack**.
3. From theSelect Templatescreen, select**Use current template**and select**Next**.
4. Decrease the value of the following parameters by 1:
    - **Maximum number of cluster nodes**
    - **Minimum number of cluster nodes**
5. Select **Next**. Click through the next pages, and then to apply the change using the**Update**button.

You can now remove one node from your cluster without AWS replacing it. To do this:

- Choose the node with the least number of running tasks.
- Shut down Confluence gracefully on the node.
- Terminate the node.

Refer to Step 4 for detailed instructions.

Troubleshooting

Disconnect a node from the cluster through the load balancer

If an error prevents you from terminating a node, try disconnecting the node from the cluster through the load balancer. In the AWS Application Load Balancer, each node is registered as a target so to disconnect a node, you'll have to de-register it. For more information on how to do this, seeTarget groups for your Application Load Balancers andRegistered targets.

Traffic is disproportionately distributed during or after upgrade

Some load balancers might use strategies that send a disproportionate amount of active users to a newly-upgraded node. When this happens, the node might become overloaded, slowing down Confluence for all users logged in to the node.

To address this, you can also temporarily disconnect the node from the cluster. This will force the load balancer to re-distribute active users between all other available nodes. Afterwards, you can add the node again to the cluster.

Node errors during rolling upgrade

If a nodes status transitions to **Error**, it means something went wrong during the upgrade. You cant finish the rolling upgrade if any node has an **Error** status. However, you can still disable Upgrade mode as long as the cluster status is still **Ready to upgrade**.

There are several ways to address this:

- Shut down Confluence gracefully on the node. This should disconnect the node from the cluster, allowing the node to transition to an **Offline** status.
- If you cant shut down Confluence gracefully, shut down the node altogether.

Once all active nodes are upgraded with no nodes in Error, you can finalize the rolling upgrade. You can investigate any problems with the problematic node afterwards and re-connect it to the cluster once you address the error.

Roll back to the original version

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

How you roll back depends on the upgrade stage you have reached. See Roll back a rolling upgrade for more information.

Mixed status with Upgrade mode disabled

If a node is in an Error state with Upgrade mode disabled, you can't enable Upgrade mode. Fix the problem or remove the node from the cluster to enable Upgrade mode.

Node won't start up

If a node is Offline or Starting for too long, you may have to troubleshoot Confluence on the node directly. See Confluence Startup Problems Troubleshooting for related information.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Upgrade a Confluence cluster through the API without downtime

This document provides guidance on how to initiate and finalize a rolling upgrade through API calls. This upgrade method is suitable for admins with the skills and automation tools to orchestrate maintenance tasks (like upgrades).

For an overview of rolling upgrades (including planning and preparation information), seeUpgrade Confluence without downtime.

API reference

The entire rolling upgrade process is governed by the following API:

```
http://<host>:<port>/rest/zdu/cluster/zdu/
```

ThisAPI has the following calls:

/zdu	Get an overview of the cluster's status.
/zdu/start	Enable upgrade mode.
/zdu/state	Get the status of the cluster.
/zdu/nodes/{nodeId}	Get an overview of a node's status, including the number of running tasks.
/zdu/cancel	Disable upgrade mode. You can only use this call if the upgrade progress is not MIXED.
/zdu/approve	Once all nodes are upgraded, finalize the rolling upgrade. This will automatically disable upgrade mode.

For detailed information about each API call, seeConfluence REST API Documentation.

Initiating a rolling upgrade

To initiate a rolling upgrade, enable rolling upgrade first. To do this, use:

```
http://<host>:<port>/rest/zdu/cluster/zdu/start
```

Upgrade mode allows your cluster to temporarily accept nodes running different Confluence versions. This lets you upgrade a node and let it rejoin the cluster (along with the other non-upgraded nodes). Both upgraded and non-upgraded active nodes work together to keep Confluence available to all users.You can disable upgrade mode as long as you havent upgraded any nodes yet.

Upgrading each node individually

Before you upgrade a node, you'll need to gracefully shut down Confluence on it.To do this, run the stop script corresponding to your operating system and configuration.Learn more about graceful Confluence shutdowns.

For example, if you installed Confluence as a service on Linux, run the following command:

```
$ sudo /etc/init.d/confluence stop
```

After upgrading Confluence on the node, wait for it to transition to an Active status first before upgrading another node.

Node statuses

To get the status of a node, use:

`http://<host>:<port>/rest/zdu/cluster/zdu/nodes/<nodeID>`

ACTIVE	Confluence is connected to the cluster and running with no errors.
STARTING	Confluence is still loading, and should transition to Active once finished.
TERMINATING	Confluence was gracefully shut down, and should transition to Offline once finished.
OFFLINE	Confluence is not responding on the node. This node will be removed from the cluster completely if it is still offline after Upgrade mode is disabled.
ERROR	Something went wrong with Confluence on the node.

Cluster statuses

To get the status of the cluster, use:

`http://<host>:<port>/rest/zdu/cluster/zdu/state`

STABLE	You can turn on Upgrade mode now.
READY_TO_UPGRADE	Upgrade mode is enabled, but no nodes have been upgraded yet. You can start upgrading your first node now.
MIXED	At least one node is upgraded, but you haven't finished upgrading all nodes yet. Your cluster has nodes running different Confluence versions. You need to upgrade all nodes to the same bug fix version to transition to the next status (**READY_TO_RUN_UPGRADE_TASKS**).
READY_TO_RUN_UPGRADE_TASKS	All nodes have node been upgraded. You can now finalize the rolling upgrade: `http://<host>:<port>/rest/zdu/cluster/zdu/approve`

Enable and disable Upgrade mode

How you roll back depends on the upgrade stage you have reached. See Roll back a rolling upgrade for more information.

Mixed status with Upgrade mode disabled

If a node is in an Error state with Upgrade mode disabled, you can't enable Upgrade mode. Fix the problem or remove the node from the cluster to enable Upgrade mode.

Troubleshooting

Node errors during rolling upgrade

If a nodes status transitions to **Error**, it means something went wrong during the upgrade. You cant finish the rolling upgrade if any node has an **Error** status. However, you can still disable Upgrade mode as long as the cluster status is still **Ready to upgrade**.

There are several ways to address this:

- Shut down Confluence gracefully on the node. This should disconnect the node from the cluster, allowing the node to transition to an **Offline** status.
- If you cant shut down Confluence gracefully, shut down the node altogether.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Once all active nodes are upgraded with no nodes in Error, you can finalize the rolling upgrade. You can investigate any problems with the problematic node afterwards and re-connect it to the cluster once you address the error.

Disconnecting a node from the cluster through the load balancer

If a node error prevents you from gracefully shutting down Confluence, try disconnecting it from the cluster through the load balancer. The following table provides guidance how to do so for popular load balancers.

NGINX	NGINX defines groups of cluster nodes through the upstream directive . To prevent the load balancer from connecting to a node, delete the node's entry from its corresponding upstream group. Learn more about the upstream directive in the ngx_http_upstream_module module.
HAProxy	With HAProxy, you candisable all traffic to the node by putting it in a `maint` state: `set server <node IP or hostname> state maint` Learn more about forcing a server's administrative state.
Apache	You can disable a node (or "worker") by setting its activation member attribute to disabled. Learn more about advanced load balancer worker properties in Apache.
Azure Application Gateway	We provide adeployment template for Confluence Data Center on Azure; this template uses the Azure Application Gateway as its load balancer. The Azure Application Gateway defines each node as a target within a backend pool. Use the **Edit backend pool** interface to remove your node's corresponding entry. Learn more about adding (and removing) targets from a backend pool.

Traffic is disproportionately distributed during or after upgrade

Some load balancers might use strategies that send a disproportionate amount of active users to a newly-upgraded node. When this happens, the node might become overloaded, slowing down Confluence for all users logged in to the node.

To address this, you can also temporarily disconnect the node from the cluster. This will force the load balancer to re-distribute active users between all other available nodes. Afterwards, you can add the node again to the cluster.

Node won't start up

If a node is Offline or Starting for too long, you may have to troubleshoot Confluence on the node directly. See Confluence Startup Problems Troubleshooting for related information.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Roll back a rolling upgrade

On this page:

- Roll back a rolling upgrade - manual and Azure deployments
- Roll back a rolling upgrade - AWS deployments

> The steps on this page only apply if you have used the rolling upgrade method to upgrade Confluence.

If something goes wrong during a rolling upgrade, you may be able to roll back to the original version.

How you roll back depends on the upgrade stage you have reached, and also how you deploy Confluence. To check the current cluster status go to > **General Configuration**> **Rolling upgrades**.

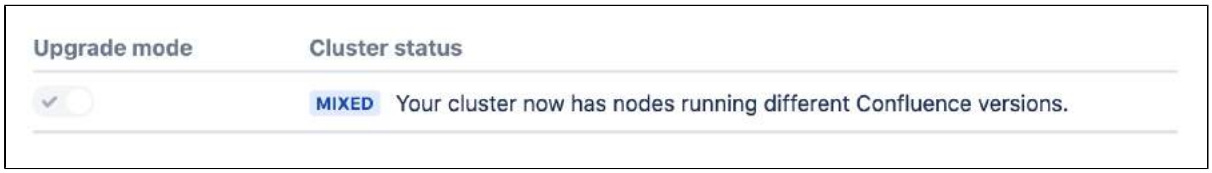

Screenshot: Cluster status in the Rolling upgrades screen

Cluster upgrade status	Action to take
READY TO UPGRADE	In the rolling upgrades screen, disable **Upgrade mode**.
MIXED READY TO RUN UPGRADE TASKS	Follow the steps below to roll back the application version on each upgraded node: • Rollback steps - manual and Azure deployments • Rollback steps - AWS deployments • Rollback steps - Kubernetes deployments Once all nodes are running the original version, the cluster status will change to Ready to upgrade.
RUNNING UPGRADE TASKS UPGRADE TASKS FAILED	You can't roll back once final cluster-wide upgrade tasks have started running. If you stop Confluence on the node running the upgrade task, another node will pick up where the stopped node left off. If upgrade tasks fail, you will need to investigate the problem, then re-run upgrade tasks from the rolling upgrades screen.
COMPLETE	You can't roll back to an earlier version, because the upgrade was finalized successfully.

Roll back a rolling upgrade - manual and Azure deployments

These instructions assume your original install directory is still available. If it's not, you may need to restore it from your backup. You don't need to restore the local home directory.

To roll back an upgraded node to its original version:

1. Access the node through a command line or SSH.
2. Shut down Confluence gracefully on the node.
3. Wait for the node to go offline. You can monitor its status on the **Node status** column of the Rolling upgrade pages Cluster overview section.
4. If you run Confluence as a service:

- On Windows,delete the new service then re-install the old serviceby running<old-install-directory>/bin/service.bat.
- On Linux, update the service to point to the old installation directory (or use symbolic links to do this).

5. Start Confluence (from the original install directory) on the node. You should not see the setup wizard.

Once all nodes are running the same version, the clusters status will revert back to Ready to upgrade. You can then turn off Upgrade mode.

Roll back a rolling upgrade - AWS deployments

To roll back the upgraded nodes to the original version:

In the AWS console, go toServices > CloudFormation. Select your deployments stack to view its Stack Details.

1. In the **Stack Details** screen, select**Update Stack**.
2. From the**Select Template**screen, select**Use current template**and select**Next**.
3. Set the**Version**parameter to your original version.
4. Select**Next**. Click through the next pages, and then to apply the change using the**Update**button.

Afterwards, terminate all nodes running the new version of Confluence. AWS will replace each with a node running the original version.Once all nodes are running the same version, the clusters status will revert back to Ready to upgrade. This will also allow you to disable Upgrade mode.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Upgrade task troubleshooting

When introducing a new feature, or making a significant change to your application, we sometimes need to transform existing data in your database or index, or change the way some data is stored.

Here's a simple example. If we stored "Sydney Australia" in a `location` column in the database, but later decide to store city and country information separately, we might use an upgrade task to take the existing data in the `location` column, and split it into a `city` and `country` column, containing "Sydney" and "Australia" respectively. Actual upgrade tasks are rarely this simplistic, but you get the idea.

On this page:

- When are upgrade tasks run?
- Troubleshooting failed upgrade tasks
- Known issues

We don't include changes that require upgrade tasks in bug fix releases, but they can be quite common in feature and platform releases. You can tell if a version has an upgrade task if the build number is different to your current version.

When are upgrade tasks run?

This depends on the type of upgrade task, and whether you are upgrading with or without downtime.

Rolling upgrade without downtime

If you're performing a rolling upgrade:

- node-specific upgrade tasks happen just prior to the application starting up on that node.
- cluster-wide upgrade tasks happen when all nodes are running the new version, and you select Run upgrade tasks and finalize upgrade in the rolling upgrades screen.

During a rolling upgrade, there's a short time where data might exist in both old and new formats.Cluster-wide upgrade tasks include tasks that transform data in the database, and may also include changes to the shared home and local home directories on each node. These tasks require all nodes to have been upgraded before they can be run.

Using the example above, a node that has not yet been upgraded would continue to write to the `location` column, while an upgraded node would write to the new `city` and `country` columns, as well as the old `location` column (to prevent data loss if you need to roll back). Once all nodes are upgraded, it's safe for us to split the existing data in the `location` field into the new `city` and `country` fields. This is often the part of the upgrade task that can take some time, depending on how much data you have.

Upgrade with downtime

When upgrading a non-clustered deployment, upgrade tasks are usually run just prior to the application starting up after the upgrade.

When upgrading a cluster with downtime (not a rolling upgrade), cluster-wide upgrade tasks are run when the first node starts up.

Troubleshooting failed upgrade tasks

If an upgrade task fails, there are a number of things you'll need to do to resolve the issue.

Check the application logs

The first step is to check the application logs. If you're running Confluence in a cluster, you may need to check the logs on more than one node.

Sometimes the cause will be obvious, such as a network timeout, not enough disk space in the local or shared home directory, or the database user / confluence user has inadequate permissions to complete the action.

Check your database configuration

The most common reason upgrade tasks fail include:

- database user does not have adequate permissions to perform the required action
- database configuration is incorrect (for example the character set or encoding)
- database version or edition not supported.

This usually results in a database error being written to the application logs.Check for KB articles about the specific problem, and confirm your setup matches the database setup specified in our documentation.

Re-run upgrade tasks

Once you've resolved any issues, you'll need to re-run the upgrade tasks. How you do this depends on whether you are upgrading with or without downtime.

- If you're performing a rolling upgrade, re-start the application on any failed nodes, then select Re-run upgrade tasks and finalize upgrade.
- If you're upgrading with downtime, re-start the application. Upgrade tasks will run prior to the application starting up.

Don't leave your application in upgrade mode

If you're performing a rolling upgrade, it's important you don't leave your cluster in an upgrade mode longer than is necessary. This is because there may be data that needs to be handled in multiple ways until the final upgrade tasks can be run.

Known issues

- Some non-enterprise editions of Microsoft SQL Server don't support online index creation. If an upgrade task needs to acquire an exclusive table lock, you may experience some performance degradation or downtime. We'll warn you if we detect that you database edition may be affected.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Supported Platforms

This page describes the additional software and infrastructure you'll need to run Confluence. Please review this info before installing Confluence. The information on this page applies to**Confluence Server and Data Center 8.4**.

- You should only use Confluence with a supported platform. Any platforms and versions not listed on this page are unsupported, which means we don't test, fix bugs or provide assistance.
- SeeEnd of Support Announcements for Confluencefor upcoming changes to supported platforms.
- Go to> **General Configuration** > **Troubleshooting and support tools** to check your instance health. It looks at things like your license validity, Tomcat version, basic database setup and more.

Definitions:

Supported - you can use**Confluence Server and Data Center 8.4**with this platform.

Limited - you can evaluate Confluence on this platform, but you can't use it to run a production Confluence site.

Deprecated - support for this platform will end in an upcoming release. SeeEnd of Support Announcements for Confluence.

On this page:

- Browsers
- Operating systems
- Databases
- Java
- Object storage
- Infrastructure

Related pages:

- Confluence Installation Guide
- Confluence Setup Guide
- Server Hardware Requirements Guide
- Supported Platforms FAQ

Browsers

Desktop browsers

Microsoft Edge (Chromium)

Chrome

Firefox

Safari (Mac only)

Mobile browsers

Chrome

Firefox

Safari (iOS only)

Android WebView

Mobile operating system (required for mobile app)

iOS 11 or later

Android 4.4 (KitKat) or later

Good to know:

- The Confluence setup wizard requires Javascript to be enabled while installing Confluence.Learn more
- Parts of Confluence won't display correctly if your browser window size is less than1024x768.
- Although some supported browsers may allow you to enable Adobe Flash in their advanced settings, we recommend leaving Flash disabled, as enabling it may expose you to security vulnerabilities.

Operating systems

Operating systems

Microsoft Windows

Known issues:

- The following operating system variants can't be used with Confluence:

Linux (most distributions)	• Windows Nano • Alpine Linux (3.5 and earlier)
MacOS / OSX (evaluation only)	**Good to know:** • You can run Confluence on 32bit or 64bit operating systems, but we only provide installers for 64bit operating systems. • You can evaluate Confluence on MacOS / OSX, but you can't install and run your production Confluence site on a Mac.

Databases

PostgreSQL	**Known issues:**

PostgreSQL 12	• Confluence will not work on MySQL variants such as:
	  • MariaDB - see CONFSERVER-29060
	  • Percona Server - see CONFSERVER-36471

PostgreSQL 13	
PostgreSQL 14	**Good to know:**
Amazon Aurora	• Amazon Aurora and Azure PostgreSQL are only supported with Confluence Data Center. • You can use Amazon's Relational Database Service (RDS) for the supported databases listed on this page. • The only supported Amazon Aurora config is a PostgreSQL-compatible clustered database with one writer replicating to zero or more readers.Learn more
PostgreSQL 12	
PostgreSQL 13	
PostgreSQL 14	
Azure PostgreSQL	
PostgreSQL 12	
PostgreSQL 13	
PostgreSQL 14	
MySQL	
MySQL 8	
Oracle	
Oracle 19c	
Microsoft SQL Server	
SQL Server 2017	
SQL Server 2019	
Azure SQL	

Java

Oracle JRE/JDK	**Known issues:**
Java 11	• There's a known issue with some Java 11 versions and TLS 1.3. We recommend Java 11.0.8 or later. • We use**Temurin**to replicate issues raised with OpenJDK. If youre using a different distribution of OpenJDK well still provide support for our products. However, if the bug is caused by a problem in Java distribution, you'll need to contact the Java distributor for help.
Java 17	
Temurin (previously AdoptOpenJDK)	
Java 11	**Good to know:**

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

✅ Java 17	• You don't need to install Java if you plan to use the installer to install Confluence, as a Temurin Java 17 JRE is bundled with Confluence. • SeeBundled Tomcat and Java versions to see which Java version was bundled with your Confluence version.

Object storage

S3 object storage	Good to know:
✅ Amazon S3	• Amazon S3 object storage is an optional attachment storage method available to anyone on a Data Center license and running Confluence in AWS. • If you're a new customer, seeS3 object storagefor setup instructions. • If you're an existing customer, you'll need to migrate your attachment data to S3 object storage from the file system or another storage method. SeeAttachment storage configurationfor steps to do this. • Even if you use S3 object storage, other non-attachment data will still be stored in your home directory.

Infrastructure

Hardware:

- You can't run Confluence on SPARC based hardware. You'll need to use x86 hardware or 64bit derivatives of x86 hardware.
- You can't use an NFS mount for your installation or home directory due to Lucene requirements. If you're installing Confluence Data Center, an NFS mount is fine for the shared home directory, but not for the local home directories.

Containerization

- You can useofficial imagesto deploy Confluence in a Docker container, or customize a Docker deployment on your own.
- We support the Atlassian Docker templates and can help with Confluence related problems. We do not provide support for Docker itself or problems with any Docker environment.

Containerization Manager

- You are recommended to useofficial helm chartsto deploy Confluence Data Center using Kubernetes, or customize a Kubernetes deployment on your own with reference to the official helm charts.
- We support the Atlassian Kubernetes helm chart and can help with Confluence Data Center product-related problems. We do not provide support for Kubernetes itself or problems with any Kubernetes environment.
- Read ourKubernetes support disclaimerand more aboutwhat we support and what we dont.

Virtualization:

- You can run Confluence and Confluence Data Center in a virtualized environment (including Docker), but our support team can't assist you with problems related to the environment itself. SeeRunning Confluence in a Virtualized Environment
- Our support team can assist you with deploying Confluence Data Center in AWS using the Cloud Formation Template or Quick Start. We won't be able to assist you if you have significantly customised the Cloud Formation Template.

Application server:

- We only support the Tomcat version that is bundled with your Confluence version. You can't run Confluence in your own application server. SeeBundled Tomcat and Java versions to see which version of Tomcat was bundled with your Confluence version.

Internet protocols:

- You can run Confluence in both IPv4 and IPv6 environments.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

- Raw IPv6 addresses are not always recognized. See the Confluence 6.9 Upgrade Notesfor limitations and known issues.

Operating system support:

- You should only install and use Confluence on operating system versions that have active vendor support. For example, you can use Confluence on any Microsoft supported version of Windows, unless specified otherwise above.

For more information see our Server Hardware Requirements Guide and System Requirements.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

End of Support Announcements for Confluence

This page is where we announce end of support for various platforms, browsers, and information on features that will be discontinued in Confluence Server.

The table below summarizes the end of support announcements for **upcoming**Confluence releases. If a platform (or version) has already reached its end of support date, it is **not** listed in the table.

What will be deprecated?	Confluence end of support
JNDI datasource connections	No support in Confluence 8.0 (announcement)
H2 database	No support in Confluence 8.0 (announcement)
PostgreSQL 11	No support in Confluence 8.0 (announcement)
MySQL 5.7	No support in Confluence 8.0 (announcement)
Java 8	No support in Confluence 8.0 (announcement)
UTF8 encoding for MySQL	No support in Confluence 8.0 (announcement)
Usage stats plugin	Removed in Confluence 8.0(announcement)
Oracle 12c	No support in Confluence 7.20 (announcement)
PostgreSQL 10	No support in Confluence 7.18 (announcement)
Microsoft SQL Server 2016	No support in Confluence 7.17(announcement)
PostgreSQL 9.6	No support in Confluence 7.14 (announcement)
Microsoft Edge 18	No support in Confluence 7.13 (announcement)
View File macros	These macros are now fully supported (updated announcement)

Most recent announcements first:

- Deprecated datasource connection (September 2022)
- Deprecated database for Confluence (July 2022)
- Deprecated Java version (May 2022)
- Deprecated database encoding for MySQL databases (May 2022)
- Deprecated database for Confluence (May 2022)
- Changes to features for Confluence (22 March 2022)
- Deprecated database for Confluence (20 December 2021)
- Deprecated database for Confluence (16 November 2021)
- Deprecated database for Confluence (13 April 2021)
- Deprecated browsers for Confluence (2 February 2021)
- Deprecated databases for Confluence (11 December 2019)
- Deprecated database for Confluence (11 December 2019)
- Deprecated databases for Confluence (14 October 2019)
- Deprecated browsers for Confluence (24 September 2019)
- Deprecated macros for Confluence (12 March 2019)
- Deprecated Gadgets in Confluence (12 March 2019)
- Changes to features in Confluence (12 March 2019)
- Deprecated databases for Confluence (2 October 2018)
- Deprecated databases for Confluence (30 January 2017)
- Deprecated macro for Confluence (31 October 2017)
- Deprecated driver for Microsoft SQL Server
- Deprecated operating system for Confluence (15 May 2017)
- Deprecated mobile browser for Confluence (3 November 2016)
- Changes to Confluence distributions (8 June 2016)
- Deprecated browsers for Confluence (8 June 2016)

- Deprecated databases for Confluence (8 June 2016)
- Deprecated macros for Confluence (13 November 2015)
- Discontinued features for Confluence (10 July 2015)
- Deprecated databases for Confluence (19 May 2015)
- 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 datasource connection (September 2022)

Atlassian will remove support for datasource connections in Confluence 8.0 as part of ongoing security and scale improvements to our product.

If you currently use a JNDI datasource connection, we recommend you connect your database with a JDBC URL instead. This will ensure an easy upgrade experience to future versions of Confluence.

How to convert a datasource to a direct JDBC connection

End of support means that Atlassian will not offer support for, or fix bugs related to, using this method of database connection on Confluence 8.0 or later.

If you have questions or concerns regarding this announcement, please email eol-announcement at atlassian dot com.

Deprecated database for Confluence (July 2022)

Atlassian will not support the following databases in Confluence 8.0:

- MySQL 5.7
- PostgreSQL 11
- H2 embedded database

If you currently use H2 database, make sure you've updated to AMPS 8.6.0 to continue using this database or HSQL, both of which are only available for Data Center testing installations only.

End of support means that Atlassian will not offer support for, or fix bugs related to, running Confluence 8.0 or later with this database.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

- Confluence 7.20 is the last version that will support MySQL 5.7, PostgreSQL 11 and H2 embedded database.
- Confluence 7.20 and earlier versions will continue to work with MySQL 5.7, PostgreSQL 11 and H2 embedded database, however, we will not fix bugs affecting this database after the end-of-life datefor your version of Confluence.
- Confluence 8.0 and later will not be tested with MySQL 5.7, PostgreSQL 11 and H2 embedded database.

Check out theSupported Platformspage for the full list of supported databases.

If you have questions or concerns regarding this announcement, please email eol-announcement at atlassian dot com.

Deprecated Java version (May 2022)

Atlassian will not support Java 8 in Confluence 8.0.

End of support means that Atlassian will not offer support for, or fix bugs related to, running Confluence 8.0 or later on this Java version.

- Confluence 7.x versions will continue to work with Java 8, however, we will not fix bugs affecting this Java version after the end-of-life datefor your version of Confluence.
- Confluence 8.0 and later versions will not be tested with Java 8.

Check out theSupported Platformspage for the list of supported Java versions.

If you have questions or concerns regarding this announcement, please email eol-announcement at atlassian dot com.

Deprecated database encoding for MySQL databases (May 2022)

Atlassian will not support UTF8 database encoding in MySQL databases in Confluence 8.0. If you run Confluence with a supported MySQL database, you should use UTF8MB4 encoding.

End of support means that Atlassian will not offer support for, or fix bugs related to running Confluence 8.0 or later with a MySQL database configured with UTF8 encoding.

If you have questions or concerns regarding this announcement, please email eol-announcement at atlassian dot com.

Deprecated database for Confluence (May 2022)

Atlassian will not support the following database in Confluence 7.20:

- Oracle 12c

End of support means that Atlassian will not offer support for, or fix bugs related to, running Confluence 7.20 or later with this database.

- Confluence 7.19 is the last version that will support Oracle 12c.
- Confluence 7.19 and earlier versions will continue to work with Oracle 12c, however, we will not fix bugs affecting this database after the end-of-life datefor your version of Confluence.
- Confluence 7.20 and later will not be tested with Oracle 12c.

Check out theSupported Platformspage for the full list of supported databases.

If you have questions or concerns regarding this announcement, please email eol-announcement at atlassian dot com.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Changes to features for Confluence (22 March 2022)

We will remove the Usage Stats plugin in Confluence 8.0.

TheConfluence Usage Statsplugin provides basic page view tracking in Confluence Server. We know that usage tracking is important to you, however the current implementation has long been disabled by default as it can have a noticeable impact on your site's performance.

Significant work is required to make this feature compatible with Confluence 8.0, and the functionality is largely superseded by Analytics, which was added in Confluence Data Center 7.11.

Analytics provides a significantly better analytics experience, and will continue to be available for Data Center customers in Confluence 8.0.

If you have questions or concerns, please comment on this issue
~~CONFSERVER-57612~~ - Plans to end support for Usage Stats CLOSED .

Deprecated database for Confluence (20 December 2021)

Atlassian will not support the following database in Confluence 7.18:

- PostgreSQL 10

End of support means that Atlassian will not offer support for, or fix bugs related to, running Confluence 7.18 or later with this database.

- Confluence 7.17 is the last version that will support PostgreSQL 10.
- Confluence 7.17 and earlier versions will continue to work with PostgreSQL 10, however we will not fix bugs affecting this database after theend-of-life datefor your version of Confluence.
- Confluence 7.18 and later will not be tested with PostgreSQL 10.

Check out theSupported Platformspage for the full list of supported databases.

If you have questions or concerns regarding this announcement, please email eol-announcement at atlassian dot com.

Deprecated database for Confluence (16 November 2021)

Atlassian will not support the following database in Confluence 7.17:

- Microsoft SQL Server 2016

End of support means that Atlassian will not offer support for, or fix bugs related to, running Confluence 7.17 or later with this database.

- Confluence 7.16 is the last version that will support Microsoft SQL Server 2016.
- Confluence 7.16 and earlier versions will continue to work with Microsoft SQL Server 2016, however we will not fix bugs affecting this database after theend-of-life datefor your version of Confluence.
- Confluence 7.17 and later will not be tested with Microsoft SQL Server 2016.

Check out theSupported Platformspage for the full list of supported databases.

If you have questions or concerns regarding this announcement, please email eol-announcement at atlassian dot com.

Deprecated database for Confluence (13 April 2021)

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Atlassian will not support the following database in Confluence 7.14:

- PostgreSQL 9.6

End of support means that Atlassian will not offer support for, or fix bugs related to, running Confluence 7.14 or later with this database.

- Confluence 7.13 is the last version that will supportPostgreSQL 9.6.
- Confluence 7.13 and earlier versions will continue to work with PostgreSQL 9.6, however we will not fix bugs affecting this database after theend-of-life datefor your version of Confluence.
- Confluence 7.14 and later will not be tested with PostgreSQL 9.6.

Check out theSupported Platformspage for the full list of supported databases.

If you have questions or concerns regarding this announcement, please email eol-announcement at atlassian dot com.

Deprecated browsers for Confluence (2 February 2021)

In January 2020 Microsoft released a new Microsoft Edge browser based on Chromium. This new version is compatible with all supported Windows versions, and replaces the previous version, now known as Microsoft Edge Legacy. Read more about the difference between the new Microsoft Edge and Microsoft Edge Legacy on the Microsoft support site.

As Microsoft have announced plans to end support for Microsoft Edge Legacy, we have also decided to end support for Microsoft Edge Legacy.

End of support means we will not fix bugs specific to Microsoft Edge Legacy, and will begin to introduce features that aren't compatible with this browser.

When is this happening?

- Confluence 7.12 is the last version to support Microsoft Edge Legacy.
- Confluence 7.13 and subsequent versions will not support Microsoft Edge Legacy.

What this means for you

We recommend switching to one of oursupported browsers, such as the new Microsoft Edge (Chromium), Google Chrome, or Mozilla Firefox.

If you have questions or concerns regarding this announcement, please email eol-announcement at atlassian dot com.

Deprecated databases for Confluence (11 December 2019)

Atlassian will not support the following databases in Confluence 7.5:

- Microsoft SQL Server 2014
- PostgreSQL 9.5

End of support means thatAtlassian will not offer support for, or fix bugs related to, running Confluence 7.5 or later with this database.

- Confluence 7.4 is the last version that will support these databases.
- Confluence 7.4 and earlier versions will continue to work withthese databases, however we will not fix bugs affecting these databases after theend-of-life datefor your version of Confluence.
- Confluence 7.5 and later will not be tested withthese databases.

Check out theSupported Platformspage for the full list of supported databases.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

If you have questions or concerns regarding this announcement, please email eol-announcement at atlassian dot com.

Deprecated database for Confluence (11 December 2019)

Atlassian will not support the following database in Confluence 7.4:

- Microsoft SQL Server 2012

End of support means thatAtlassian will not offer support for, or fix bugs related to, running Confluence 7.4 or later with this database.

- Confluence 7.3 is the last version that will support Microsoft SQL Server 2012.
- Confluence 7.3 and earlier versions will continue to work with Microsoft SQL Server 2012, however we will not fix bugs affecting this database after theend-of-life datefor your version of Confluence.
- Confluence 7.4 and later will not be tested with Microsoft SQL Server 2012.

Check out theSupported Platformspage for the full list of supported databases.

If you have questions or concerns regarding this announcement, please email eol-announcement at atlassian dot com.

Deprecated databases for Confluence (14 October 2019)

Atlassian will not support the following databases in Confluence 7.4:

- PostgreSQL 9.4
- MySQL 5.6
- Oracle 12c R1

End of support means thatAtlassian will not offer support for, or fix bugs related to, running Confluence 7.4 or later with these databases.

- Confluence 7.3 is the last version that will support these databases.
- Confluence 7.3 and earlier versions will continue to work with these databases, however we will not fix bugs affecting these databases after theend-of-life datefor your version of Confluence.
- Confluence 7.4 and later will not be tested with these databases.

Check out theSupported Platformspage for the full list of supported databases.

If you have questions or concerns regarding this announcement, please email eol-announcement at atlassian dot com.

Deprecated browsers for Confluence (24 September 2019)

In 2015 Microsoft released Edge as the browser to supersede Internet Explorer, and in recent times Microsoft has discouraged the use of Internet Explorer as a default browser. To allow us to continue to take advantage of modern web standards to deliver improved functionality and the best possible user experience across all of our products, we have decided to end support for Internet Explorer 11.

End of support means we will not fix bugs specific to Internet Explorer 11, and will begin to introduce features that aren't compatible with this browser.

When is this happening?

- Confluence 7.4 is the last version to support Internet Explorer. Confluence 7.4 will be anEnterprise release.
- Confluence 7.5 and subsequent versions will not support Internet Explorer 11.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

What this means for you

We recommend switching to one of our supported browsers, such as Microsoft Edge, Google Chrome, or Mozilla Firefox.

If you have questions or concerns regarding this announcement, please email eol-announcement at atlassian dot com.

Deprecated macros for Confluence (12 March 2019)

We will end support for the following macros in Confluence 7.0, and hide them from the macro browser.Any existing instances of these macros will still work, but you won't be able to insert these macros into the editor using the macro browser:

- IM Presence macro
- Netwok macro
- Search results macro
- Space details macro

End of support means Atlassian will not fix bugs related to these macros in Confluence 7.0 or later versions. We will remove these macro entirely in a future Confluence release, and will provide more information at that time.

To check whether a macro is used in your site, go to >**General Configuration** >**Macro Usage**. Some macros will be listed under the system app that provides them.

IM Presence macro

TheIM Presence Macroshows when a given user is online in a selected chat service. The macro only supports a small number of chat services, and we feel that most modern chat tools provide better ways to see this information.

If you have questions or concerns, please comment on this issue ~~CONFSERVER-57596~~ CLOSED .

Network macro

TheNetwork Macroallows you display the people a particular user is following, or people who are following that user. Following someone is a useful way to get notifications about their activity, and this network information is also available on each user's profile page.

If you have questions or concerns, please comment on this issue ~~CONFSERVER-57597~~ CLOSED .

Search results macro

TheSearch Results Macroallows you to display the results of a keyword search on a page. We are making some great changes to Search over the next few releases, and have observed that this macro is rarely used.

If you have questions or concerns, please comment on this issue ~~CONFSERVER-57598~~ CLOSED .

Space details macro

TheSpace Details Macroallows you to display basic information about the current space on a page. This information is available at all times from Space Tools > Overview.

If you have questions or concerns, please comment on this issue ~~CONFSERVER-57599~~ CLOSED .

Deprecated Gadgets in Confluence (12 March 2019)

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

We will end support for the following Gadgets in Confluence 7.0,and hide them from the macro browser.Any existing instances of these gadgets will still work, but you won't be able to insert these gadgets into the editor using the macro browser:

- Activity Stream
- Confluence Page Gadget
- Confluence Quick Nav Gadget
- News gadget

End of support means Atlassian will not fix bugs related to these gadgets in Confluence 7.0 or later versions. We will remove these gadgets entirely in a future Confluence release, and will provide more information at that time.

Gadgetswere designed to allow you to display information dynamically from sources like iGoogle or Jira, for example, in Confluence. The first gadgets were introduced in Confluence 3.1, and much of the technology they were based on is now superseded or obsolete. Since then we have also implemented a number of better ways to display dynamic information using macros and other integration points.

Activity Stream gadget

The Activity stream gadget shows a list of recently changed content in your site. We recommend using the Recently Updated macro as an alternative in Confluence.

Confluence Page Gadget

This gadget displays the contents of a Confluence page. We recommend using the Include Page macro as an alternative in Confluence.

Confluence Quick Nav Gadget

This gadget provides a search field that can be used to search for page titles in Confluence. We recommend using the Livesearch macro as an alternative in Confluence.

News gadget

This gadget previously displayed blogs and other news from Atlassian. It has not been displaying content for some time. We will remove this gadget completley in 7.0.

If you have questions or concerns, please comment on this issue ~~CONFSERVER-57614~~ CLOSED .

Changes to features in Confluence (12 March 2019)

Shortcut links

Shortcut Linkswere introduced in Confluence 2.3 and provided a quick way to add links to websites in wiki markup. Shortcut links can only be configured by an administrator, are not easily discoverable, and seldom used by end users.

Thanks for your feedback on CONFSERVER-57610 NOT BEING CONSIDERED

We've heard you, and will not end support for Shortcut links in Confluence 7.0.

Trackback and external referrers

We will remove the trackback and referrers features completley in Confluence 7.0.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Trackbackenables Confluence to send and receive trackback pings when pages are linked to.External Referrers appear on the**Page Information**view of a page, and list clicks from external websites to the page. Trackback is no longer widely used in modern websites, and because it relies on accepting unauthenticated requests to a particular URL, is a spam vector.

If you have questions or concerns, please comment on this issue ~~CONFSERVER-57611~~ CLOSED .

Orphaned pages screen

We will remove the Orphaned pages screen in the default theme in Confluence 7.0.

The Orphaned pages screen provided a list of all pages that Confluence considersorphaned pages(not a child of a space homepage, and not linked to by any other page). Since the introduction of the Confluence 5 default theme, the orphaned pages screen has been less useful because it's always possible to see all pages in a space via**Space Tools**>**Reorder pages**.

If you have questions or concerns, please comment on this issue ~~CONFSERVER-57601~~ CLOSED .

Hipchat integration

We have discontinued development on all chat products. Hipchat Cloud services were shut down in February 2019, and Hipchat Data Center and Server will both reach end of life within the next year.

We will end support for all bundled Hipchat plugins in Confluence 7.0. These will be disabled by default for new installations. This will have no impact on existing installations, and can be easily enabled if required.

End of support means that Atlassian will not fix bugs related to Hipchat integration in Confluence 7.0 or later versions.

If you have questions or concerns, please comment on this issue ~~CONFSERVER-57602~~ CLOSED .

Deprecated databases for Confluence (2 October 2018)

Atlassian will end support for**PostgreSQL 9.3**in Confluence 6.13.End of support means thatAtlassian will not offer support for, or fix bugs related to, running Confluence 6.13 or later with this database.

- Confluence 6.12 is the last version that will support PostgreSQL 9.3.
- Confluence 6.12 and earlier versions will continue to work with PostgreSQL 9.3, however we will not fix bugs affecting these databases after the end-of-life datefor your version of Confluence.
- Confluence 6.13 and later will not be tested with PostgreSQL 9.3.

Check out theSupported Platformspage for the full list of supported databases.

If you have questions or concerns regarding this announcement, please email eol-announcement at atlassian dot com.

Deprecated databases for Confluence (30 January 2017)

Atlassian will end support for **PostgreSQL 9.2** in Confluence 6.8.End of support means thatAtlassian will not offer support for, or fix bugs related to, running Confluence 6.8 or later with this database.

- Confluence 6.7 is the last version that will support PostgreSQL 9.2.
- Confluence 6.7 and earlier versions will continue to work with PostgreSQL 9.2, however we will not fix bugs affecting these databases after the end-of-life datefor your version of Confluence.
- Confluence 6.8 and later will not be tested with PostgreSQL 9.2.

Check out theSupported Platformspage for the full list of supported databases.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

If you have questions or concerns regarding this announcement, please email eol-announcement at atlassian dot com.

Deprecated macro for Confluence (31 October 2017)

We will end support for the **JUnit Report macro** with the release of Confluence 6.6. This macro is used to display the results of JUnit tests on a Confluence page and, based on our research, is rarely used.

End of support means that Atlassian will not fix bugs related to this macro past the support end date for your version of Confluence. We will remove this macro entirely in a future Confluence release, and will provide more information at that time.

To check whether this macro is used in your site, go to > **General Configuration**> **Macro Usage**. The JUnit Report macro will be listed under **Advanced Macros** if it's used.

If you have questions or concerns, please comment on this issue
~~CONFSERVER-53942~~ - Plans to remove JUnit macro CLOSED .

Deprecated driver for Microsoft SQL Server

We are replacing the open source jTDS driver for SQL Server with the official Microsoft JDBC Driver for SQL Server. This new driver is bundled with Confluence 6.4 and later.

Atlassian will end support for the jTDS driver in Confluence 6.6. End of support means that Atlassian will not offer support for, or fix bugs related to, installing and running Confluence 6.6 or later with this driver.

- Confluence 6.5.x will be the last major release to bundle the jTDS driver.
- Confluence 6.5.x and earlier versions will continue to be supported with the jTDS driver, until their support end date.
- Confluence 6.6.x will not bundle or support the jTDS driver. We'll provide plenty of information on how to migrate to the new driver at that time.

If you have questions or concerns regarding this announcement, please email eol-announcement at atlassian dot com.

Deprecated operating system for Confluence (15 May 2017)

Atlassian will end support for the Oracle Solaris operating systemin Confluence 6.3.End of support means thatAtlassian will not offer support for, or fix bugs related to, installing and running Confluence 6.3 or later on this operating system.

- Confluence 6.2.x will be the last major release that can be installed on Solaris.
- Confluence 6.2.x and earlier versions will continue to be supported on Solaris, until their support end date.

Check out theSupported Platformspage for the full list of supported operating systems.

If you have questions or concerns regarding this announcement, please email eol-announcement at atlassian dot com.

Deprecated mobile browser for Confluence (3 November 2016)

Atlassian will end support for the default browser provided with Android4.0.3 (Ice Cream Sandwich)in Confluence 6.0.End of support means thatAtlassian will not fix bugs related to this browser past the support end date,except for security related issues. This means:

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

- Confluence 5.10 will be the last major release that supportsthe default browser provided with Android4.0.3 (Ice Cream Sandwich).
- Confluence 5.10.x and earlier versions will continue to work onthe default browser provided with Android4.0.3 (Ice Cream Sandwich).

With the release of Confluence 6.0 we have added support for the default browser provided with current Android versions from 4.4 (KitKat) and later. Check out theSupported Platformspage for the full list of supported browsers.

If you have questions or concerns regarding this announcement, please email eol-announcement at atlassian dot com.

Changes to Confluence distributions (8 June 2016)

To help us bring you new Confluence Server releases faster, we are considering only providing 64-bit installers. Confluence 5.10 would be the last Confluence release to provide a 32-bit installer.

Q: Can I upgrade using the 64-bit installer?

Yes. If you installed Confluence using the 32-bit installer on a 64-bit operating system, you will be able to upgrade using the 64-bit installer.

Q: What if I am not able to use the 64-bit installer?

We'd love to hear from you to better understand how this change would impact you. Comment on this issue ~~CONFSERVER-42817~~ - Planned deprecation of 32-bit installers CLOSED or contact us directly at`eol-announcement at atlassian dot com`.

Deprecated browsers for Confluence (8 June 2016)

Atlassian will end support for Internet Explorer 10 in Confluence 6.0.End of support means thatAtlassian will not fix bugs related to Internet Explorer 10 past the support end date,except for security related issues.

This changeallows us to use more modern browser technologies to give you the best user experience in Confluence.Check out theSupported Platformspage for the full list of supported browsers.

If you have questions or concerns regarding this announcement, please email eol-announcement at atlassian dot com.

Internet Explorer 10 (IE10) end of support notes

- Confluence 5.10 will be the last major release that supportsInternet Explorer 10.
- Confluence 5.10.x and earlier versions will continue to work onInternet Explorer 10.
- No Confluence releases after 5.10.x will be tested with Internet Explorer 10.

Deprecated databases for Confluence (8 June 2016)

This section announces the end of Atlassian support for certain databases for Confluence. End of support means that Atlassian will not fix bugs related to the specified database past the support end date for your version of Confluence.

The details are below. Please refer to the list ofsupported platformsfor details of platform support for Confluence. If you have questions or concerns regarding this announcement, please email`eol-announcement at atlassian dot com`.

End of life announcement for database support

Database	Support End Date

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

MySQL 5.5	After Confluence 5.10.x

Notes:

- Confluence 5.10 is the last version that will support MySQL 5.5.
- Confluence 5.10 and previously-released versions will continue to work with the database version listed above, however we will not fix bugs affecting these databases after the end-of-life datefor your version of Confluence.
- No Confluence releases after 5.10.x will be tested with the database listed above.

Deprecated macros for Confluence (13 November 2015)

Update 22 January 2019
We know from your feedback that the existing View File macros provide important functionality that the newer file upload and preview experience does not. For this reason, we've decided to **reverse the decision to stop supporting these macros**.

This means from Confluence 6.14, we will fix bugs relating to these macros, and will not remove the macros from Confluence.

With the release of Confluence **5.9** we will be ending support for the following macros, known collectively as the 'View File' macros:

- Office Excel
- Office Word
- Office PowerPoint
- PDF

End of support means that Atlassian will not fix bugs related to these macros past the support end date for your version of Confluence. We plan to remove these macros in a future Confluence release, and will provide plenty of information to help you make the transition when the time comes.

The View File macros will still be available in future Confluence releases (including Confluence 5.9, 5.10 and later), but we recommend inserting Office and PDF files as a thumbnail or link, and using the preview to view the file in full, as it provides a much better way to display Office and PDF files on your pages. See Display Files and Images for more info.

If you have any questions or concerns, please comment on this issue
~~CONFSERVER-39829~~ - Plans to remove the view file macros CLOSED

Discontinued features for Confluence (10 July 2015)

Status updates

As part of our work to make Confluence simpler and easier to use we've decided to remove the Status Updates feature in**Confluence 5.9**. This includes the ability to:

- update your status
- see other people's status via their profile or the User Status List macro.

Our research tells us that this feature isn't widely used, and we believe thatHipChatgives your team much better ways to share their status.

We'll provide more information at the time of the Confluence 5.9 release. If you have questions or concerns, please comment on this issue ~~CONFSERVER-38253~~ - Plans to remove status updates CLOSED .

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

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 favorite documentation theme features, including the ability to add:

- a header and footer
- custom content to the sidebar.

These new additions to the default theme are available in Confluence 5.9. As these fields will continue to use wiki markup, you will be able to drop your existing wiki markup straight from the Documentation theme into the default theme.

To help you switch themes we've put together a FAQ and step-by-step guidewhich covers everything from how to turn on the default theme, find out which spaces are using the theme, and what to do if the Documentation theme is the global theme for your whole site.

If you have any questions or concerns please comment on this issue
~~CONFSERVER-38256~~ - Plans to remove the documentation theme CLOSED .

Deprecated databases for Confluence (19 May 2015)

This section announces the end of Atlassian support for certain databases for Confluence. End of support means that Atlassian will not fix bugs related to the specified database past the support end date for your version of Confluence.

The details are below. Please refer to the list of supported platforms for details of platform support for Confluence. If you have questions or concerns regarding this announcement, please email `eol-announcement at atlassian dot com`.

End of Life Announcement for Database Support

Database	Support End Date
Microsoft SQL 2008	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.

DeprecatedTomcatplatform for Confluence (1 May 2015)

This section announces the end of Atlassian support for Tomcat 7.0.x for Confluence. Aspreviously announced, we now only support the version of Tomcat that is bundled with your version of Confluence.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

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 affectingTomcat 7.0.x after the end-of-life date for your version of Confluence.
- Confluence 5.8 will not be tested with Tomcat 7.0.x.

Deprecated Web Browsers for Confluence (20 April 2015)

Atlassian will end support for Internet Explorer 9 in the next major release after Confluence 5.8.x.End of support means thatAtlassian will not fix bugs related to Internet Explorer 9 past the support end date,except for security related issues.

This change allows us to use more modern browser technologies to give you the best user experience in Confluence. Check out theSupported Platformspage for the full list of supported browsers.

If you have questions or concerns regarding this announcement, please email eol-announcement at atlassian dot com.

Internet Explorer 9 (IE9) End of Support Notes

- Confluence 5.8 will be the last major release that supportsInternet Explorer 9
- Confluence 5.8.x and earlier versions will continue to work onInternet Explorer 9
- No Confluence releases after 5.8.x will be tested with Internet Explorer 9

Deprecated Java platform for Confluence (27 January 2015)

This section announces the end of Atlassian support for Java 7 for Confluence. Please note that Oracle is planning to stop providing public updates for JRE 7 inApril 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.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

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 bethe version of Tomcat that is bundled with each release.

Confluence 5.6 will be the last Confluence release to provide an EAR/WAR edition.

Q: Do I need to use the installer?
No,the removal of the EAR/WAR distribution does not force you to use the installer. You can still use the standalone distribution, which doesn't have an install script - it's just a copy of Tomcat with Confluence configured inside it. Essentially it's a directory that you unpack and then run yourself.

Q: What if a security problem is found in the bundled version of Tomcat?
Our security team monitors vulnerabilities in all our dependencies, including Tomcat, and fixes continue to follow ourSecurity Bugfix Policy. If at any time you become aware of a vulnerability we've missed, please report it as described inHow to report a security issue.

If you have more questions or concerns regarding this announcement, please contact us ateol-announcement at atlassian dot com.

Deprecated databases for Confluence (12 June 2014)

This section announces the end of Atlassian support for certain databases for Confluence. End of support means that Atlassian will not fix bugs related to the specified database past the support end date for your version of Confluence.

The details are below. Please refer to the list of supported platforms for details of platform support for Confluence. If you have questions or concerns regarding this announcement, please email eol-announcement at atlassian dot com.

End of Life Announcement for Database Support

Database	Support End Date
PostgreSQL 8.4	With the release of Confluence 5.7
PostgreSQL 9.0	
PostgreSQL 9.1	
MySQL 5.1	

Notes:

- Confluence 5.6 is the last version that will support the database versions listed above.
- Confluence 5.6 and previously-released versions will continue to work with the database versions listed above, however we will not fix bugs affecting these databases after the end-of-life date for your version of Confluence.
- Confluence 5.7 has not been tested with the databases listed above.

DeprecatedTomcatplatform for Confluence (22 April 2014)

This section announces the end of Atlassian support for Tomcat 6.0.x for Confluence.

End of support means that Atlassian will not fix bugs related to the specified version of Tomcat, past the support end date for your version of Confluence. The details are below. Please refer to the list of supported platforms for details of platform support for Confluence. If you have questions or concerns regarding this announcement, please email eol-announcement at atlassian dot com.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

End of Life Announcement for Tomcat 6.0.x Support

Platform	Support End Date
Tomcat 6.0.x	When Confluence 5.6 is released, due in mid 2014

Tomcat 6.0.x notes:

- Confluence 5.5 is the last major version that will support Tomcat 6.0.x. The Confluence 5.5.x bug-fix releases will also continue to support Tomcat 6.0.x.
- Confluence 5.5.x and previously-released versions will continue to work with Tomcat 6.0.x. However, we will not fix bugs affectingTomcat 6.0.x after the end-of-life date for your version of Confluence.
- Confluence 5.6 will not be tested with Tomcat 6.0.x.

Deprecated Databases for Confluence (2 December 2013)

This section announces the end of Atlassian support for certain databases for Confluence. End of support means that Atlassian will not fix bugs related to the specified database past the support end date for your version of Confluence.

The details are below. Please refer to the list of supported platforms for details of platform support for Confluence. If you have questions or concerns regarding this announcement, please email `eol-announcement at atlassian dot com`.

End of Life Announcement for Database Support

Database	Support End Date
PostgreSQL 8.3	When Confluence 5.5 is released, due in early 2014

PostgreSQL 8.3 notes:

- Confluence 5.4 is the last version that will support PostgreSQL 8.3.
- Confluence 5.4 and previously-released versions will continue to work with PostgreSQL 8.3. However, we will not fix bugs affecting PostgreSQL 8.3 after the end-of-life date for your version of Confluence.
- Confluence 5.5 will not be tested with PostgreSQL 8.3.

Deprecated Web Browsers for Confluence (24 September 2013)

To allow us to dedicate resources to providing the best experience on modern browsers, Confluence 5.5 will be the**last release that supports Internet Explorer 8 (IE8)**. The reasons behind this decision are to enable us to provide the best user experience to our customers, accelerate our pace of innovation and give us the ability to utilize modern browser technologies.

End of support means that Atlassian will not perform any maintenance on Confluence related to IE8 after the final release of Confluence 5.5.x, except for security related issues. In order to minimize the impact on you and the way your company uses Confluence, we have provided this announcement as early as possible, and hope that the subsequent 6 month period will give you adequate time to prepare for this change without disruption.

Atlassian will continue to supportInternet Explorer 9 (IE9) and Internet Explorer 10 (IE10) as well asthe latest versions of Chrome, Firefox and Safari. For further information,please refer to theSupported Platformspage. If you have questions or concerns regarding this announcement, please email eol-announcement at atlassian dot com.

Deprecated Databases for Confluence (13 August 2013)

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

This section announces the end of Atlassian support for certain databases for Confluence. End of support means that Atlassian will not fix bugs related to the specified database past the support end date for your version of Confluence.

The details are below. Please refer to the list of supported platforms for details of platform support for Confluence. If you have questions or concerns regarding this announcement, please email `eol-announcement at atlassian dot com`.

End of Life Announcement for Database Support

Database	Support End Date
MS SQL 2005	When Confluence 5.3 is released, due in late 2013

MS SQL 2005 notes:

- Confluence 5.2 is the last version that will support MS SQL 2005.
- Confluence 5.2 and previously-released versions will continue to work with MS SQL 2005. However, we will not fix bugs affecting MS SQL 2005 after the end-of-life date for your version of Confluence.
- Confluence 5.3 will not be tested with MS SQL 2005.

Deprecated Tomcat platform for Confluence (29 August 2012)

This section announces the end of Atlassian support for Tomcat 5.5.x for Confluence. Please note: Apache have announced that support for Apache Tomcat 5.5.x will end on 30 September 2012:End of life for Apache Tomcat 5.5.x.

End of support means that Atlassian will not fix bugs related to the specified version of Tomcat, past the support end date for your version of Confluence. The details are below. Please refer to the list of supported platforms for details of platform support for Confluence. If you have questions or concerns regarding this announcement, please email `eol-announcement at atlassian dot com`.

End of Life Announcement for Tomcat 5.5.x Support

Platform	Support End Date
Tomcat 5.5.x	When Confluence 5.0 is released, due in early 2013

Tomcat 5.5.x notes:

- Confluence 4.3 is the last major version that will support Tomcat 5.5.x. The Confluence 4.3.x bug-fix releases will also continue to support Tomcat 5.5.x.
- Tomcat 6.0.x will still be supported in Confluence 5.0.
- Confluence 4.3.x and previously-released versions will continue to work with Tomcat 5.5.x. However, we will not fix bugs affectingTomcat 5.5.x after the end-of-life date for your version of Confluence.
- Confluence 5.0 will not be tested with Tomcat 5.5.x.

Deprecated Java platform for Confluence (6 August 2012)

This section announces the end of Atlassian support for Java 6 for Confluence. Please note that Oracle has announced the end of public updates for Java 6: Java SE 6 End of Public Updates Notice.

End of support means that Atlassian will not fix bugs related to the specified version of Java, past the support end date for your version of Confluence. The details are below. Please refer to the list of supported platforms for details of platform support for Confluence. If you have questions or concerns regarding this announcement, please email `eol-announcement at atlassian dot com`.

End of Life Announcement for Java 6 Support

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Platform	Support End Date
Java 6 (JRE and JDK 1.6)	When Confluence 5.0 is released, due in early 2013

Java 6 notes:

- Confluence 4.3 is the last major version that will support Java 6. The Confluence 4.3.x bug-fix releases will also continue to support Java 6.
- Java 7 (JRE and JDK 1.7) will still be supported in Confluence 5.0.
- Confluence 4.3.x and previously-released versions will continue to work with Java 6. However, we will not fix bugs affecting Java 6 after the end-of-life date for your version of Confluence.
- Confluence 5.0 will not be tested with Java 6.

Deprecated Databases for Confluence (1 May 2012)

This section announces the end of Atlassian support for certain databases for Confluence. End of support means that Atlassian will not fix bugs related to the specified database past the support end date for your version of Confluence.

The details are below. Please refer to the list of supported platforms for details of platform support for Confluence. If you have questions or concerns regarding this announcement, please email `eol-announcement at atlassian dot com`.

End of Life Announcement for Database Support

Database	Support End Date
PostgreSQL 8.2	When Confluence 4.3 is released, due in mid 2012

PostgreSQL 8.2 notes:

- Confluence 4.2 is the last version that will support version 8.2 of PostgreSQL.
- Versions 8.3, 8.4 and 9.0 will still be supported in Confluence 4.3.
- Confluence 4.2 and previously-released versions will continue to work with PostgreSQL 8.2. However, we will not fix bugs affecting PostgreSQL 8.2 after the end-of-life date for your version of Confluence.
- Confluence 4.3 will not be tested with PostgreSQL 8.2.

Deprecated Databases for Confluence (13 March 2012)

This section announces the end of Atlassian support for certain databases for Confluence. End of support means that Atlassian will not fix bugs related to the specified database past the support end date for your version of Confluence.

The details are below. Please refer to the list of supported platforms for details of platform support for Confluence. If you have questions or concerns regarding this announcement, please email `eol-announcement at atlassian dot com`.

End of Life Announcement for Database Support

Database	Support End Date
DB2	When Confluence 4.3 is released, due in mid 2012

DB2 notes:

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

- Confluence 4.2 is the last version that will support DB2.
- From Confluence 4.3, no versions of DB2 will be supported.
- Confluence 4.2 and previously-released versions will continue to work with DB2. However, we will not fix bugs affecting DB2 after the end-of-life date for your version of Confluence.
- Confluence 4.3 will not be tested with DB2.
- For help with moving from DB2 to a supported database, please refer to the list of supported databases and the guide to migrating to another database.

Deprecated Operating Systems for Confluence (21 July 2011)

This section announces the end of Atlassian support for certain operating systems for Confluence. End of support means that Atlassian will not fix bugs related to running Confluence server on that operating system past the support end date.

We will stop supporting the following operating systems from Confluence 4.0, due in late 2011:

- Mac OS X (as a Confluence server platform).

The details are below. Please refer to the list of supported platforms for details of platform support for Confluence. If you have questions or concerns regarding this announcement, please email `eol-announcement at atlassian dot com`.

End of Life Announcement for Operating System Support

Operating System	Support End Date
Mac OS X (as a Confluence server platform)	When Confluence 4.0 releases, due in late 2011

- **Mac OS X Notes:**
  - Atlassian intends to end support for Mac OS X (as a server platform) in Confluence 4.0 (due for release in late 2011). Confluence 3.5 is the last version that will support Mac OS X.
  - The Sun/Oracle JDK/JRE 1.6 is the only JDK platform officially supported by Atlassian. This means that Apple Mac OS X is not a supported operating system for the Confluence server, as the Sun/Oracle JDK does not run on Mac OS X.
  - Accessing Confluence as a user from Mac OS X via a compatible web browser will still be supported for the forseeable future.

Deprecated Databases for Confluence (7 January 2011)

This section announces the end of Atlassian support for certain database versions for Confluence. End of support means that Atlassian will not fix bugs related to certain database versions past the support end date.

We will **stop supporting the following database versions** from Confluence 4.0, due in late 2011:

- MySQL 5.0.

The details are below. Please refer to the list of supported platforms for details of platform support for Confluence. If you have questions or concerns regarding this announcement, please email `eol-announcement at atlassian dot com`.

End of Life Announcement for Database Support

Database	Support End Date
MySQL (version 5.0 only)	When Confluence 4.0 releases, due in late 2011

- **MySQL Notes:**

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

- 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)

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

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)

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

This section announces the end of Atlassian support for certain database versions for Confluence. End of support means that Atlassian will not fix bugs related to certain database versions past the support end date.

We will **stop supporting the following database versions**:

- From Confluence 3.4, due in the second half of 2010, Confluence will no longer support Oracle 10g (i.e. Oracle 10.1 and Oracle 10.2).
  Note, Oracle 11g (i.e. Oracle 11.1 and Oracle 11.2) will still be supported.

We have made these decisions in line with Oracle's decision to stop support for Oracle 10g, as per the "Oracle Database (RDBMS) Releases Support Status Summary [ID 161818.1]" article on the Oracle Support site (note, you will need an Oracle Support account to find and view the article). This also will reduce the testing time required for each release and help us speed up our ability to deliver market-driven features. We are committed to helping our customers understand this decision and assist them in upgrading to Oracle 11g if needed.

The details are below. Please refer to the Supported Platforms for more details regarding platform support for Confluence. If you have questions or concerns regarding this announcement, please email `eol-announcement at atlassian dot com`.

End of Life Announcement for Database Support

Database	Support End Date
Oracle (version 10.1 and 10.2 only)	When Confluence 3.4 releases, due in the second half of 2010

- **Oracle (version 10.1 and 10.2 only) End of Support Notes:**
  - Atlassian intends to end support for Oracle 10.1 and Oracle 10.2 in Confluence 3.4 (due to release in the second half of 2010), with the final support for these platforms in Confluence 3.3 . Oracle 11.1 and Oracle 11.2 will still be supported.
  - 'Support End Date' means that Confluence 3.3 and previous released versions will continue to work with the Oracle 10.1 and Oracle 10.2. However, we will not fix bugs affecting Oracle 10.1 or Oracle 10.2 past the support end date.
  - Confluence 3.4 (due to release in the second half of 2010) will not be tested with Oracle 10.1 and Oracle 10.2.

Deprecated Web Browsers for Confluence (6 July 2010)

This section announces the end of Atlassian support for certain web browser versions for Confluence. End of support means that Atlassian will not fix bugs related to certain web browser versions past the support end date.

We will **stop supporting the following web browser versions**:

- From Confluence 3.4, due in the second half of 2010, Confluence will no longer support Safari 3 or Safari 3.1.
  Note, Safari 4 will still be supported.

The details are below. Please refer to the Supported Platforms for more details regarding platform support for Confluence. If you have questions or concerns regarding this announcement, please email `eol-announcement at atlassian dot com`.

End of Life Announcement for Web Browser Support

Web Browser	Support End Date
Safari (version 3 and 3.1 only)	When Confluence 3.4 releases, due in the second half of 2010

- **Safari (version 3 and 3.1 only) End of Support Notes:**
  - Atlassian intends to end support for Safari 3 and Safari 3.1 in Confluence 3.4 (due to release in the second half of 2010), with the final support for these platforms in Confluence 3.3. Safari 4 will still be supported.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

- 'Support End Date' means that Confluence 3.3 and previous released versions will continue to work with the Safari 3 and Safari 3.1. However, we will not fix bugs affecting Safari 3 and Safari 3.1 past the support end date.
- Confluence 3.4 (due to release in the second half of 2010) will not be tested with Safari 3 and Safari 3.1.

Deprecated Databases for Confluence (24 March 2010)

This section announces the end of Atlassian support for certain database versions for Confluence. End of support means that Atlassian will not fix bugs related to certain database versions past the support end date.

We will **stop supporting the following database versions**:

- From Confluence 3.3, due in Q3 2010, Confluence will no longer support DB2 8.2.
  Note, DB2 9.7 will still be supported.

We are reducing our database support to reduce the amount of testing time and help us speed up our ability to deliver market-driven features. We are committed to helping our customers understand this decision and assist them in upgrading to DB2 9.7 if needed.

The details are below. Please refer to the Supported Platforms for more details regarding platform support for Confluence. If you have questions or concerns regarding this announcement, please email `eol-announcement at atlassian dot com`.

End of Life Announcement for Database Support

Database	Support End Date
DB2 (version 8.2 only)	When Confluence 3.3 releases, due Q3 2010

- **DB2 (version 8.2 only) End of Support Notes:**
  - Atlassian intends to end support for DB2 8.2 in Q3 2010, with the final support for these platforms in Confluence 3.2. DB2 9.7 will still be supported.
  - 'Support End Date' means that Confluence 3.2 and previous released versions will continue to work with the DB2 8.2. However, we will not fix bugs affecting DB2 8.2 past the support end date.
  - Confluence 3.3 (due to release in Q3 2010) will not be tested with DB2 8.2.

Deprecated Application Servers for Confluence (27 January 2010)

This section announces the end of Atlassian support for certain application servers for Confluence. End of support means that Atlassian will not fix bugs related to certain application servers past the support end date.

We will **stop supporting the following application servers**:

- From Confluence 3.2, due late Q1 2010, Confluence will no longer support JBoss application servers.
- From Confluence 3.3, due in Q3 2010, Confluence will no longer support Oracle WebLogic, IBM WebSphere or Caucho Resin.

We are reducing our application server platform support to reduce the amount of testing time and help us speed up our ability to deliver market-driven features. We are committed to helping our customers understand this decision and assist them in migrating to Tomcat, our supported application server.

The details are below. Please refer to the Supported Platforms for more details regarding platform support for Confluence. If you have questions or concerns regarding this announcement, please email `eol-announcement at atlassian dot com`.

End of Life Announcement for Application Server Support

Application Servers	Support End Date
JBoss 4.2.2	When Confluence 3.2 releases, due late Q1 2010

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Oracle WebLogic 9.2	When Confluence 3.3 releases, due Q3 2010
IBM WebSphere 6.1	When Confluence 3.3 releases, due Q3 2010
Caucho Resin 3.0, 3.1.6, 3.1.7	When Confluence 3.3 releases, due Q3 2010

- **JBoss End of Support Notes:**
  - 'Support End Date' means that Confluence 3.1 and previous released versions will continue to work with stated application servers. However, we will not fix bugs affecting JBoss application servers.
  - Confluence 3.2 will not support JBoss application servers.

- **WebLogic, WebSphere and Resin End of Support Notes:**
  - Atlassian intends to end support for Oracle WebLogic, IBM WebSphere, and Caucho Resin in Q3 2010, with the final support for these platforms in Confluence 3.2.
  - 'Support End Date' means that Confluence 3.2 and previous released versions will continue to work with the stated application servers. However, we will not fix bugs affecting Oracle WebLogic, IBM WebSphere, and Caucho Resin application servers past the support end date.
  - Confluence 3.3 (due to release in Q3 2010) will only be tested with and support Tomcat 5.5.20+ and 6.0.
  - If you have concerns with this end of support announcement, please email `eol-announcement at atlassian dot com`.

Why is Atlassian doing this?

We have chosen to standardize on Tomcat, because it is the most widely used application server in our user population. It is fast, robust, secure, well-documented, easy to operate, open source, and has a huge community driving improvements. It is the de facto industry standard, with several companies available that specialize in providing enterprise grade support contracts for it, ranging from customizations to 24/7 support.

Deprecated Java Platforms for Confluence (27 January 2010)

This section announces the end of Atlassian support for certain Java Platforms for Confluence.

We will **stop supporting the following Java Platforms**:

- From Confluence 3.3, due Q3 2010, support for Java Platform 5 (JDK/JRE 1.5) will end.

We are ending support for Java Platform 5, in line with the Java SE Support Roadmap (i.e. "End of Service Life" for Java Platform 5 dated October 30, 2009). We are committed to helping our customers understand this decision and assist them in updating to Java Platform 6, our supported Java Platform.

The details are below. Please refer to the Supported Platforms for more details regarding platform support for Confluence. If you have questions or concerns regarding this announcement, please email `eol-announcement at atlassian dot com`.

End of Life Announcement for Java Platform Support

Java Platform	Support End Date
Java Platform 5 (JDK/JRE 1.5)	When Confluence 3.3 releases, due Q3 2010

- **Java Platform 5 End of Support Notes:**
  - Atlassian intends to end support for Java Platform 5 in Q3 2010.
  - 'Support End Date' means that Confluence 3.2.x and previous released versions will continue to work with Java Platform 5 (JDK/JRE 1.5), however we will not fix bugs related to Java Platform 5 past the support end date.
  - Confluence 3.3 will only be tested with and support Java Platform 6 (JDK/JRE 1.6).
  - If you have concerns with this end of support announcement, please email `eol-announcement at atlassian dot com`.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

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.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Bundled Tomcat and Java versions

This page lists the specific versions of Apache Tomcat and Adopt OpenJDK or Eclipse Temurin OpenJDK that we bundle with Confluence. This information is useful if you want to check whether your Confluence version might be using a Tomcat or Java version that's affected by a specific issue, vulnerability, or bug.

We also list the specific Java versions we use when testing Confluence, which can be handy if you don't run Confluence with the bundled JRE.

Confluence 8.4

Confluence	Tomcat	Bundled JRE	Tested JDKs
8.4.0	9.0.73	Eclipse Temurin 17.0.7_7	Eclipse Temurin 11.0.19_7 Eclipse Temurin 17.0.7_7

Confluence 8.3

Confluence	Tomcat	Bundled JRE	Tested JDKs
8.3.2	9.0.73	Eclipse Temurin 17.0.7_7	Eclipse Temurin 11.0.19_7 Eclipse Temurin 17.0.7_7
8.3.1	9.0.73	Eclipse Temurin 17.0.6_10	Eclipse Temurin 11.0.19_7 Eclipse Temurin 17.0.7_7
8.3.0	9.0.73	Eclipse Temurin 17.0.6_10	Eclipse Temurin 11.0.18_10 Eclipse Temurin 17.0.6_10

Confluence 8.2

Confluence	Tomcat	Bundled JRE	Tested JDKs
8.2.0	9.0.71	Eclipse Temurin 11.0.17_8	Eclipse Temurin 11.0.18_10 Eclipse Temurin 17.0.6_10
8.2.1	9.0.71	Eclipse Temurin 11.0.17_8	Eclipse Temurin 11.0.18_10 Eclipse Temurin 17.0.6_10
8.2.2	9.0.73	Eclipse Temurin 11.0.17_8	Eclipse Temurin 11.0.18_10 Eclipse Temurin 17.0.6_10
8.2.3	9.0.73	Eclipse Temurin 11.0.17_8	Eclipse Temurin 11.0.18_10 Eclipse Temurin 17.0.6_10

Confluence 8.1

Confluence	Tomcat	Bundled JRE	Tested JDKs
8.1.0	9.0.65	Eclipse Temurin 11.0.16.1_1	Eclipse Temurin 11.0.17_8 Eclipse Temurin 17.0.4_8
8.1.1	9.0.65	Eclipse Temurin 11.0.17_8	Eclipse Temurin 11.0.17_8 Eclipse Temurin 17.0.4_8
8.1.3	9.0.65	Eclipse Temurin 11.0.17_8	Eclipse Temurin 11.0.18_10 Eclipse Temurin 17.0.6_10
8.1.4	9.0.65	Eclipse Temurin 11.0.17_8	Eclipse Temurin 11.0.18_10 Eclipse Temurin 17.0.6_10

Note: Confluence 8.1.2 was an internal release.

Confluence 8.0

Confluence	Tomcat	Bundled JRE	Tested JDKs
8.0.0	9.0.65	Eclipse Temurin 11.0.16.1_1	Eclipse Temurin 11.0.16.1_1 Eclipse Temurin 17.0.4_8
8.0.1	9.0.65	Eclipse Temurin 11.0.16.1_1	Eclipse Temurin 11.0.16.1_1 Eclipse Temurin 17.0.4_8
8.0.2	9.0.65	Eclipse Temurin 11.0.16.1_1	Eclipse Temurin 11.0.16.1_1 Eclipse Temurin 17.0.4_8
8.0.3	9.0.65	Eclipse Temurin 11.0.16.1_1	Eclipse Temurin 11.0.17_8 Eclipse Temurin 17.0.4_8
8.0.4	9.0.65	Eclipse Temurin 11.0.16.1_1	Eclipse Temurin 11.0.17_8 Eclipse Temurin 17.0.4_8

Confluence 7.20

Confluence	Tomcat	Bundled JRE	Tested JDKs
7.20.0	9.0.65	Eclipse Temurin 11.0.16.1_1	Eclipse Temurin 11.0.16.1_1 Eclipse Temurin 8u322b06
7.20.1	9.0.65	Eclipse Temurin 11.0.16.1_1	Eclipse Temurin 11.0.16.1_1 Eclipse Temurin 8u322b06
7.20.2	9.0.65	Eclipse Temurin 11.0.16.1_1	Eclipse Temurin 11.0.16.1_1 Eclipse Temurin 8u322b06
7.20.3	9.0.65	Eclipse Temurin 11.0.16.1_1	Eclipse Temurin 11.0.16.1_1 Eclipse Temurin 8u322b06

Confluence 7.19

Confluence	Tomcat	Bundled JRE	Tested JDKs
7.19.0	9.0.63	Eclipse Temurin 11.0.14.1_1	Oracle JDK8u321 Oracle JDK11.0.14 Eclipse Temurin 8u322b06
7.19.1	9.0.63	Eclipse Temurin 11.0.14.1_1	Oracle JDK8u321 Oracle JDK11.0.14 Eclipse Temurin 8u322b06
7.19.2	9.0.65	Eclipse Temurin 11.0.14.1_1	Oracle JDK8u321 Oracle JDK11.0.14 Eclipse Temurin 8u322b06
7.19.3	9.0.65	Eclipse Temurin 11.0.16.1_1	Oracle JDK8u321 Oracle JDK 11.0.16 Eclipse Temurin 8u322b06
7.19.4	9.0.65	Eclipse Temurin 11.0.16.1_1	Oracle JDK8u321 Oracle JDK 11.0.16 Eclipse Temurin 8u322b06
7.19.5	9.0.65	Eclipse Temurin 11.0.16.1_1	Oracle JDK8u321 Oracle JDK 11.0.14 Eclipse Temurin 8u322b06

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

7.19.6	9.0.65	Eclipse Temurin 11.0.16.1_1	Oracle JDK8u321 Oracle JDK 11.0.14 Eclipse Temurin 8u322b06
7.19.7	9.0.71	Eclipse Temurin 11.0.17_8	Oracle JDK8u321 Oracle JDK 11.0.14 Eclipse Temurin 8u362b09
7.19.8	9.0.73	Eclipse Temurin 11.0.17_8	Oracle JDK8u321 Oracle JDK 11.0.14 Oracle JDK 17.0.6 Eclipse Temurin 8u362b09 Eclipse Temurin 11.0.18_10 Eclipse Temurin 17.0.6_10
7.19.9	9.0.73	Eclipse Temurin 11.0.17_8	Oracle JDK8u321 Oracle JDK 11.0.14 Oracle JDK 17.0.6 Eclipse Temurin 8u362b09 Eclipse Temurin 11.0.18_10 Eclipse Temurin 17.0.6_10
7.19.10	9.0.73	Eclipse Temurin 11.0.17_8	Oracle JDK8u321 Oracle JDK 11.0.14 Oracle JDK 17.0.6 Eclipse Temurin 8u362b09 Eclipse Temurin 11.0.19_7 Eclipse Temurin 17.0.7_7
7.19.11	9.0.76	Eclipse Temurin 11.0.19_7	Oracle JDK8u321 Oracle JDK 11.0.14 Oracle JDK 17.0.6 Eclipse Temurin 8u362b09 Eclipse Temurin 11.0.19_7 Eclipse Temurin 17.0.7_7

Confluence 7.18

Confluence	Tomcat	Bundled JRE	Tested JDKs
7.18.0	9.0.58	Eclipse Temurin 11.0.14.1_1	Oracle JDK8u321 Oracle JDK11.0.14 Eclipse Temurin 8u322b06
7.18.1	9.0.63	Eclipse Temurin 11.0.14.1_1	Oracle JDK8u321 Oracle JDK11.0.14 Eclipse Temurin 8u322b06
7.18.2	9.0.63	Eclipse Temurin 11.0.14.1_1	Oracle JDK8u321 Oracle JDK11.0.14 Eclipse Temurin 8u322b06
7.18.3	9.0.63	Eclipse Temurin 11.0.14.1_1	Oracle JDK8u321 Oracle JDK11.0.14 Eclipse Temurin 8u322b06

Confluence 7.17

Confluence	Tomcat	Bundled JRE	Tested JDKs
7.17.0	9.0.58	Eclipse Temurin 11.0.12_7	Oracle JDK8u321 Oracle JDK11.0.14 Eclipse Temurin 8u322b06

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

7.17.1	9.0.58	Eclipse Temurin 11.0.12_7	Oracle JDK8u321 Oracle JDK11.0.14 Eclipse Temurin 8u322b06
7.17.2	9.0.58	Eclipse Temurin 11.0.14.1_1	Oracle JDK8u321 Oracle JDK11.0.14 Eclipse Temurin 8u322b06
7.17.3	9.0.58	Eclipse Temurin 11.0.14.1_1	Oracle JDK8u321 Oracle JDK11.0.14 Eclipse Temurin 8u322b06
7.17.4	9.0.63	Eclipse Temurin 11.0.14.1_1	Oracle JDK8u321 Oracle JDK11.0.14 Eclipse Temurin 8u322b06
7.17.5	9.0.63	Eclipse Temurin 11.0.14.1_1	Oracle JDK8u321 Oracle JDK11.0.14 Eclipse Temurin 8u322b06

Confluence 7.16

Confluence	Tomcat	Bundled JRE	Tested JDKs
7.16.0	9.0.54	Eclipse Temurin 11.0.12_7	Oracle JDK 8u301 Oracle JDK 11.0.12 Eclipse Temurin 8u302b08
7.16.1	9.0.54	Eclipse Temurin 11.0.12_7	Oracle JDK8u321 Oracle JDK11.0.14 Eclipse Temurin 8u322b06
7.16.2	9.0.54	Eclipse Temurin 11.0.12_7	Oracle JDK8u321 Oracle JDK11.0.14 Eclipse Temurin 8u322b06
7.16.3	9.0.58	Eclipse Temurin 11.0.12_7	Oracle JDK8u321 Oracle JDK11.0.14 Eclipse Temurin 8u322b06
7.16.4	9.0.63	Eclipse Temurin 11.0.14.1_1	Oracle JDK8u321 Oracle JDK11.0.14 Eclipse Temurin 8u322b06
7.16.5	9.0.63	Eclipse Temurin 11.0.14.1_1	Oracle JDK8u321 Oracle JDK11.0.14 Eclipse Temurin 8u322b06

Confluence 7.15

Confluence	Tomcat	Bundled JRE	Tested JDKs
7.15.0	9.0.54	Eclipse Temurin 11.0.12_7	Oracle JDK 8u301 Oracle JDK 11.0.12 Eclipse Temurin 8u302b08
7.15.1	9.0.54	Eclipse Temurin 11.0.12_7	Oracle JDK 8u301 Oracle JDK 11.0.12 Eclipse Temurin 8u302b08
7.15.2	9.0.63	Eclipse Temurin 11.0.14.1_1	Oracle JDK8u321 Oracle JDK11.0.14 Eclipse Temurin 8u322b06

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

7.15.3	9.0.63	Eclipse Temurin 11.0.14.1_1	Oracle JDK8u321 Oracle JDK11.0.14 Eclipse Temurin 8u322b06

Confluence 7.14

Confluence	Tomcat	Bundled JRE	Tested JDKs
7.14.0	9.0.45	Adopt OpenJDK 11.0.8_10	Oracle JDK 8u301 Oracle JDK 11.0.12 Eclipse Temurin 8u302b08
7.14.1	9.0.54	Eclipse Temurin 11.0.12_7	Oracle JDK 8u301 Oracle JDK 11.0.12 Eclipse Temurin 8u302b08
7.14.2	9.0.54	Eclipse Temurin 11.0.12_7	Oracle JDK 8u301 Oracle JDK 11.0.12 Eclipse Temurin 8u302b08
7.14.3	9.0.54	Eclipse Temurin 11.0.12_7	Oracle JDK 8u301 Oracle JDK 11.0.12 Eclipse Temurin 8u302b08
7.14.4	9.0.63	Eclipse Temurin 11.0.14.1_1	Oracle JDK8u321 Oracle JDK11.0.14 Eclipse Temurin 8u322b06

Confluence 7.13

Confluence	Tomcat	Bundled JRE	Tested JDKs
7.13.0	9.0.45	Adopt OpenJDK 11.0.8_10	Oracle JDK 8u291 Oracle JDK 11.0.11 Eclipse Temurin 8u292b10
7.13.1	9.0.45	Adopt OpenJDK11.0.12_7	Oracle JDK 8u301 Oracle JDK 11.0.12 Eclipse Temurin 8u302b08
7.13.2	9.0.45	Eclipse Temurin 11.0.12_7	Oracle JDK 8u301 Oracle JDK 11.0.12 Eclipse Temurin 8u302b08
7.13.3	9.0.45	Eclipse Temurin 11.0.12_7	Oracle JDK 8u301 Oracle JDK 11.0.12 Eclipse Temurin 8u302b08
7.13.4	9.0.45	Eclipse Temurin 11.0.12_7	Oracle JDK 8u301 Oracle JDK 11.0.12 Eclipse Temurin 8u302b08
7.13.5	9.0.58	Eclipse Temurin 11.0.12_7	Oracle JDK 8u301 Oracle JDK 11.0.12 Eclipse Temurin 8u302b08
7.13.6	9.0.58	Eclipse Temurin 11.0.14.1_1	Oracle JDK8u321 Oracle JDK11.0.14 Eclipse Temurin 8u322b06
7.13.7	9.0.63	Eclipse Temurin 11.0.14.1_1	Oracle JDK8u321 Oracle JDK11.0.14 Eclipse Temurin 8u322b06

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

7.13.8	9.0.63	Eclipse Temurin 11.0.14.1_1	Oracle JDK8u321 Oracle JDK11.0.14 Eclipse Temurin 8u322b06
7.13.9	9.0.63	Eclipse Temurin 11.0.14.1_1	Oracle JDK8u321 Oracle JDK11.0.14 Eclipse Temurin 8u322b06
7.13.10	9.0.65	Eclipse Temurin 11.0.16.1_1	Oracle JDK 8u321 Oracle JDK 11.0.14 Eclipse Temurin 8u322b06
7.13.11	9.0.65	Eclipse Temurin 11.0.16.1_1	Oracle JDK 8u321 Oracle JDK 11.0.14 Eclipse Temurin 8u322b06
7.13.12	9.0.65	Eclipse Temurin 11.0.16.1_1	Oracle JDK 8u321 Oracle JDK 11.0.14 Eclipse Temurin 8u322b06
7.13.13	9.0.65	Eclipse Temurin 11.0.16.1_1	Oracle JDK 8u321 Oracle JDK 11.0.14 Eclipse Temurin 8u322b06
7.13.14	9.0.65	Eclipse Temurin 11.0.16.1_1	Oracle JDK8u321 Oracle JDK 11.0.14 Eclipse Temurin 8u322b06
7.13.15	9.0.71	Eclipse Temurin 11.0.17_8	Oracle JDK8u321 Oracle JDK 11.0.14 Eclipse Temurin 8u362b09
7.13.16	9.0.73	Eclipse Temurin 11.0.17_8	Oracle JDK8u321 Oracle JDK 11.0.14 Oracle JDK 17.0.6 Eclipse Temurin 8u362b09 Eclipse Temurin 11.0.18_10 Eclipse Temurin 17.0.6_10
7.13.17	9.0.73	Eclipse Temurin 11.0.17_8	Oracle JDK8u321 Oracle JDK 11.0.14 Oracle JDK 17.0.6 Eclipse Temurin 8u362b09 Eclipse Temurin 11.0.18_10 Eclipse Temurin 17.0.6_10
7.13.18	9.0.73	Eclipse Temurin 11.0.17_8	Oracle JDK8u321 Oracle JDK 11.0.14 Oracle JDK 17.0.6 Eclipse Temurin 8u362b09 Eclipse Temurin 11.0.19_7 Eclipse Temurin 17.0.7_7

Confluence 7.12

Confluence	Tomcat	Bundled JRE	Tested JDKs
7.12.0	9.0.40	Adopt OpenJDK 11.0.8_10	Oracle JDK 8u261 Oracle JDK 11.0.8 Adopt OpenJDK 8u265-b01
7.12.1	9.0.40	Adopt OpenJDK 11.0.8_10	Oracle JDK 8u261 Oracle JDK 11.0.8 Adopt OpenJDK 8u265-b01

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

7.12.2	9.0.45	Adopt OpenJDK 11.0.8_10	Oracle JDK 8u261 Oracle JDK 11.0.8 Adopt OpenJDK 8u265-b01
7.12.3	9.0.45	Adopt OpenJDK 11.0.8_10	Oracle JDK 8u261 Oracle JDK 11.0.8 Adopt OpenJDK 8u265-b01
7.12.4	9.0.45	Adopt OpenJDK 11.0.8_10	Oracle JDK 8u291 Oracle JDK 11.0.11 Adopt OpenJDK8u292b10
7.12.5	9.0.45	Adopt OpenJDK 11.0.8_10	Oracle JDK 8u291 Oracle JDK 11.0.11 Eclipse Temurin 8u292b10

Confluence 7.11

Confluence	Tomcat	Bundled JRE	Tested JDKs
7.11.0	9.0.40	Adopt OpenJDK 11.0.8_10	Oracle JDK 8u261 Oracle JDK 11.0.8 Adopt OpenJDK 8u265-b01
7.11.1	9.0.40	Adopt OpenJDK 11.0.8_10	Oracle JDK 8u261 Oracle JDK 11.0.8 Adopt OpenJDK 8u265-b01
7.11.2	9.0.40	Adopt OpenJDK 11.0.8_10	Oracle JDK 8u261 Oracle JDK 11.0.8 Adopt OpenJDK 8u265-b01
7.11.3	9.0.40	Adopt OpenJDK 11.0.8_10	Oracle JDK 8u261 Oracle JDK 11.0.8 Adopt OpenJDK 8u265-b01
7.11.6	9.0.45	Adopt OpenJDK 11.0.8_10	Oracle JDK 8u291 Oracle JDK 11.0.11 Eclipse Temurin 8u292b10

Confluence 7.10

Confluence	Tomcat	Bundled JRE	Tested JDKs
7.10.0	9.0.40	Adopt OpenJDK 11.0.8_10	Oracle JDK 8u261 Oracle JDK 11.0.8 Adopt OpenJDK 8u265-b01
7.10.1	9.0.40	Adopt OpenJDK 11.0.8_10	Oracle JDK 8u261 Oracle JDK 11.0.8 Adopt OpenJDK 8u265-b01
7.10.2	9.0.40	Adopt OpenJDK 11.0.8_10	Oracle JDK 8u261 Oracle JDK 11.0.8 Adopt OpenJDK 8u265-b01

Confluence 7.9

Confluence	Tomcat	Bundled JRE	Tested JDKs
7.9.0	9.0.37	Adopt OpenJDK 11.0.7_10	Oracle JDK 8u261 Oracle JDK 11.0.8 Adopt OpenJDK 8u265-b01

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

7.9.1	9.0.37	Adopt OpenJDK 11.0.7_10	Oracle JDK 8u261 Oracle JDK 11.0.8 Adopt OpenJDK 8u265-b01
7.9.2	9.0.37	Adopt OpenJDK 11.0.8_10	Oracle JDK 8u261 Oracle JDK 11.0.8 Adopt OpenJDK 8u265-b01

Confluence 7.8

Confluence	Tomcat	Bundled JRE	Tested JDKs
7.8.0	9.0.37	Adopt OpenJDK 11.0.7_10	Oracle JDK 8u251 Oracle JDK 11.0.8 Adopt OpenJDK 8u252-b09
7.8.1	9.0.37	Adopt OpenJDK 11.0.7_10	Oracle JDK 8u251 Oracle JDK 11.0.8 Adopt OpenJDK 8u252-b09
7.8.3	9.0.37	Adopt OpenJDK 11.0.7_10	Oracle JDK 8u261 Oracle JDK 11.0.8 Adopt OpenJDK 8u265-b01

Confluence 7.7

Confluence	Tomcat	Bundled JRE	Tested JDKs
7.7.2	9.0.33	Adopt OpenJDK 11.0.7_10	Oracle JDK 8u251 Oracle JDK 11.0.7 Adopt OpenJDK 8u232-b09
7.7.3	9.0.33	Adopt OpenJDK 11.0.7_10	Oracle JDK 8u251 Oracle JDK 11.0.8 Adopt OpenJDK 8u252-b09
7.7.4	9.0.37	Adopt OpenJDK 11.0.7_10	Oracle JDK 8u251 Oracle JDK 11.0.8 Adopt OpenJDK 8u252-b09

Note: 7.7.0 and 7.7.1 were internal releases.

Confluence 7.6

Confluence	Tomcat	Bundled JRE	Tested JDKs
7.6.0	9.0.33	Adopt OpenJDK 11.0.7_10	Oracle JDK 8u251 Oracle JDK 11.0.7 Adopt OpenJDK 8u232-b09
7.6.1	9.0.33	Adopt OpenJDK 11.0.7_10	Oracle JDK 8u251 Oracle JDK 11.0.7 Adopt OpenJDK 8u232-b09
7.6.2	9.0.33	Adopt OpenJDK 11.0.7_10	Oracle JDK 8u251 Oracle JDK 11.0.7 Adopt OpenJDK 8u232-b09

Confluence 7.5

Confluence	Tomcat	Bundled JRE	Tested JDKs

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

7.5.0	9.0.33	Adopt OpenJDK 11.0.5_10	Oracle JDK 8u251 Oracle JDK 11.0.7 Adopt OpenJDK 8u232-b09
7.5.1	9.0.33	Adopt OpenJDK 11.0.7_10	Oracle JDK 8u251 Oracle JDK 11.0.7 Adopt OpenJDK 8u232-b09
7.5.2	9.0.33	Adopt OpenJDK 11.0.7_10	Oracle JDK 8u251 Oracle JDK 11.0.7 Adopt OpenJDK 8u232-b09

Confluence 7.4

Confluence	Tomcat	Bundled JRE	Tested JDKs
7.4.0	9.0.33	Adopt OpenJDK 11.0.5_10	Oracle JDK 8u221 Oracle JDK 11.0.5 Adopt OpenJDK 8u232-b09
7.4.1	9.0.33	Adopt OpenJDK 11.0.5_10	Oracle JDK 8u251 Oracle JDK 11.0.7 Adopt OpenJDK 8u232-b09
7.4.2	9.0.33	Adopt OpenJDK 11.0.7_10	Oracle JDK 8u251 Oracle JDK 11.0.7 Adopt OpenJDK 8u232-b09
7.4.3	9.0.33	Adopt OpenJDK 11.0.7_10	Oracle JDK 8u251 Oracle JDK 11.0.7 Adopt OpenJDK 8u232-b09
7.4.4	9.0.33	Adopt OpenJDK 11.0.7_10	Oracle JDK 8u251 Oracle JDK 11.0.8 Adopt OpenJDK 8u252-b09
7.4.5	9.0.33	Adopt OpenJDK 11.0.7_10	Oracle JDK 8u261 Oracle JDK 11.0.8 Adopt OpenJDK 8u265-b01
7.4.6	9.0.33	Adopt OpenJDK 11.0.7_10	Oracle JDK 8u261 Oracle JDK 11.0.8 Adopt OpenJDK 8u265-b01
7.4.7	9.0.40	Adopt OpenJDK 11.0.8_10	Oracle JDK 8u261 Oracle JDK 11.0.8 Adopt OpenJDK 8u265-b01
7.4.8	9.0.40	Adopt OpenJDK 11.0.8_10	Oracle JDK 8u261 Oracle JDK 11.0.8 Adopt OpenJDK 8u265-b01
7.4.9	9.0.40	Adopt OpenJDK 11.0.8_10	Oracle JDK 8u261 Oracle JDK 11.0.8 Adopt OpenJDK 8u265-b01
7.4.10	9.0.45	Adopt OpenJDK 11.0.8_10	Oracle JDK 8u291 Oracle JDK 11.0.11 Adopt OpenJDK8u292b10
7.4.11	9.0.45	Adopt OpenJDK 11.0.8_10	Oracle JDK 8u291 Oracle JDK 11.0.11 Eclipse Temurin 8u292b10

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

7.4.12	9.0.45	Eclipse Temurin 11.0.12_7	Oracle JDK 8u301 Oracle JDK 11.0.12 Eclipse Temurin 8u302b08
7.4.13	9.0.45	Eclipse Temurin 11.0.12_7	Oracle JDK 8u301 Oracle JDK 11.0.12 Eclipse Temurin 8u302b08
7.4.14	9.0.45	Eclipse Temurin 11.0.12_7	Oracle JDK 8u301 Oracle JDK 11.0.12 Eclipse Temurin 8u302b08
7.4.15	9.0.45	Eclipse Temurin 11.0.12_7	Oracle JDK8u321 Oracle JDK11.0.14 Eclipse Temurin 8u322b06
7.4.16	9.0.58	Eclipse Temurin 11.0.12_7	Oracle JDK8u321 Oracle JDK11.0.14 Eclipse Temurin 8u322b06
7.4.17	9.0.63	Eclipse Temurin 11.0.14.1_1	Oracle JDK8u321 Oracle JDK11.0.14 Eclipse Temurin 8u322b06
7.4.18	9.0.63	Eclipse Temurin 11.0.14.1_1	Oracle JDK8u321 Oracle JDK11.0.14 Eclipse Temurin 8u322b06

Note: Adopt OpenJDK 11.0.3 and 11.0.4 had known issues in Linux and Windows in 7.4.0.

Confluence 7.3

Confluence	Tomcat	Bundled JRE	Tested JDKs
7.3.1	9.0.27	Adopt OpenJDK 11.0.5_10	Oracle JDK 8u221 Oracle JDK 11.0.5 Adopt OpenJDK 8u232-b09
7.3.2	9.0.27	Adopt OpenJDK 11.0.5_10	Oracle JDK 8u221 Oracle JDK 11.0.5 Adopt OpenJDK 8u232-b09
7.3.3	9.0.27	Adopt OpenJDK 11.0.5_10	Oracle JDK 8u221 Oracle JDK 11.0.5 Adopt OpenJDK 8u232-b09
7.3.4	9.0.33	Adopt OpenJDK 11.0.5_10	Oracle JDK 8u221 Oracle JDK 11.0.5 Adopt OpenJDK 8u232-b09
7.3.5	9.0.33	Adopt OpenJDK 11.0.5_10	Oracle JDK 8u221 Oracle JDK 11.0.5 Adopt OpenJDK 8u232-b09

Confluence 7.2

Confluence	Tomcat	Bundled JRE	Tested JDKs
7.2.0	9.0.27	Adopt OpenJDK 8u202b08	Oracle JDK 8u202 Oracle JDK 11.0.1 Adopt OpenJDK11.0.1+13

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

7.2.1	9.0.27	Adopt OpenJDK 8u202b08	Oracle JDK 8u202 Oracle JDK 11.0.1 Adopt OpenJDK11.0.1+13
7.2.2	9.0.27	Adopt OpenJDK 8u202b08	Oracle JDK 8u202 Oracle JDK 11.0.1 Adopt OpenJDK11.0.1+13

Note: Java 11 is supported, but not bundled in Confluence 7.2.

Confluence 7.1

Confluence	Tomcat	Bundled JRE	Tested JDKs
7.1.0	9.0.22	Adopt OpenJDK 8u202b08	Oracle JDK 8u202 Oracle JDK 11.0.1 Adopt OpenJDK11.0.1+13
7.1.1	9.0.27	Adopt OpenJDK 8u202b08	Oracle JDK 8u202 Oracle JDK 11.0.1 Adopt OpenJDK11.0.1+13
7.1.2	9.0.27	Adopt OpenJDK 8u202b08	Oracle JDK 8u202 Oracle JDK 11.0.1 Adopt OpenJDK11.0.1+13

Note: Java 11 is supported, but not bundled in Confluence 7.1.

Confluence 7.0

Confluence	Tomcat	Bundled JRE	Tested JDKs
7.0.1	9.0.22	Adopt OpenJDK 8u202b08	Oracle JDK 8u202 Oracle JDK 11.0.1 Adopt OpenJDK11.0.1+13
7.0.2	9.0.22	Adopt OpenJDK 8u202b08	Oracle JDK 8u202 Oracle JDK 11.0.1 Adopt OpenJDK11.0.1+13
7.0.3	9.0.22	Adopt OpenJDK 8u202b08	Oracle JDK 8u202 Oracle JDK 11.0.1 Adopt OpenJDK11.0.1+13
7.0.4	9.0.22	Adopt OpenJDK 8u202b08	Oracle JDK 8u202 Oracle JDK 11.0.1 Adopt OpenJDK11.0.1+13
7.0.5	9.0.27	Adopt OpenJDK 8u202b08	Oracle JDK 8u202 Oracle JDK 11.0.1 Adopt OpenJDK11.0.1+13

Confluence 6.15

Confluence	Tomcat	Bundled JRE	Tested JDKs
6.15.0	9.0.12	Adopt OpenJDK 8u192b12	Oracle JDK 8u202
6.15.1	9.0.12	Adopt OpenJDK 8u192b12	Oracle JDK 8u202
6.15.2	9.0.12	Adopt OpenJDK 8u202b08	Oracle JDK 8u202
6.15.3	9.0.17	Adopt OpenJDK 8u202b08	Oracle JDK 8u202

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

6.15.4	9.0.19	Adopt OpenJDK 8u202b08	Oracle JDK 8u202
6.15.5	9.0.19	Adopt OpenJDK 8u202b08	Oracle JDK 8u202
6.15.6	9.0.19	Adopt OpenJDK 8u202b08	Oracle JDK 8u202
6.15.7	9.0.21	Adopt OpenJDK 8u202b08	Oracle JDK 8u202
6.15.8	9.0.22	Adopt OpenJDK 8u222b10	Oracle JDK 8u202
6.15.9	9.0.22	Adopt OpenJDK 8u202b08	Oracle JDK 8u202
6.15.10	9.0.22	Adopt OpenJDK 8u202b08	Oracle JDK 8u202

Note: There was a known issue with Adopt OpenJDK 8u222b10, which was bundled with Confluence 6.15.8 ~~CONFSERVER-58784~~ CLOSED .

Confluence 6.14

Confluence	Tomcat	Bundled JRE	Tested JDKs
6.14.0	9.0.12	Adopt OpenJDK 8u192b12	Oracle JDK 8u202
6.14.1	9.0.12	Adopt OpenJDK 8u192b12	Oracle JDK 8u202
6.14.2	9.0.12	Adopt OpenJDK 8u192b12	Oracle JDK 8u202
6.14.3	9.0.12	Adopt OpenJDK 8u202b08	Oracle JDK 8u202

Confluence 6.13

Confluence	Tomcat	Bundled JRE	Tested JDKs
6.13.0	9.0.12	Oracle JDK 1.8.0_192	-
6.13.1	9.0.12	Oracle JDK 1.8.0_192	-
6.13.2	9.0.12	Adopt OpenJDK 8u192b12	Oracle JDK 1.8.0_192
6.13.3	9.0.12	Adopt OpenJDK 8u192b12	Oracle JDK 1.8.0_192
6.13.4	9.0.12	Adopt OpenJDK 8u202b08	Oracle JDK 8u202
6.13.5	9.0.19	Adopt OpenJDK 8u202b08	Oracle JDK 8u202
6.13.6	9.0.19	Adopt OpenJDK 8u202b08	Oracle JDK 8u202
6.13.7	9.0.22	Adopt OpenJDK 8u222b10	Oracle JDK 8u202
6.13.8	9.0.22	Adopt OpenJDK 8u202b08	Oracle JDK 8u202
6.13.9	9.0.22	Adopt OpenJDK 8u202b08	Oracle JDK 8u202
6.13.10	9.0.22	Adopt OpenJDK 8u202b08	Oracle JDK 8u202
6.13.11	9.0.22	Adopt OpenJDK 8u202b08	Oracle JDK 8u202
6.13.12	9.0.33	Adopt OpenJDK 8u252b09	Oracle JDK 8u251
6.13.13	9.0.33	Adopt OpenJDK 8u252b09	Oracle JDK 8u251
6.13.15	9.0.33	Adopt OpenJDK 8u252b09	Oracle JDK 8u251
6.13.17	9.0.33	Adopt OpenJDK 8u252b09	Oracle JDK 8u261

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

6.13.18	9.0.33	Adopt OpenJDK 8u252b09	Oracle JDK 8u261
6.13.19	9.0.40	Adopt OpenJDK 8u265b01	Oracle JDK 8u261
6.13.20	9.0.40	Adopt OpenJDK 8u265b01	Oracle JDK 8u261
6.13.21	9.0.40	Adopt OpenJDK 8u265b01	Oracle JDK 8u261
6.13.23	9.0.40	Adopt OpenJDK 8u265b01	Oracle JDK 8u291

Note: There was a known issue with Adopt OpenJDK 8u222b10, which was bundled with Confluence 6.13.7 ~~CONFSERVER-58784~~ CLOSED .

Note: Confluence 6.13.14 and 6.13.16 were internal releases.

Confluence 6.12

Confluence	Tomcat	Bundled JRE
6.12.0	9.0.11	Oracle JDK 1.8.0_181
6.12.1	9.0.11	Oracle JDK 1.8.0_181
6.12.2	9.0.12	Oracle JDK 1.8.0_181
6.12.3	9.0.12	Oracle JDK 1.8.0_181
6.12.4	9.0.12	Oracle JDK 1.8.0_181

Confluence 6.11

Confluence	Tomcat	Bundled JRE
6.11.0	9.0.10	Oracle JDK 1.8.0_162
6.11.1	9.0.10	Oracle JDK 1.8.0_162
6.11.2	9.0.11	Oracle JDK 1.8.0_181

Confluence 6.10

Confluence	Tomcat	Bundled JRE
6.10.0	9.0.8	Oracle JDK 1.8.0_162
6.10.1	9.0.8	Oracle JDK 1.8.0_162
6.10.2	9.0.10	Oracle JDK 1.8.0_162
6.10.3	9.0.19	Oracle JDK 1.8.0_202

Confluence 6.9

Confluence	Tomcat	Bundled JRE
6.9.0	8.0.51	Oracle JDK 1.8.0_162
6.9.1	8.0.51	Oracle JDK 1.8.0_162
6.9.2	8.0.52	Oracle JDK 1.8.0_162

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

6.9.3	8.0.52	Oracle JDK 1.8.0_162

Confluence 6.8

Confluence	Tomcat	Bundled JRE
6.8.0	8.0.50	Oracle JDK 1.8.0_162
6.8.1	8.0.50	Oracle JDK 1.8.0_162
6.8.2	8.0.51	Oracle JDK 1.8.0_162
6.8.3	8.0.51	Oracle JDK 1.8.0_162
6.8.4	8.0.52	Oracle JDK 1.8.0_162
6.8.5	8.0.52	Oracle JDK 1.8.0_162

Confluence 6.7

Confluence	Tomcat	Bundled JRE
6.7.0	8.0.48	Oracle JDK 1.8.0_162
6.7.1	8.0.48	Oracle JDK 1.8.0_162
6.7.2	8.0.48	Oracle JDK 1.8.0_162
6.7.3	8.0.50	Oracle JDK 1.8.0_162

Confluence 6.6

Confluence	Tomcat	Bundled JRE
6.6.0	8.0.47	Oracle JDK 1.8.0_152
6.6.1	8.0.48	Oracle JDK 1.8.0_162
6.6.2	8.0.48	Oracle JDK 1.8.0_162
6.6.3	8.0.50	Oracle JDK 1.8.0_162
6.6.4	8.0.50	Oracle JDK 1.8.0_162
6.6.5	8.0.51	Oracle JDK 1.8.0_162
6.6.6	8.0.51	Oracle JDK 1.8.0_162
6.6.7	8.0.52	Oracle JDK 1.8.0_162
6.6.8	8.0.53	Oracle JDK 1.8.0_162
6.6.9	8.0.53	Oracle JDK 1.8.0_181
6.6.10	8.0.53	Oracle JDK 1.8.0_192
6.6.11	8.0.53	Oracle JDK 1.8.0_192
6.6.12	8.0.53	Oracle JDK 1.8.0_192
6.6.13	8.0.53	Oracle JDK 1.8.0_202

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

6.6.14	8.0.53	Oracle JDK 1.8.0_202
6.6.15	8.0.53	Oracle JDK 1.8.0_202
6.6.16	8.0.53	Oracle JDK 1.8.0_202
6.6.17	8.0.53	Oracle JDK 1.8.0_202

Confluence 6.5

Confluence	Tomcat	Bundled JRE
6.5.0	8.0.47	Oracle JDK 1.8.0_144
6.5.1	8.0.47	Oracle JDK 1.8.0_144
6.5.2	8.0.47	Oracle JDK 1.8.0_152
6.5.3	8.0.51	Oracle JDK 1.8.0_162

Confluence 6.4

Confluence	Tomcat	Bundled JRE
6.4.0	8.0.43	Oracle JDK 1.8.0_131
6.4.1	8.0.43	Oracle JDK 1.8.0_131
6.4.2	8.0.46	Oracle JDK 1.8.0_144
6.4.3	8.0.47	Oracle JDK 1.8.0_144

Confluence 6.3

Confluence	Tomcat	Bundled JRE
6.3.0	8.0.43	Oracle JDK 1.8.0_131
6.3.1	8.0.43	Oracle JDK 1.8.0_131
6.3.2	8.0.43	Oracle JDK 1.8.0_131
6.3.3	8.0.43	Oracle JDK 1.8.0_131
6.3.4	8.0.43	Oracle JDK 1.8.0_131

Confluence 6.2

Confluence	Tomcat	Bundled JRE
6.2.0	8.0.41	Oracle JDK 1.8.0_112
6.2.1	8.0.43	Oracle JDK 1.8.0_131
6.2.2	8.0.43	Oracle JDK 1.8.0_131
6.2.3	8.0.43	Oracle JDK 1.8.0_131
6.2.4	8.0.43	Oracle JDK 1.8.0_131

Confluence 6.1

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Confluence	Tomcat	Bundled JRE
6.1.0	8.0.41	Oracle JDK 1.8.0_112
6.1.1	8.0.41	Oracle JDK 1.8.0_112
6.1.2	8.0.41	Oracle JDK 1.8.0_112
6.1.3	8.0.41	Oracle JDK 1.8.0_112
6.1.4	8.0.43	Oracle JDK 1.8.0_131

Confluence 6.0

Confluence	Tomcat	Bundled JRE
6.0.1	8.0.36	Oracle JDK 1.8.0_102
6.0.2	8.0.36	Oracle JDK 1.8.0_102
6.0.3	8.0.39	Oracle JDK 1.8.0_112
6.0.4	8.0.39	Oracle JDK 1.8.0_112
6.0.5	8.0.39	Oracle JDK 1.8.0_112
6.0.6	8.0.41	Oracle JDK 1.8.0_112
6.0.7	8.0.41	Oracle JDK 1.8.0_112

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

On this page:

- Confluence 8.4
- Confluence 8.3
- Confluence 8.2
- Confluence 8.1
- Confluence 8.0
- Confluence 7.20
- Confluence 7.19
- Confluence 7.18
- Confluence 7.17
- Confluence 7.16
- Confluence 7.15
- Confluence 7.14
- Confluence 7.13
- Confluence 7.12
- Confluence 7.11
- Confluence 7.10
- Confluence 7.9
- Confluence 7.8
- Confluence 7.7
- Confluence 7.6
- Confluence 7.5
- Confluence 7.4
- Confluence 7.3
- Confluence 7.2
- Confluence 7.1
- Confluence 7.0
- Confluence 6.15
- Confluence 6.14
- Confluence 6.13
- Confluence 6.12
- Confluence 6.11
- Confluence 6.10
- Confluence 6.9
- Confluence 6.8
- Confluence 6.7
- Confluence 6.6
- Confluence 6.5
- Confluence 6.4
- Confluence 6.3
- Confluence 6.2
- Confluence 6.1
- Confluence 6.0

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Supported Platforms FAQ

Q: How does Atlassian choose which JRE versions, application servers and databases to support?

For application servers and databases, we try to pick a good cross-section of open source options and popular commercial platforms. We then choose which JRE versions to support based on the recommended environments for these servers.

Q: What is a supported platform?

A supported platform is one that:

- Confluence is regularly tested on during the development cycle
- One that is available within Atlassian for support technicians and developers to reproduce problems
- Bugs raised against it will be given a high priority

Supporting a platform means we know how to get Confluence running in that environment and can troubleshoot Confluence issues within it. It does not mean we have any particular expertise beyond that. As such, we may not be able to provide assistance with customizing or tuning that application server or database. (Atlassian support is not a substitute for a good database administrator.)

Q: Can I get assistance with running Confluence on a platform that is *not* supported?

If you are running Confluence on an unsupported platform, then we can not guarantee providing any support for it. Furthermore, we will recommend that you switch to a platform which is supported.

Q: If you write your application to standards like J2EE, JDBC and SQL, doesn't that mean it should run on any compliant server?

Confluence is a complicated application and we commonly encounter interesting edge-cases where different servers have interpreted the specifications differently. Then again, each server has its own different collection of bugs.

Q: How can I get Atlassian to support Confluence on a new platform?

Supporting a new platform involves a significant investment of time by Atlassian, both up-front costs to set up new testing environments and fix any issues we might encounter and the ongoing costs involved in maintaining the application against this new environment in the future. As such, supporting a new platform is not something we will do unless we know there is significant demand for it.

Please be aware that your interest alone will not be enough for us to add support for your application server or database. We would need to see a significant number of votes on the issue raised in our public Jira site or a significant level of interest in our forums, before considering supporting that platform.

Q: My organization has standardized on an operating environment that Confluence does not support. What can I do?

In this situation, you have the following two options:

1. Run Confluence in the unsupported environment, with the caveats mentioned above.
2. Make an exception to your standardized operating environment and set up Confluence based on its supported platforms.

Migrate your Confluence site

Whether you're ready to make the move to cloud, or need the deployment and administrative flexibility of Data Center, we have everything you need to migrate successfully.

- Upgrade from Confluence Server to Data Center
- Migrate from Confluence Cloud to Server
- Migrating Confluence Between Servers
- Move to a non-clustered installation
- From Confluence Evaluation through to Production Installation
- Cloud Migration Assistant for Confluence

Considering a move to cloud? Check our the Server to cloud migration guide.

Upgrade from Confluence Server to Data Center

If you're a current Confluence Server customer looking to upgrade to Confluence Data Center, this page will help you create a free trial license and set up Data Center. There are several ways to get started with Confluence Data Center, depending on your current setup.

If youre installing Confluence Data Center for the first time with no existing Confluence Server data to migrate, check out how to install a Confluence Data Center trial.

On this page:

Set up Data Center
Upgrade to Data Center
Review and upgrade your apps
Upgrade your Confluence license
Set up your cluster

Set up Data Center

Things you should know about when setting up your Data Center:

Its your Confluence license that determines the type of Confluence you have: Server or Data Center. Confluence will auto-detect the license type when you enter your license key, and automatically unlock any license-specific features.

To upgrade from Confluence Server to Confluence Data Center, you will need a Data Center license. You can either purchase a full Data Center license or get a free trial license for 30 days. When your 30-day trial finishes, youll have the option to purchase a Data Center license and carry on using Confluence Data Center without losing any data youve created during the trial. If you decide Confluence Data Center is not for you, you can easily revert to your existing Server license.

> Note that as of February 15, 2024 PT, your Server products will reach the end of support.

See ourSupported Platformspage for information on the database, Java, and operating systems you'll be able to use. These requirements are the same for Server and Data Center deployments.

Apps extend what your team can do with Atlassian applications, so it's important to make sure that your team can still use their apps after migrating to Data Center.When you switch to Data Center, you'll be required to switch to the Data Center compatible version of your apps, if one is available.

SeeEvaluate apps for Data Center migrationfor more information.

To use Confluence Data Center you must:

- Have a Data Center license (you canpurchase a Data Center licenseor create an evaluation license atmy.atlassian.com)
- Use asupportedexternal database, operating system and Java version
- Use OAuth authentication if you haveapplication linksto other Atlassian products (such as Jira)

If you plan to run Confluence Data Center in a cluster there are some additional infrastructure requirements. SeeClustering with Confluence Data Centerfor more information.

> Theres a known issue during setup where a load balancer (or proxy) pings the server and breaks Confluence installation or migration to Data Center. See
>
> **CONFSERVER-61189** - Opening the base URL multiple times during Data Center migration will break the migration process. GATHERING IMPACT
>
> During installation, you need to disable load balancer health checks and make sure you dont open multiple tabs that point to the same Confluence URL.

Upgrade to Data Center

Review and upgrade your apps

If you have any apps installed on your site, you'll need to upgrade to the Data Center app version, if one is available. To avoid any impact to your apps, we recommend you do this before you enter your Confluence Data Center license key.Learn more about upgrading Server apps when you migrate to Data Center

Upgrade your Confluence license

To upgrade from Confluence Server to Confluence Data Center:

1. Go to > **General Configuration**
2. From the sidebar select**License details**.
3. Enter your Confluence Data Center license key.

Data Center features such as read-only mode, SAML single sign-on, and CDN will now be available.

Set up your cluster

If your organization requirescontinuous uptime, scalability, and performance under heavy load, you'll want to run Confluence Data Center with multiple nodes in a cluster.

To find out more about clustering, including infrastructure requirements, see Clustering with Confluence Data Center.

If you're ready to set up your cluster now, head toSet up a Confluence Data Center cluster.

Looking to migrate all your Atlassian applications to Data Center? Weve got you covered:

- Upgrade from Bitbucket Serverto Bitbucket Data Center
- Migrate to Crowd Data Center
- Migrate to Confluence Data Center
- Migrate to Jira Data Center

Considering moving to cloud? Plan your cloud migration.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Migrate from Confluence Cloud to Server

Important changes to our server and Data Center products

Weve ended sales for new Server licenses, and will end support for server on February 15, 2024. Were continuing our investment in Data Center with several key improvements. Learn what this means for you

This page is for people who are currently using Confluence Cloud, and wish to move to Confluence Server (a self-managed Confluence site).

Not moving from Cloud to Server?

These resources will help you plan your migration from:

- Confluence Server to Cloud
- Confluence Server to Data Center
- Confluence Server to Server

On this page:

- Before you begin
  - Minimum Confluence version
  - Features and app availability
  - Templates
  - Team Calendars and Questions data
  - Migration approach
  - Infrastructure and database choice
  - Licenses
  - Account visibility
- Migration steps
  - Step 1: Check your apps
  - Step 2: Install Confluence Server or Data Center
  - Step 3: Export your Confluence Cloud Site
  - Step 4: Import your Confluence Cloud site export file
  - Step 5: Recover sys admin permissions
  - Step 6: Install any apps
  - Step 7: Check your application links
- Troubleshooting

Before you begin

There's a few things to understand before you begin this process. Ready to migrate? Skip to the migration steps

Minimum Confluence version

You can migrate from Confluence Cloud to **Confluence Server / Data Center 6.0 or later** only. You can't import Cloud data (either the whole site or individual spaces) into any earlier versions of Confluence.

We recommend installing either latest version of Confluence, or the latest Enterprise Release. TheConfluence Upgrade Matrix will help you choose the right version for your organisation.

Features and app availability

Some Cloud features won't be available in Confluence Server or Data Center. The navigation and user experience will also be different in some places. The core functionality of Confluence is the same however.

Marketplace apps are not automatically migrated. When you set up your Confluence Server or Data Center site, you'll need to reinstall each of your apps.

Not all apps are available for both Cloud and Server / Data Center. When planning your migration, we recommend you check that your essential apps are available for Server / Data Center in theAtlassian Marketplaceand make a list of the ones you'll need to reinstall.

Templates

All pages that were created from a template will be migrated.

However, any custom templates you may have created in your Confluence Cloud site will not be migrated. You'll need to re-create your templates once your migration is complete.

You should also be aware that the range of built-in templates (known as blueprints) is much smaller in Confluence Server and Data Center, so some of the default templates you have previously used may not be available.See the full list of blueprints

Team Calendars and Questions data

Confluence Questions and Team Calendars data can't be migrated, as there is currently no way to export this data from Confluence Cloud.

Migration approach

You can choose to migrate your entire site in one go, or to import your team's content, space by space.

A **full site migration** involves a full site export (backup), and importing this file into Confluence Server or Data Center. Users and groups are included in this export. All spaces will be migrated, including archived spaces and personal spaces.

SeeMigration steps below to find out how to do this.

A **space by space migration** involves exporting each space individually, and importing these files into Confluence Server or Data Center one at a time. This means you can choose which spaces you want to migrate, or migrate in stages over time. Users and groups are not automatically migrated. If you've connected Confluence Server or Data Center to an external user directory, or have already populated your new site with user accounts, we'll attempt to attribute content to the right people on import.

SeeImport a space from Confluence Cloud if you plan to migrate your spaces one by one.

Infrastructure and database choice

Where you choose to host Confluence Server or Data Center is up to you. SeeSupported Platforms to find out which operating systems and databases are supported.

You can use any database listed on theSupported Platformspage, but you don't already have a database server, we recommend PostgreSQL, which is what Confluence Cloud is run on.

Licenses

You will need a new license to migrate to Confluence Server or Data Center. Your existing Confluence Cloud license can't be used. You can get a new license athttps://my.atlassian.com. You'll also need new licenses for any paid Marketplace apps.

Account visibility

In Confluence Cloud, people can choose not to make their profile information visible. This means when a Cloud site is imported into Server, user account information such as their full name, may not be included.

As long as you are logged in as a Site Admin when you complete the site export, email addresses will always be included, and used as the username when the user accounts are created. Users can then log in, and update their profile to provide the missing information.

Migration steps

This page will guide you through a **full site migration**. SeeImport a space from Confluence Cloudif you plan to migrate your spaces one by one.

Step 1: Check your apps

To check your apps are compatible:

1. In Confluence Cloud, go to**Settings**>**Manage Apps**.
2. Make a note of all**User-installed** apps.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

3. Go tohttps://marketplace.atlassian.comand look up each app to see if a Server or Data Center edition is available.

Step 2: Install Confluence Server or Data Center

The way you do this depends on whether you're migrating to Server or Data Center, and how you plan to host the application.

SeeConfluence Installation Guide for links to all the installation options.

Step 3: Export your Confluence Cloud Site

To export your Confluence Cloud site:

1. Log in to Confluence Cloud as a **Site Admin**
2. In Confluence Cloud, go to **Settings** > **Backup Manager**
3. Follow the prompts to back up the site, and download the XML file.

The file will include all spaces and pages (including attachments), and all your users and groups.

Step 4: Import your Confluence Cloud site export file

Unless your site export file is quite small (less than 25mb) we recommend importing via the home directory method.

> ⚠ The import will overwrite all spaces, pages, and user accounts in your site - including your administrator account. You'll recover that account in the next step.
>
> You should back up your database, home directory, and installation directory before you begin, in case you need to roll back.

To import a site from the home directory:

1. Copy your export file to`<confluence-home>/restore.`
   (If you're not sure where this directory is located, the path is listed in the**Backup and Restore**screen)
2. Go to> **General Configuration**>**Backup and Restore**.
3. Select your site export file under**Import from home directory**
4. Make sure **Build Index**is checked so that your index is created automatically.
5. Choose**Import**.

SeeRestore a Site for more information about the import process.

Step 5: Recover sys admin permissions

When you import a site export file, all user accounts are overwritten, including the system administrator account that was created when you installed Confluence. Your existing Cloud Site Admin account will not automatically have system administrator permissions for Confluence Server or Data Center.

To recover system administrator permissions:

1. Stop Confluence.
2. Edit`<installation-directory>/bin/setenv.sh`or`setenv.bat`and add the following system property, replacing`<your-password>`with a unique, temporary password.

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

   SeeConfiguring System Propertiesfor more information on using system properties.
3. Start Confluence**manually**(don't start Confluence as a service).
4. Log in to Confluence with the username**recovery_admin**and the temporary password you specified in the system property.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

5. Go to > **User Management** > **Add Users**.
6. Enter the details for your new system administrator account and hit **Save**. Make sure to use a strong password.
7. Choose **Edit Groups** and select the `confluence-administrators` group. This is a super-group that has system administrator permissions.
8. Log out, and confirm that you can successfully log in with your new account.
9. Stop Confluence.
10. Edit`<installation-directory>/bin/setenv.sh`or`setenv.bat`and remove the system property.
11. Restart Confluence using your usual method (manually or by starting the service).

SeeRestore Passwords To Recover Admin User Rights for more information on this process.

Step 6: Install any apps

To re-install your apps:

1. Log in to Confluence Server or Data Center as an administrator.
2. Go to > **Manage apps**.
3. Follow the prompts to search for or upload the apps you identified in step 1. You'll need to purchase new licenses for these apps.

Remember that Team Calendars and Questions data is not included in your export, and cannot be migrated from Cloud at this time.

Step 7: Check your application links

If you had multiple Cloud products, such JIRA Software,you may need to make some changes to the application links.

To remove or update application links:

1. Go to > **General Configuration**> **Application Links**.
2. Follow the prompts to check and update any application links that are now pointing to the wrong place.

If you're unable to remove the Jira Cloud application link from your Confluence after the import, you'll need to remove those references directly from the Confluence database. SeeAlternative Methods of Deleting Application Links in Confluence.

Troubleshooting

There are a few known issues that you might encounter when importing your Cloud site.

Can't load pages in your new site

If you experience problems loading pages after the import, head to > **General Configuration**to check your base URL as the port may have changed.

User management admin screens are missing

This is a fairly uncommon problem caused by a dark feature flag that is included in your Cloud site export file. See ~~CONFSERVER-35177~~ CLOSED for a workaround.

Jira issues macros are broken

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

If your Confluence Cloud site has macros that depend on the Application Links back to a Jira Cloud instance, and you are migrating Jira as well, these references will need to be updated to work properly. See APL-1144 NEEDS TRIAGE for a workaround.

You can also edit the XML file prior to importing it into Confluence Server or Data Center, or by bulk editing those references in Confluence database. SeeHow to bulk update JIRA Issue Macro to point to a different JIRA instance.

User mentions are broken

When a page with user mentions is migrated from Confluence Cloud to Confluence Server or Data Center using either site or space migration, the mentions display as "broken link". For a Cloud-to-Cloud migration via Confluence Data Center, the mentions display as "@unlicensed user". See CONFSERVER-79583 IN PROGRESS for the workaround.

Broken anchor links

Confluence Cloud replaces anchor macros with web links that are not compatible with Data Center versions after migration. See CONFSERVER-79006 GATHERING IMPACT for the workaround.

Users' favourites (starred pages, or saved for later) are missing

If you find that some of your users' favorites (pages saved for later) are missing due to CONFSERVER-36348 GATHERING IMPACT . See How to restore missing favorites after import from XML for more information.

Full-width space templates cause failed or incomplete space imports

If you import an XML space from Confluence Cloud containing space templates with the full-width property it will result in a failed import, or only a partial import. See CONFSERVER-80146 GATHERING IMPACT for a workaround.

Some user accounts are missing or created without user details

Users in Confluence Cloud have the ability to change their profile visibility settings. To ensure all user data is included in the export, ask a site admin to perform the export.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Migrating Confluence Between Servers

This page describes how to move Confluence between physical servers using the same or a different operating system.

It doesn't coverdatabase migrationorupgradingyour Confluence version. We suggest you do each of these steps separately.

On this page:

- Transferring Confluence to another server

Transferring Confluence to another server

To transfer Confluence to another server you will copy the home and install folders straight into an identical external database and user management setup. If your new server is using a different operating system there may be some additional changes at step 4.

1. Run the Confluence installer on your new server
2. Shut down Confluence on both your old and new servers
3. If you're using Oracle or MySQL, copythe driversfrom your old server to the new one
4. Delete the contents of the home directory on your new Confluence server, then copy in the contents of the home directory from your old Confluence server.
5. Make any additional changes required for your environment.

   If the path to your home directory is different on the new server open the `Confluence_install_directory/confluence/WEB-INF/classes` directory and edit `confluence-init.properties` by changing the line starting with '`confluence.home=`'.

   If you have also moved your database from one server to another you can change the JDBC URL in`<confluence.home>/confluence.cfg.xml`if you are using a direct JDBC connection or in the definition of your datasource (if you are connecting via a datasource).

   If you're migrating from **Windows to Linux**, you'll need to replace the backslashes with forward slashes in the following lines in`confluence.cfg.xml`:

   ```
   <property name="attachments.dir">${confluenceHome}/attachments</property>
   <property name="lucene.index.dir">${localHome}/index</property>
   <property name="webwork.multipart.saveDir">${localHome}/temp</property>
   ```

   If you're migrating from **Linux to Windows**, you'll need to replace the forward slashes with backslashes in the following lines in`confluence.cfg.xml`:

   ```
   <property name="attachments.dir">${confluenceHome}\attachments</property>
   <property name="lucene.index.dir">${localHome}\index</property>
   <property name="webwork.multipart.saveDir">${localHome}\temp</property>
   ```

6. Copy the `<confluence-install>/conf/server.xml`file from your old server to the same location on your new server
7. If you use a data source, ensure the data source points to the new database. See Configuring a datasource connection.
8. Start Confluence, then head to**General configuration**>**License Details**to add your license key

Westrongly recommend you perform arebuild of your content indicesafter performing a migration, to ensureConfluence search works as expected.

Move to a non-clustered installation

This page outlines how to switch from a clustered Confluence deployment to a non-clustered deployment. In these instructions we'll assume that you'll use one of your existing cluster nodes as your new, non-clustered installation.

If you have a valid Server license and want to move from Data Center to Server, see Moving from Data Center to Server.

Run Confluence in a cluster with one node

If you no longer need clustering for high availability or managing load, you can simply reduce the number of application nodes in your cluster to one.There are some advantages to this setup, as it is very easy to add more nodes if you require them in future, but there is a small performance overhead as Confluence will still operate as a cluster.

Move to a non-clustered installation

If you no longer need clustering, and want to avoid the overhead that comes from running a cluster with just one node, you can go back to a non-clustered (sometimes known as standalone) Data Center installation.

In these instructions we'll assume that you'll use one of your existing cluster nodes as your new, non-clustered installation. You'll also need to make some infrastructure changes as part of the switch. We recommend completing this process in a staging environment, and running a set of functional tests, integration tests, and performance tests, before making these changes in production.

Terminology

In this guide we'll use the following terminology:

- Installation directory The directory where you installed Confluence.
- Local home directory The home or data directory stored locally on each cluster node (if Confluence is not running in a cluster, this is simply known as the home directory).
- Shared home directory The directory you created that is accessible to all nodes in the cluster via the same path.

1. Shut down Confluence

Make sure read only mode is turned off, then stop Confluence on all cluster nodes before you proceed.

2. Configure your load balancer

Configure your load balancer to redirect traffic away from all Confluence nodes, except the node you plan to keep.

If you no longer need your load balancer, you can remove it at this step.

3. Move items in the cluster shared home back to local home

To move everything back to your local home:

1. Create a directory called`/shared-home`in the`<local home>` directory on the node you plan to keep (if you removed this directory when you set up clustering).
2. Move the following directories and files from your`<shared home>`directory to the`<local home>/shared-home`directory
    - config
    - confluence.cfg.xml
    - dcl-document
    - dcl-document_hd
    - dcl-thumbnail
3. Move the remaining contents of your <shared home> directory to the root of your <local home> directory. Make sure your attachments directory is moved as part of this step.

Your cluster's shared home directory should now be empty.

⚠Make sure you don't accidentally overwrite the`confluence.cfg.xml`in your local home directory. The`confluence.cfg.xml`file from your shared home directory doesn't contain the same parameters as the one in your local home directory.

From Confluence 7.12, you can choose to skip this step and keep your existing shared home directory. For example, this may be beneficial if you're using elastic storage for the `<shared home>/attachments` directory and want to keep that setup.

4. Modify cluster properties

1. Take a backup of the existing`<local home>/confluence.cfg.xml`
2. Edit`<local home>/confluence.cfg.xml`
3. Change the`setupType` parameter from `cluster` to`custom`:

```
<setupType>custom</setupType>
```

4. Remove all cluster properties that begin with`confluence.cluster`.

   Here are some example cluster properties that should be removed. These will vary depending on how you configured your cluster.

   ```
   confluence.cluster
   confluence.cluster.address
   confluence.cluster.home
   confluence.cluster.interface
   confluence.cluster.join.type
   confluence.cluster.name
   ```

   ⚠ If you chose to keep your shared home directory at the previous step, do not remove the `confluence.cluster.home` property, or Confluence will not know where to find your shared home, or attachments directory.

5. Save the file.

5. Start Confluence

Restart Confluence.

To confirm you're now running a standalone installation, go to > **General Configuration** > **Clustering**.

The active cluster should no longer appear. Instead, you'll see information about getting started with clustering, and the option to enable cluster mode.

Additional steps if you have a Synchrony cluster

If you also have a Synchrony cluster, but would prefer to let Confluence manage Synchrony for you, you'll need to make some additional changes.

SeeMigrate from a standalone Synchrony cluster to managed Synchrony. This guide assumes you're running Confluence in a cluster, but the steps are similar for a non-clustered installation.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

From Confluence Evaluation through to Production Installation

Important changes to our server and Data Center products

Weve ended sales for new Server licenses, and will end support for server on February 15, 2024. Were continuing our investment in Data Center with several key improvements. Learn what this means for you

On this page:

- Step 1. Set up your evaluation Confluence site
- Step 2. Add users and content to your evaluation site
- Step 3. Look for interesting Marketplace apps as part of your evaluation
- Step 4. Set up your production Confluence site

Related pages:

- Supported Platforms
- Add and Invite Users
- Getting Started as Confluence Administrator
- Confluence installation and upgrade guide

So, you want to try Confluence on an evaluation installation, then move to a production installation when you are ready? This page gives an overview of the steps to follow.

Assumptions:

- This page starts with telling you how to install an evaluation Confluence site. If you have already finished evaluating Confluence, you can safely skip steps 1 to 3.
- Your production installation will be an installed version of Confluence, **not** a Confluence Cloud site.
- You will evaluate Confluence on an installed version too, not a Confluence Cloud site.

If you are using Confluence Cloud to evaluate Confluence, please refer to the following guide when you want to move to an installed version: Migrate from Confluence Cloud to Server.

Step 1. Set up your evaluation Confluence site

If you have already set up an evaluation Confluence site, you can skip this step.

Below is a summary of the installation and setup procedure, focusing on the choice of database.

To install Confluence:

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: Go to >**User management**.

Step 3. Look for interesting Marketplace apps as part of your evaluation

If you have finished evaluating Confluence, you can skip this step.

Apps, also called plugins or add-ons, provide additional features that you can install into your Confluence site. Some of them are provided free of charge. Many of the commercial apps are available free for an evaluation period.

You can browse and download app on the Atlassian Marketplace. You can also find apps via the Confluence user interface, which interacts with the Atlassian Marketplace for you.

To find useful apps via the Confluence user interface:

1. Go to > **Manage apps**.
2. Choose **Find new add-ons**.

Step 4. Set up your production Confluence site

When you are ready to move from an evaluation site to a production site, you need to migrate to a production-ready database. This involves installing a new Confluence site with a new database, and instructing Confluence to copy the data from your evaluation site to the new site. You will also need to check some important configuration settings, and define your backup strategy. The instructions below lead you through all the steps required.

Migrating your data to a production database:

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.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

- Decide on proxy setup and other settings that determine where Confluence fits into your network. See Web Server Configuration.
- Consider setting up a secure connection via SSL. SeeRunning Confluence Over SSL or HTTPS.
- Read our guidelines on security. See Best Practices for Configuring Confluence Security.
- Decide whether you will manage your users in Confluence or connect to an external LDAP directory. See Configuring User Directories.
- Decide whether you want to allow public (anonymous) access to your site. SeeSetting Up Public Access.
- Set up your permission scheme. See Permissions and restrictions.
- Connect Confluence to Jira applications such as Jira Software or Jira Service Management or other applications. See Linking to Another Application.

Defining your backup strategy:

By default, Confluence will create daily XML backups of your content and user data. This is suitable when you are evaluating Confluence. When you move to a production site, you need more robust backup procedures and technologies. See Production Backup Strategy.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Cloud Migration Assistant for Confluence

robots	noindex
description	Step-by-step instructions on how to use the Confluence Cloud Migration Assistant to move your data from Confluence Server or Data Center to Confluence Cloud.

robots	noindex
robots	noindex

Before you migrate, check your cloud organization

Were currently rolling out changes that may affect your migration experience. From yourorganization at admin.atlassian.com, if the **Users** list and **Groups** list are under the **Directory** tab, you have the improved user management experience. This means that the users and groups across sites will be merged under the organization. Read more about how groups and permissions are migrated. If you have any concerns, contact support.

For a test migration or UAT, we recommend that your test cloud site is not part of the organization that also hosts your prod site. The prod site should be hosted in a different organization. This is to ensure smooth migration of the relevant users and groups.

The Confluence Cloud Migration Assistant is an app that helps you easily move content, users, and groups from Confluence Server or Data Center to Confluence Cloud. Built and maintained by Atlassian, the app is free to install and use.

With the app, you can choose what you want to move to the cloud, start migrating at your convenience, and monitor the progress of everything throughout the migration process.

Important changes to our server and Data Center products

Weve ended sales for new Server licenses, and will end support for server on February 15, 2024. Were continuing our investment in Data Center with several key improvements. Learn what this means for you

When to use the Confluence Cloud Migration Assistant

- When you want to move users or data from Confluence Server or Data Center to Confluence Cloud.
- When you want to assess your apps before moving from Confluence Server or Data Center to Confluence Cloud.
- When you want to run a test or trial migration from Confluence Server or Data Centerto Confluence Cloud.
- When the Atlassian Support team has recommended using the app.

The Confluence Cloud Migration Assistant will not work for Jira products. You can download the Jira Cloud Migration Assistant for Jira migrations to cloud.

Before you begin

Make sure you have reviewed the server to cloud migration guide. This guide will walk you through the migration process step-by-step and help you identify what to look out for.

Before attempting a test or production migration, ensure you've completed all of the steps for the Confluence Cloud Migration Assistant in thepre-migration checklist. The checklist will help you prepare yourself and your data for migration, and ensure you avoid common sources of migration failure.

Install the Confluence Cloud Migration Assistant app

If your Confluence Server site is version 6.13 or above you won't need to install anything because it comes pre-installed, although you may be asked to update the app.

To install the app on versions 5.10 to 6.12:

1. In Confluence Server go to > **Manage apps**.
2. Choose**Find new add-ons**.
3. Search for the**Confluence Cloud Migration Assistant**app.
4. Choose**Install**and you're all set.

Alternatively, you can install it from theAtlassian Marketplace.

Once installed, you can access the migration assistant by going to**Confluence Administration**>look for the**Atlassian Cloud**category > select **Migration Assistant.**

If your Confluence Server site is behind a firewall, you'll need to allow access to the domain: **atlassian.com**

Use the migration assistant to assess your apps

Carrying out an assessment of your apps helps you to establish which apps are needed for a migration.

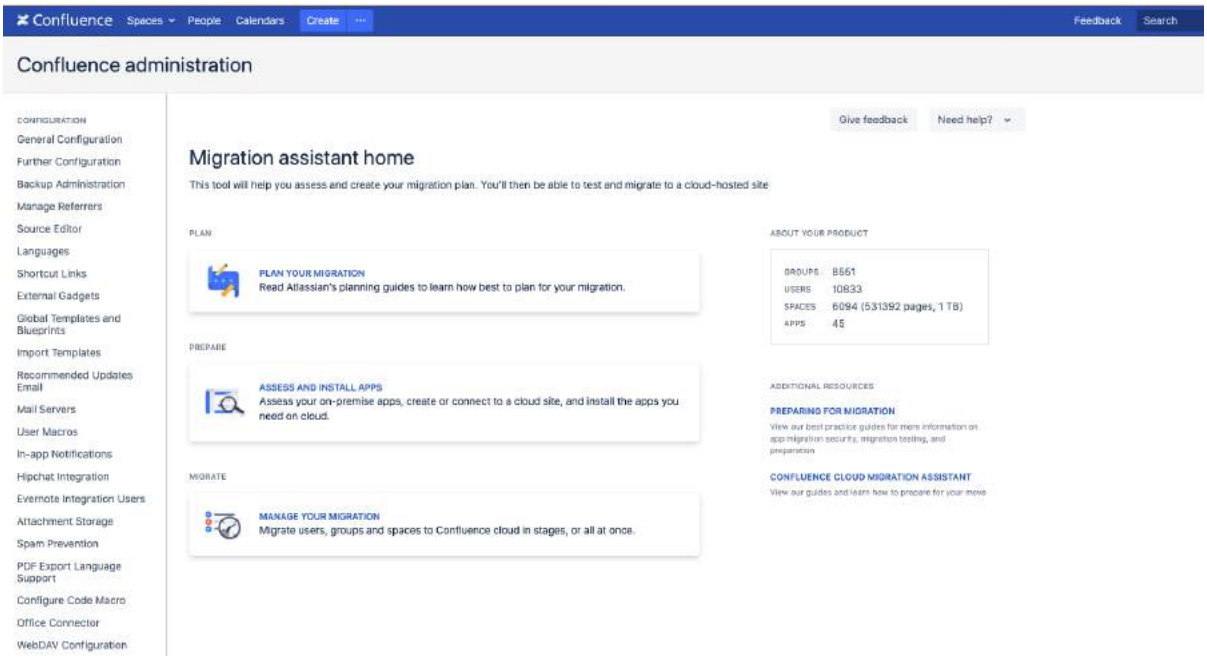

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

You can find step-by-step instructions for this process in Assessing and migrating apps with the Confluence Cloud Migration Assistant.

Check for possible data conflicts in your cloud site

You can reduce the risk of running into issues, or the migration failing, if you conduct somemanualchecks in your server and cloud sites.

1. Check for groupconflicts

Make sure that there are**no groups already in your cloud site with the same name as groups from your server site,**unless you are intentionally trying to merge them.

If we find a group in your server site that has the samename as a groupin your cloud site (either Jira or Confluence), we will merge the users from the server group into the cloud group.The server group users will inherit the permissions of the cloud group.This also applies to groups with Jira product access that have the same name as a Confluence group you are migrating. This is because all users and groups are managed in a central location in your cloud site.

If you dont want this to happen, youll need to make sure all groups across server and cloudhave unique namesb efore running your migration.

The following groups manage admin access and are **blacklisted.** They will not be migrated at all: "site-admins", "system-administrators", "atlassian-addons", "atlassian-addons-admin". Users in these groups will still be migrated; if you want them to be in one of the blacklisted groups youll need to manually add them after migration.

2. Check for space keyconflicts

Before migrating,check that there are no spaces with the samespace keybetween your server and cloud sites.

If a space from your server site has the samespace keyas a space in your cloud site your migration will fail. This is because every space in Confluence Cloud must have a unique space key. If you find aconflict youcan:

- delete duplicate spaces from yourcloudorserversites
- reset your cloud site
- choose not to migrate these spaces

If the migration assistant finds a conflict, the space will not migrate.

If a space key conflict is caused by a previous test migration you canreset your cloud sitebefore migrating.

Use the app to set up and run your migration

Once you have the app installed, there are five key steps to set up and run your migration from server or Data Center to cloud:

1. Connect to cloud
2. Choose what to migrate
3. Check for errors
4. Review your migration
5. Migrate

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

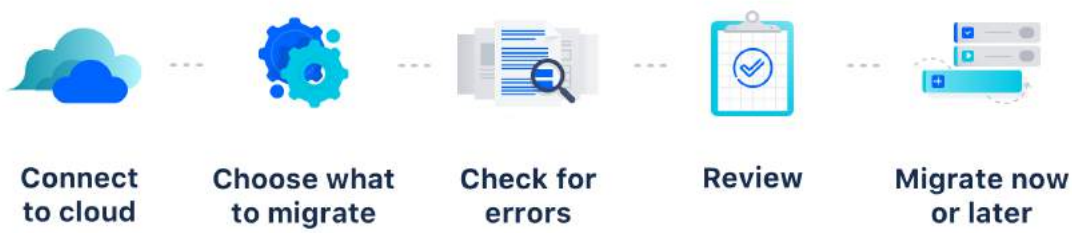

The sections below describe each step in detail and explain some common errors that you may come across. If you have technical questions or issues while using the migration assistant, get in touch with our support team.

Running a test migration

We strongly recommend doing a trial run of your migration to a test or staging site before running your final migration. Check out our guidance on testing your migration.

1. Connect to your destination Confluence Cloud site

Youll be asked to add a name for your migration and choose which cloud site you would like to migrate to. You need to be an admin in both your server and the destination cloud sites.

If you have already connected a cloud site, you should see it in the dropdown. If there is nothing there, you will need to either connect a new cloud site or sign up for a new cloud license.

When youre ready to go, check the box to allow Atlassian to move your data from your server site to your cloud site. If youre unable to grant Atlassian this access, you wont be able to migrate with the migration assistant and will need to do a space import instead.

If your Confluence Server site is behind a firewall, you'll need to allow access to the domain:**atlassian.com**. You also might need to allow access tootheer Atlassian domains.

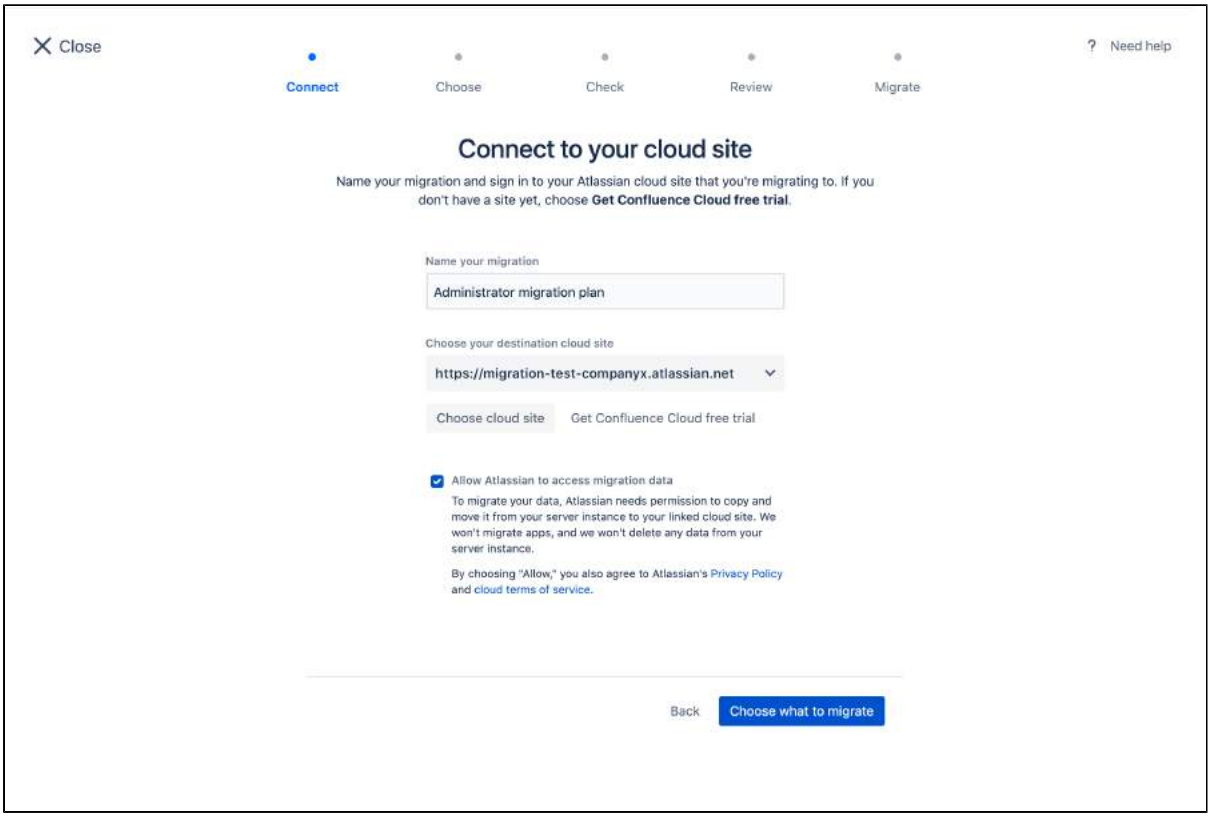

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

2. Choose what to migrate

You can migrate everything together or break it up into different stages.

You can choose:

- all or some of your users and groups
- which individual spaces (and their attachments) you'd like to migrate

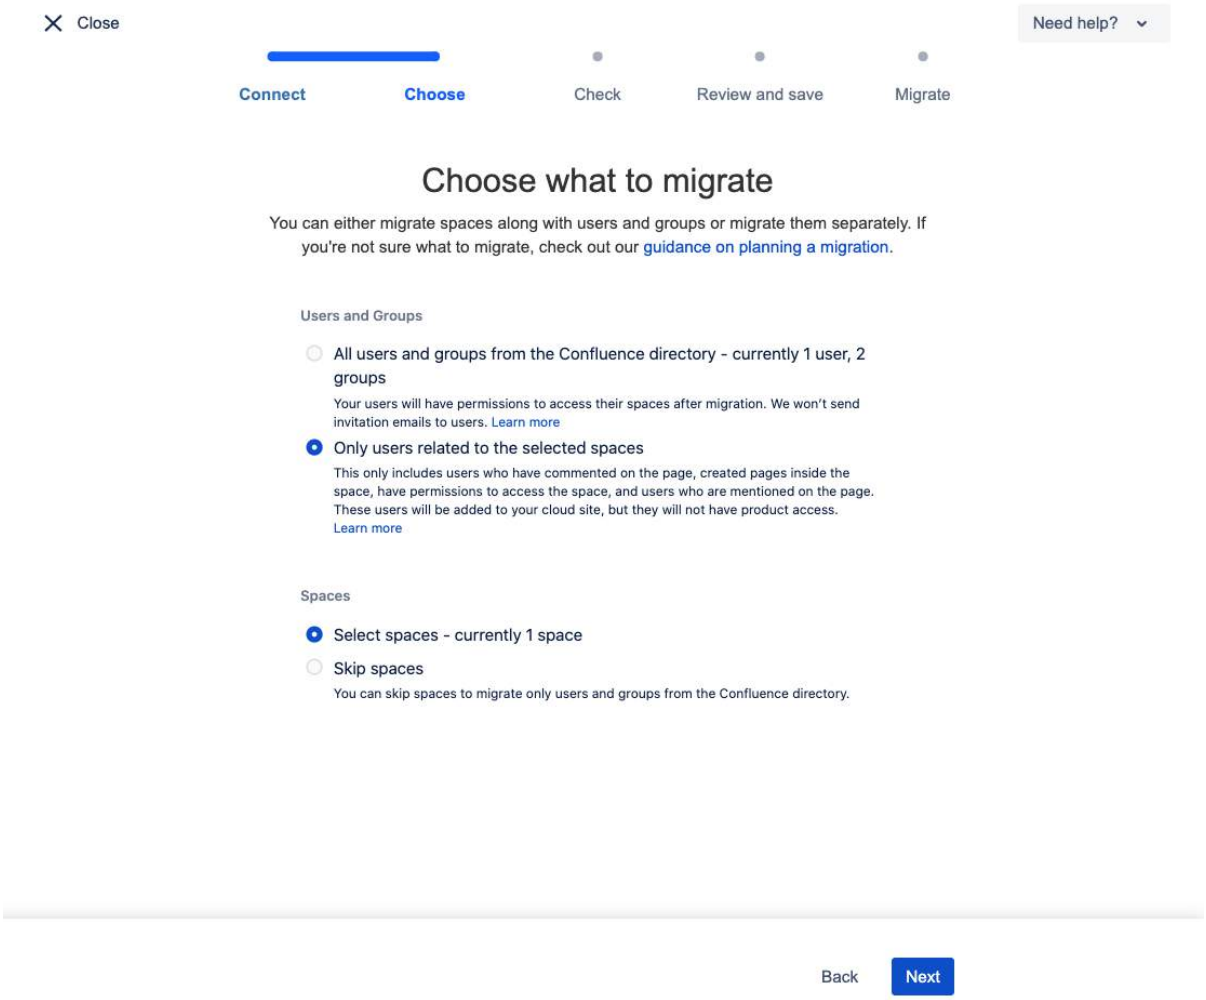

Users and groups

You can choose to either migrate all or some of your users.

If you choose the migrate your users, the first time you do so all your users will be added to your cloud site. Every migration, after the first, we will just link your data to the users that already exist in cloud. If you have a large userbase we suggest following our recommendations.

When you migrate your users, they will be added to their groups when they get to cloud. You will need to review and approve group permissions after you migrate. When you approve group permissions, your users will be given Confluence access and will be added to your bill.

We wont send an invitation to your users. To invite your users you can choose to send an invitation from the **Ad ministration** space after you have migrated, or send a link for them to log in themselves.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

When you select **Only users related to the selected spaces**under users and groups, we will still migrate some user data connected to the spaces you are migrating. This is to make sure that mentions, comments, and page history stay active.

User data that will be migrated every time includes:

- full name
- username (discarded after migration)
- email address

We will only migrate this information for **users directly connected to the spaces** you are migrating. We will not give these users product access or add them to any groups. They will appear in your cloud site user list.

If you choose to migrate users later, their product and group access will be updated.

Also, if you choose not to migrate users and groups and you have a space permission granted by a group that don't exist in cloud, the Confluence Cloud Migration Assistant will not migrate the respective space permission. To avoid this scenario, we recommend you to create the specific group in the cloud site before migration.

Other things to be aware of when migrating users and groups:

- Users are migrated using email address as the source of truth. On subsequent migrations, the migration assistant will link users by email address rather than re-migrating them. Check out ourtips for migrating a large number of users.
- You must validate all your user accounts (email addresses) before migrating to cloud. Migrating unknown user accounts can potentially allow unauthorized access to your cloud sites. For example, if you had users in your server instance with emails that you dont own, say email@example.com, you might be inviting someone who owns @example.com to your site in cloud.
- Confluence Cloud is subscription-based and billed on a per-user basis. If you plan to migrate your users, make sure you check the licensing options available.
- If you use an external user management system, we recommend synchronizing it with your local directory **before** migrating. This is to make sure that your users and groups are up to date before you transfer any data.
- Users with **disabled** status in your server site will be migrated as **active but without any product access**. This means they **will not be counted** as active Confluence users for billing purposes.
- If we find a group in your server site that has the same name as a group in your cloud site, we will merge the users from the server group into the cloud group.
- Global settings and global site permissions are not migrated with this tool. Youll need to set these manually after migration.
- If you have users that already exist in your destination cloud site and you choose to migrate users with this app, the following will occur:
  - If a user has **product accessin cloud**, but has **disabled status in your server site**, they will continue to have product access in cloud after migration.
  - If a user **does not have product access in cloud**, but is **enabled in your server site**, they will be granted product access through the migration process.

If you use Confluence as a knowledge base for Jira Service Management (formerly Jira Service Desk), your Jira Service Management users may also be migrated along with your Confluence users. This will happen if you can see your Jira Service Management users in the **cwd_user** table in Confluence.

Spaces

If you want to migrate all or some of your spaces choose **Select spaces** from the options. You will then be able to select what spaces you want to migrate. If you arent migrating any spaces you will be taken straight to check for errors.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

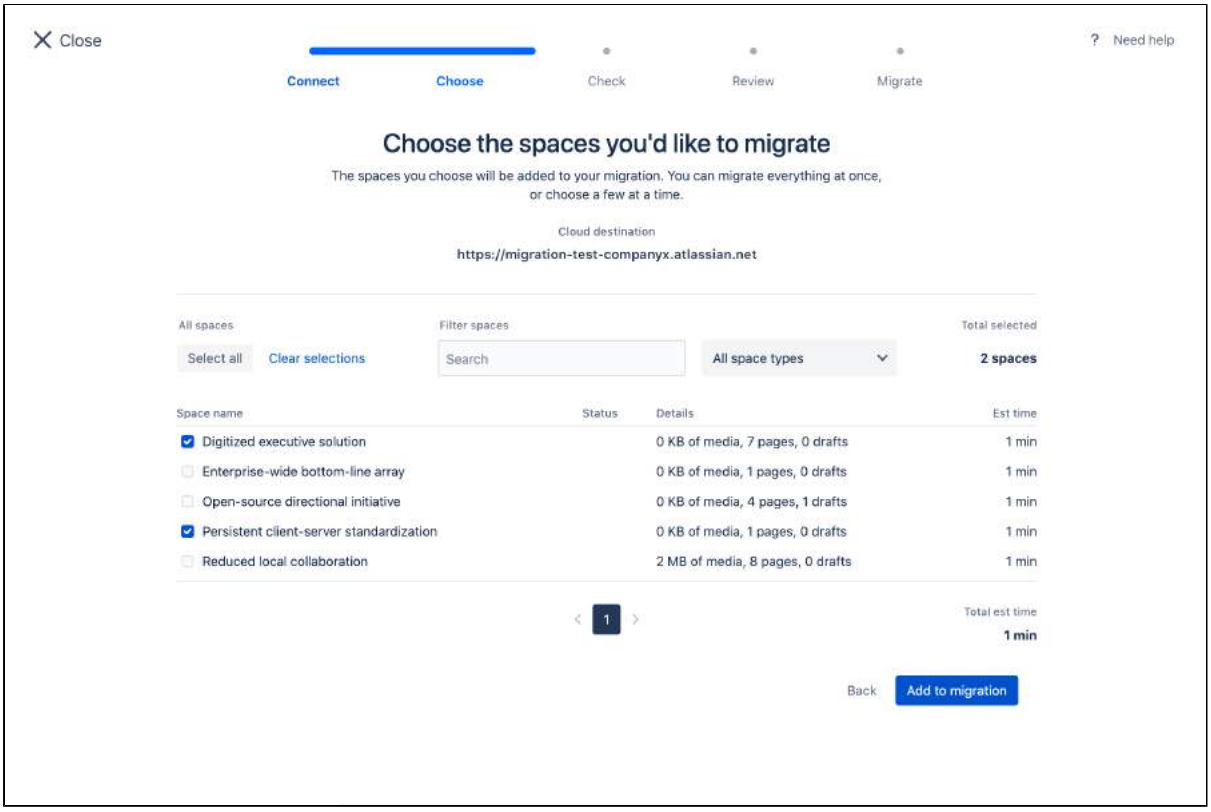

Select the spaces you want to add to your migration. You can filter the list or search for particular spaces, or click **Select all** if you want to migrate everything at once. You wont be able to migrate spaces with space keys that already exist in your Confluence Cloud destination site.

If a space has a MIGRATED status, we have detected that you have already migrated this space to the same cloud site.

If a space has a QUEUED status, it has already been added to a migration that is waiting to be run.

> If you have lots of spaces and attachments or you are on Data Center, you might want to break the migration up into a few smaller migrations. The migration assistant can be slow to load and process tasks when there is a lot to manage.

When you've chosen all your spaces, select**Add to migration**.

3. Check for errors

In this step, the Confluence Cloud Migration Assistant will review your migration and check for common errors. It will check if your:

- migration assistant app is up to date
- users have valid and unique email addresses
- groups will merge through the migration process
- spaces already exist in your cloud site
- spaces are publicly available and searchable online

> You may also encounter other issues during the migration process; this step only checks for the issues mentioned here.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

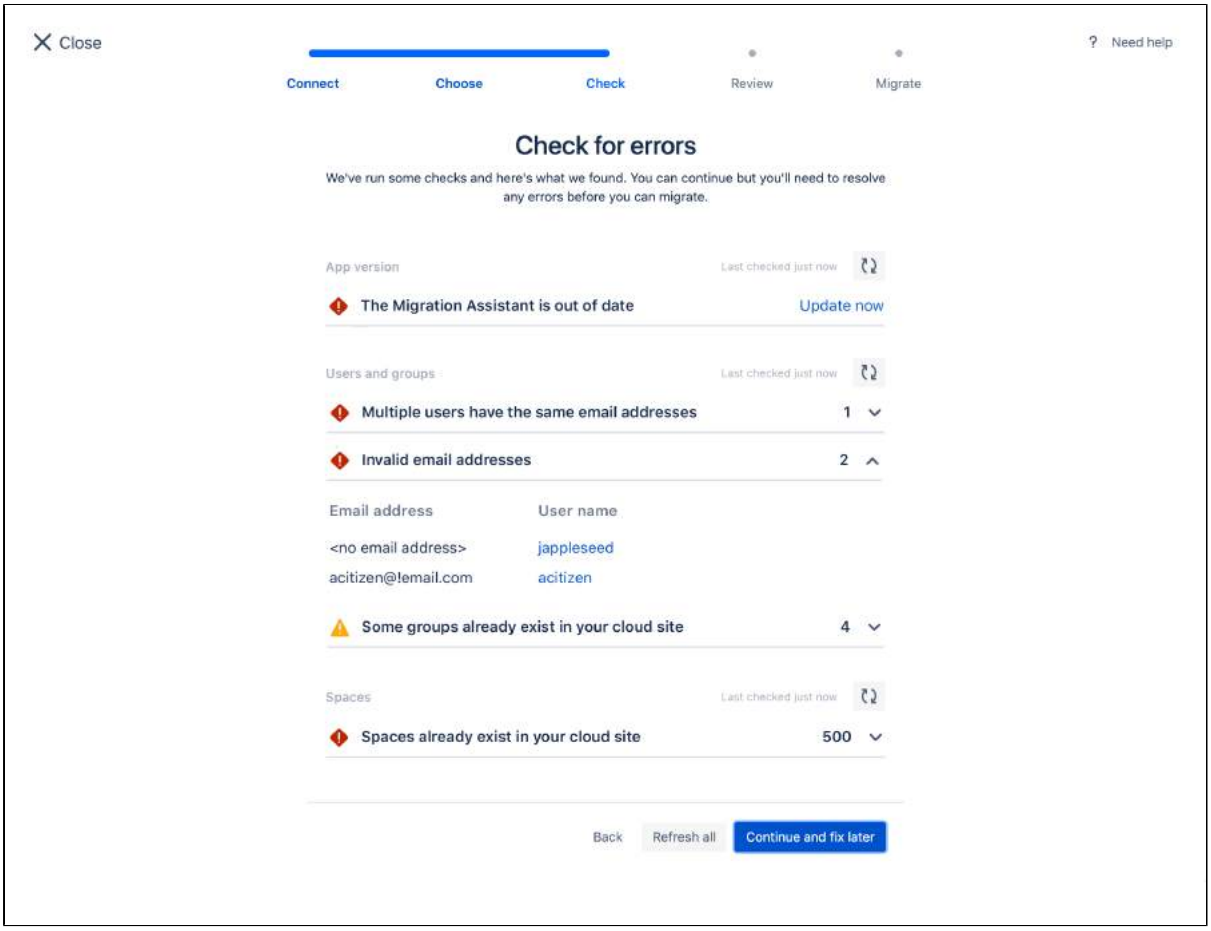

If there is a green tick then the check has passed. If you get a warning sign then you can continue, but you need to be aware of a potential issue.

If a check comes back with a red error then you will need to resolve the error before you can run your migration.

If you decide to **Continue and fix later**, you can come back to view the errors once you have saved your migration.

Updating the app

The migration assistant may be out of date. If you get this error, youll need to update it before running any migrations.

Users and groups errors

All users will need to have a valid and unique email address. If we detect invalid emails or multiple users with the same email, you will get an error. You will need to fix these email addresses before you can run your migration.

If you have chosen to migrate all users, we will check to see if you have any groups with the same name already in your cloud site. If we find groups with the same name, we will merge the users from the server group into the cloud group with the same name. You can continue with your migration without fixing this issue, but its important to check that this wont cause permission escalation.

> The following groups manage admin access and are **blacklisted.** They will not be migrated at all: "site-admins", "system-administrators", "atlassian-addons", "atlassian-addons-admin". Users in these groups will still be migrated; if you want them to be in one of the blacklisted groups youll need to manually add them after migration.

Space errors

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

If youre migrating spaces we will check to see if there will be any space key conflicts. If you get an error you can:

- delete duplicate spaces from your cloud or server sites
- reset your cloud site
- choose not to migrate these spaces by removing them from your migration

You will need to resolve any space key conflicts before you can run your migration.

4. Review your migration

This is the final step in setting up your migration.

If everything looks correct and you want to start your migration, click **Run**. If you would like to start your migration later or you still have errors to fix, click **Save**. If you choose to run your migration, it will still be saved to your dashboard. There, you can view the progress and details of all your migrations.

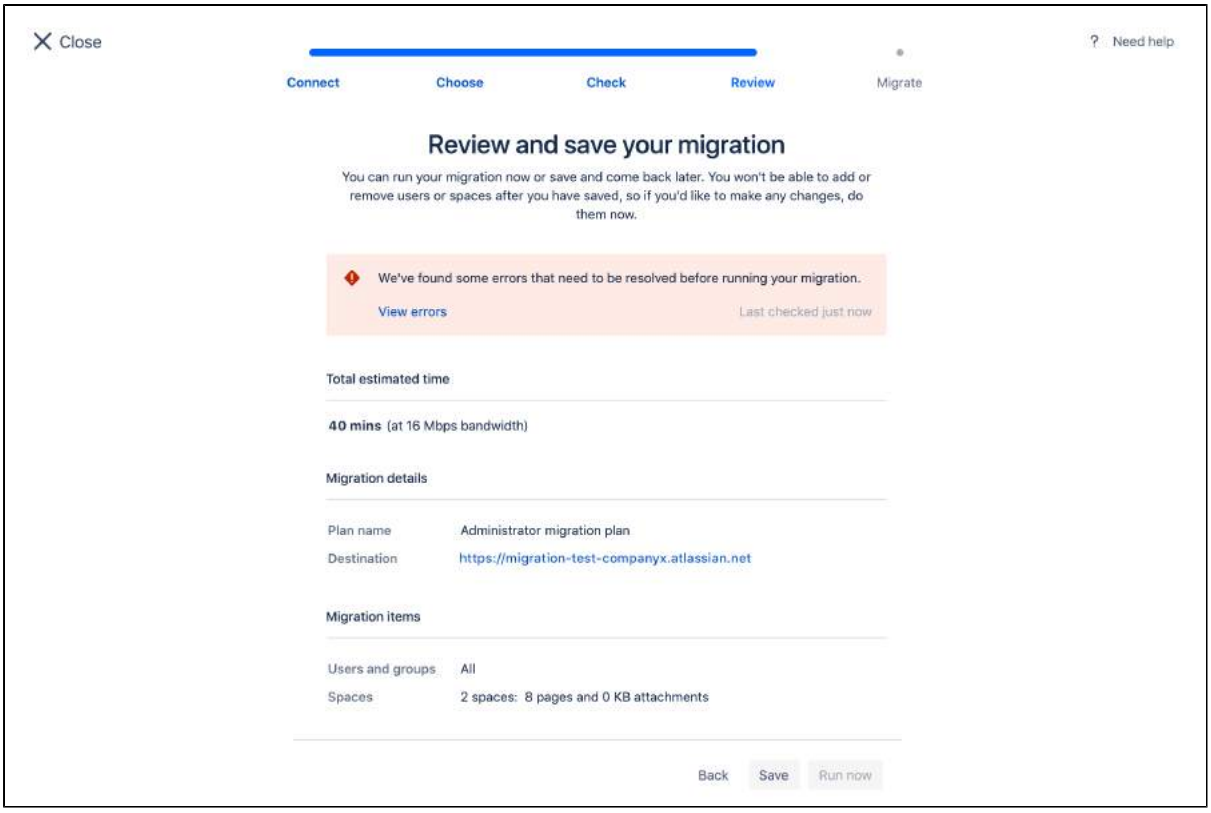

5. Manage your migrations

Your saved migration will be listed on the migration dashboard, where you can **view details** or **run** it. You can also check the status of a migration, monitor the progress, stop a migration that's currently running, or create a new one.

You can create as many migrations as you need. At this time, migrations can't be edited or deleted, so if you create a migration that can't be used, just create a new one.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

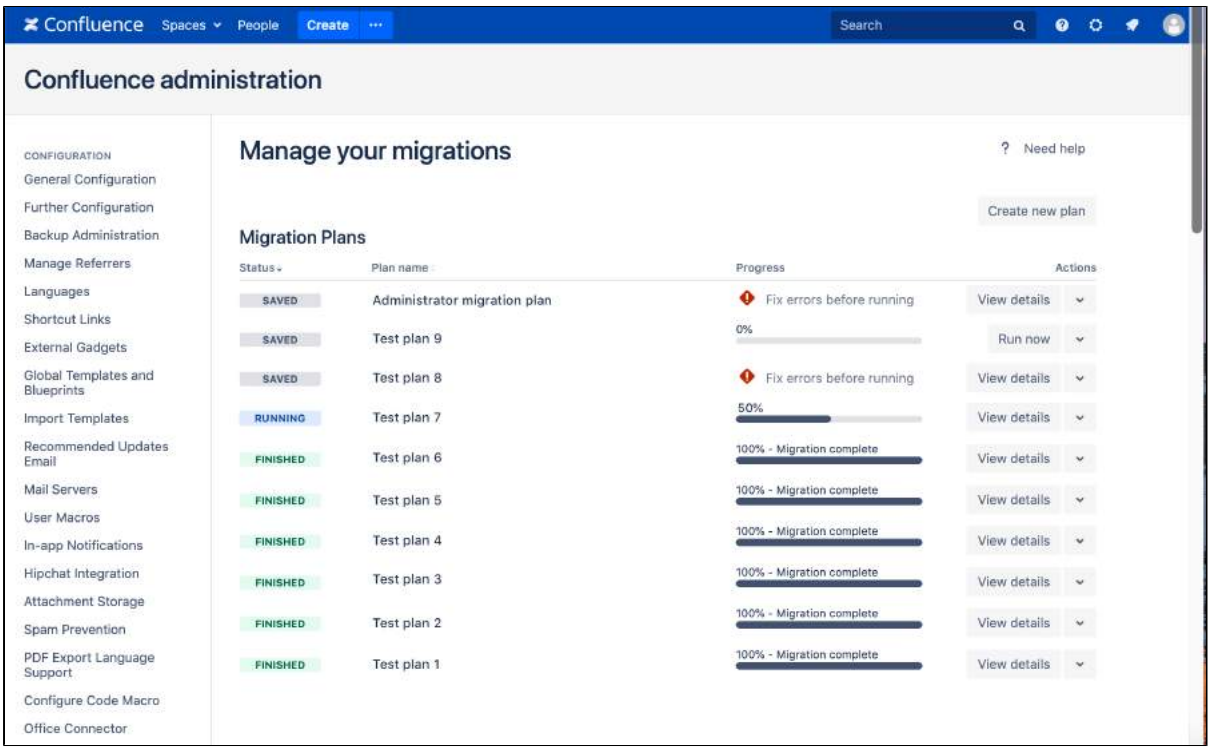

Status definitions

SAVED Your migration is saved and ready to run.

RUNNING Your migration is currently in progress.

FINISHED All tasks in your migration have been completed.

STOPPED Your migration has been manually stopped. Once stopped, it can't be resumed. Any step already in progress will first need to finish before the migration is shown as fully stopped. Some users, groups, and spaces may already have been migrated to your Confluence Cloud site.

FAILED We were unable to complete the migration. This might be because a space key already exists in the destination site, or the migration hit an unexpected error. Some users, groups, and spaces may already have been migrated to your Confluence Cloud site.

After migrating

After migrating spaces, it may take a while for them to appear in the space directory. However, you can still access them via a direct link.

Depending on the type of migration, there may be some things you need to do once your migration is finished.

Users and groups

To make sure your users and groups are set up correctly:

- Review members of groups and approve their permissions by going to **Review imported groups**. (If you have the Free plan, permissions cant be modified; users and groups retain the same permissions that they had on your original site.)
- Add users to the generic groups if necessary. The generic groups are: "site-admins", "system-administrators", "atlassian-addons", "atlassian-addons-admin".
- If you use an external user management system, check that your users have synced correctly.
- When you are ready, invite your users. Go to **Administration > Users > Show details** and then **Resend invite**. When they first log in they may be prompted to set a new password and add personal details.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

> If you have the improved user management experience, go to **Administration > Directory > Users > Show details** and then **Resend invite**.

- We recommend providing some training or sending an onboarding email to your users to help them get familiar with their new cloud workspace.

Spaces

To check that your spaces have migrated successfully:

- Review content and spaces, or ask your users to review their own content.
- Check for any instances of *Former User.* This means that we were unable to match content to a user.
- Link your other Atlassian products by going to **Settings > Application links.**
- Use the Jira macro repair to update any links to Jira. On your cloud site go to **Settings > Jira macro repair** and follow the steps.

> Confluence short links like *https://confluence.example.com/x/PywS* may not work after migrating. Replacing them with internal links (or full URLs if theyre not in your Confluence site)\before migrating should solve this issue.

You can then install any apps you wish to use and onboard your users.

For the full overview of post-migration actions check out the server to cloud migration guide.

More information and support

We have a number of channels available to help you with your migration.

- For more migration planning information and FAQs, visit the Atlassian Cloud Migration Center.
- Have a technical issue or need more support with strategy and best practices? Get in touch.
- Looking for peer advice? Ask the Atlassian Community.

- Want expert guidance? Work with an Atlassian Partner.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Confluence Data Center documentation

Data Center is our self-managed edition of Confluence built for enterprises. It provides the deployment flexibility and administrative control you need to manage mission-critical Confluence sites.Learn more aboutC onfluence Data Centeron our website.

Data Center architecture

You can deploy Confluence Data Center in two ways.

Non-clustered (single node)

Run the Confluence Data Center application on a single server. (Available for Confluence 7.2 and later).

This allows you to take advantage of Data Center-only features without adding to your infrastructure.

Learn more about Data Center features.

Clustered

Run Confluence Data Center in a cluster with multiple application nodes, and a load balancer to direct traffic.

Clustering is designed for large, or mission-critical, Confluence sites, allowing you to provide high availability, and maintain performance as you scale.

Learn more about clustering with Data Center.

Get started

Install or upgrade Confluence Data Center

- Install Confluence Data Center from scratch
- Upgrade from Confluence Server to Data Center

Clustering with Confluence Data Center

- Learn about clustering architecture and requirements
- Set up a Data Center cluster
- Add or remove application nodes
- Turn off clustering (revert to a non-clustered Data Center installation)
- Troubleshoot a clustering issue

Want to see what's included with Data Center? Head to theConfluence Server and Data Center feature comparison.

You can purchase a Data Center license or create an evaluation license at my.atlassian.com

Getting Started with Confluence Data Center

Data Center is our self-managed edition of Confluence built for enterprises. It provides the deployment flexibility and administrative control you need to manage mission-critical Confluence sites.

You can run Confluence Data Center documentation in a cluster, or as standalone (non-clustered) installation.

This guide covers**clustered** Data Center deployments.

On this page:

- 1. Define your requirements
- 2. Provision your infrastructure
- 3. Plan your deployment
- 4. Install and configure Confluence Data Center
- 5. Maintain and scale your Confluence cluster

1. Define your requirements

Before you get started, its a good idea to define your organizations goals and requirements. If you need high availability, scalability, and performance under heavy load, you'll want to run Confluence Data Center in a cluster.

To prepare, we recommend assessing:

- the number of users you have
- the amount of data you have
- your expected usage patterns
- the resources your organization has allocated to maintain your Confluence site.

For more information about disaster recovery for Confluence, head toConfluence Data Center disaster recovery.

Our sizing and performance benchmarks can help you assess your expected load, andpredict performance:

- Confluence Data Center load profiles
- Confluence Data Center performance
- Infrastructure recommendations for enterprise Confluence instances on AWS

2. Provision your infrastructure

Once you've identified your organization's needs, you can prepare your clustered environment.Read ourClustering with Confluence Data Centerfor importanthardware and infrastructure considerations.

To help you get started with clustering, we've provided a Confluence Data Center sample deployment and monitoring strategy.

We've also provided some general advice about node sizing and load balancers, to help you find your feet if this is your first clustered environment:

- Node sizing overview for Atlassian Data Center
- Load balancer configuration options
- Traffic distribution with Atlassian Data Center

3. Plan your deployment

If you're new to Confluence, you can try out Confluence Data Center bydownloading a free trial. This can help you identify dependencies and plan your path to production.

Migrating from Confluence Server to Confluence Data Center?Read throughthese guides to help minimize disruption during the switch:

- Moving to Confluence Data Center
- Atlassian Data Center migration plan
- Atlassian Data Center migration checklist

It's also important to take aninventory of your third-party apps (also known as add-ons) to make sure they're compatible with Data Center. Using a large number of add-ons can degrade performance, soit's a good idea to remove any add-ons that aren't crucial to functionality.

Find out how to evaluate add-ons for Data Center migration.

4. Install and configure Confluence Data Center

Once your environment is ready, it's time to install and configure Confluence Data Center in a cluster.

How you install depends on your environment:

- **Your own hardware** seeInstalling Confluence Data Center
- **Kubernetes** - seeRunning Data Center products on a Kubernetes cluster
- **Azure** seeGetting started with Confluence Data Center on Azure
- **AWS (Amazon Web Services)** seeRunning Confluence Data Center in AWS

If you're migrating from Confluence Server to Confluence Data Center,follow the instructions outlined inUpgrade from Confluence Server to Data Center.

Before deploying Confluence Data Center to production, we recommend thoroughly testing the installation. Head to ourData Center migration plan for detailed advice about testing and launching to production.

5. Maintain and scale your Confluence cluster

Once you've deployed your Confluence Data Center cluster in production, here are some resources for monitoring the health of the cluster, and scaling it up to accomodate moreusers:

- Tools for monitoring your Data Center application

Ready to grow? Read up on scaling and adding nodes to your new Confluence Data Centercluster:

- Scaling with Atlassian Data Center
- Adding or removing Confluence Data Center nodes

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Confluence Server and Data Center feature comparison

Important changes to our server and Data Center products

Weve ended sales for new Server licenses, and will end support for server on February 15, 2024. Were continuing our investment in Data Center with several key improvements. Learn what this means for you

If you manage your own Confluence site (it's not hosted by Atlassian), you'll have either a**Confluence Server** or**Confluence Data Center**license. If we manage Confluence for you, you'll have a **Confluence Cloud** license.

Your Confluence license determines which features and infrastructure choices are available.

We want all teams to get the most out of Confluence, so the core features are available for everyone including creating pages, working together, and organizing your work.

Feature comparison

Heres a summary of available features for Confluence Server and Confluence Data Center. If youre interested in having Atlassian host and manage your products, see how a cloud plan compares on ourConfluence features page.

Core features	Server license	Data Center license
Create spaces Create spaces to store your team or project work.	✓	✓
Create pages Create pages and blog posts, and work on them with your team.	✓	✓
Collaborative editing Up to 12 people can work on the same page at the same time. Learn more	✓	✓
Browser and mobile Use your browser, or use the iOS or Android app. Learn more	✓	✓
Team calendars Create and view calendars from your organization. Learn more	✗	✓ 7.11+
Analytics Track engagement with all the content in your site. Learn more	✗	✓ 7.11+
User management		
External user directories Store users in Active Directory, Crowd, Jira or another LDAP directory. Learn more	✓	✓
Single sign-on Use a SAML or OpenID Connect identity provider for authentication and single-sign on. Learn more	✓ via Crowd	✓ 7.12+
Just-in-time user provisioning Create and update accounts automatically when someone logs in through SAML SSO or OpenID Connect SSO.Learn more	✗	✓ 7.7+

Multiple identity providers Use more than one IdP, and disable login methods you don't want to use (such as basic authentication).Learn more	✗	✓ 7.16+
Advanced permissions management Inspect user and group permissions for auditing and troubleshooting purposes. Learn more	✗	✓ 7.3+
High availability and performance at scale		
Clustering Run Confluence on multiple nodes high availability. Learn more	✗	✓ 5.6+
Content Delivery Network (CDN) support Improve geo-performance for distributed teams. Learn more	✗	✓ 7.0+
Infrastructure and Control		
Read-only mode Limit what users can do in your site while you perform maintenance. Learn more	✗	✓ 6.10+
Sandboxed processes Run resource intensive tasks in external sandboxes for greater stability. Learn more	✗	✓ 6.12+
Rate limiting Control how many external REST API requests users and automations can make. Learn more	✗	✓ 7.3+
Retention rules Set rules to automatically delete historical page versions and purge deleted items from the trash. Learn more	✗	✓ 7.16+
Rolling upgrades Upgrade to the latest bug fix update of the same feature release with no downtime. Learn more	✗	✓ 7.9+
OAuth 2.0 Configure Confluence as an OAuth 2.0 provider, allowing external applications to access Confluence.Learn more	✓ 7.17+	✓ 7.17+
Object storage Scale your data requirements efficiently by configuring Amazon S3 object storage with your instance. This is optional. Learn more	✗	✓ 8.1+
Business intelligence and monitoring		
Advanced auditing Access a wider range of audit events, and integrate with third-party logging systems. Learn more	✗	✓ 7.5+
Data pipeline Export all Confluence data for analysis in your preferred business intelligence platform.Learn more	✗	✓ 7.12+
Application monitoring Use JMX to monitor installed apps and geta deeper insight into how apps affect your instance.Learn more	✓ 7.17+	✓ 7.17+
Deployment options		

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Your own hardware Run Confluence on your own physical servers, virtualized servers, or in the data center of your choice.	✓	✓
AWS Quick Start Use our Cloud Formation Templates to deploy Confluence on AWS. Learn more	✗	✓ 6.1+
Azure template Use our template to deploy Confluence on Azure. Learn more	✗	✓ 6.6+
Kubernetes Helm charts Use our Helm charts to deploy Confluence on Kubernetes.Learn more	✗	✓ 7.13+

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Clustering with Confluence Data Center

Confluence Data Center allows you to run a cluster of multiple Confluence nodes, providing high availability, scalable capacity, and performance at scale.

This guide describes the benefits of clustering, and provides you an overview of what youll need to run Confluence in a clustered environment, including infrastructure and hardware requirements.

Ready to get started? See Set up a Confluence Data Center cluster

On this page

- Is clustering right for my organization?
- Clustering architecture
- Infrastructure and hardware requirements
- App compatibility
- Ready to get started?

Is clustering right for my organization?

Clustering is designed for enterprises with large or mission-critical Data Center deployments that require continuous uptime, instant scalability, and performance under high load.

There are a number of benefits to running Confluence in a cluster:

- **High availability and failover:** If one node in your cluster goes down, the others take on the load, ensuring your users have uninterrupted access to Confluence.
- **Performance at scale:** each node added to your cluster increases concurrent user capacity, and improves response time as user activity grows.
- **Instant scalability:** add new nodes to your cluster without downtime or additional licensing fees. Indexes and apps are automatically synced.
- **Disaster recovery:** deploy an offsite Disaster Recovery system for business continuity, even in the event of a complete system outage. Shared application indexes get you back up and running quickly.
- **Rolling upgrade**: upgrade to the latest bug fix update of your feature release without any downtime. Apply critical bug fixes and security updates to your site while providing users with uninterrupted access to Confluence.

Clustering is only available with a Data Center license. Learn more about upgrading to Data Center.

Clustering architecture

The basics

A Confluence Data Center cluster consists of:

- Multiple identical application nodes running Confluence Data Center.
- A load balancer to distribute traffic to all of your application nodes.
- A sharedfilesystem that stores attachments, and other shared files.
- A database that all nodes read and write to.

All application nodes are active and process requests. A user will access the same Confluence node for all requests until their session times out, they log out, or a node is removed from the cluster.

The image below shows a typical configuration:

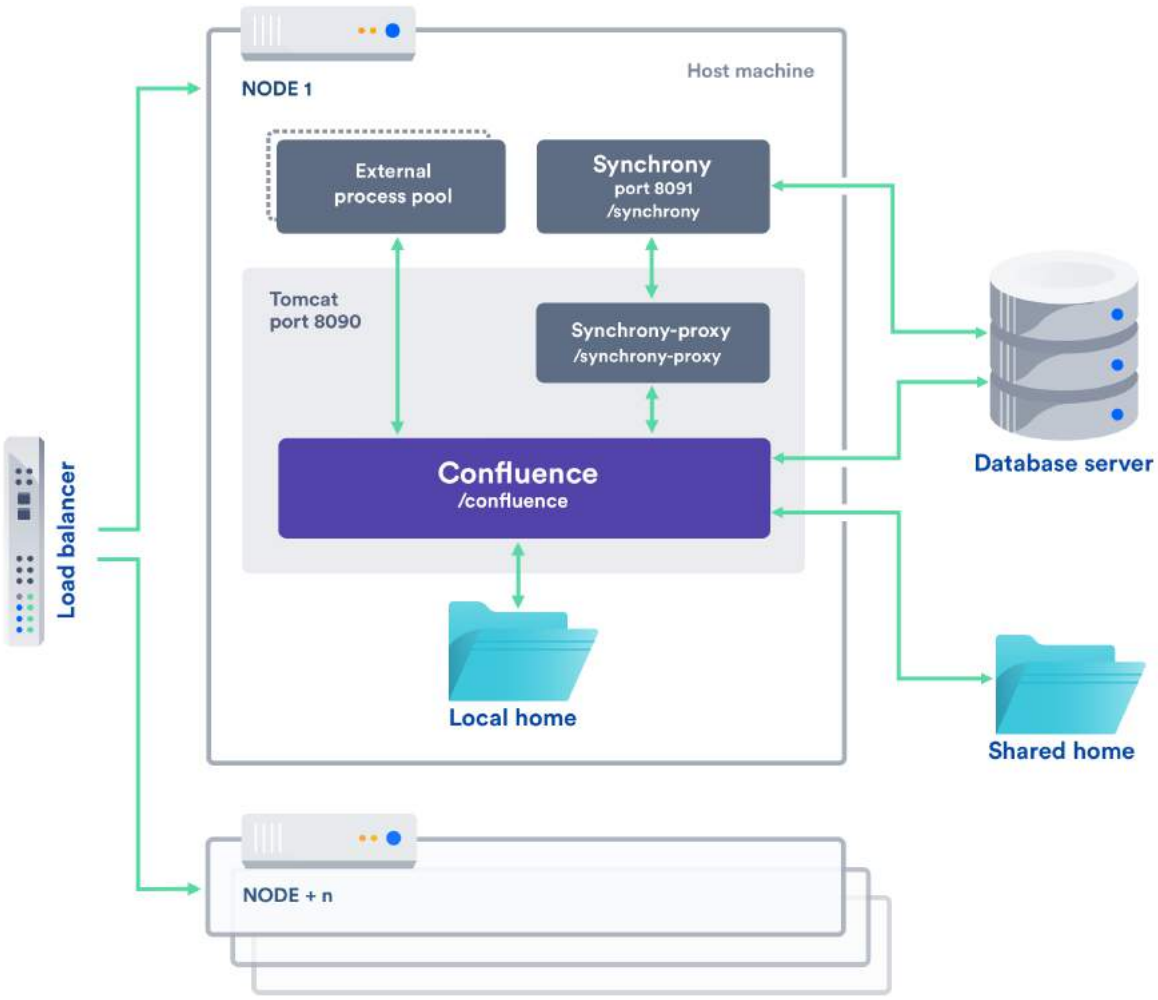

Licensing

Your Data Center license is based on the number of users in your cluster, rather than the number of nodes. This means you can scale your environment without additional licensing fees for new servers or CPU.

You can monitor the available license seats in the License Details page in the admin console.

If you wanted to automate this process (for example to send alerts when you are nearing full allocation) you can use the REST API.

The following GET requests require an authenticated user with system administrator permissions. The requests return JSON.

`<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

Your Confluence license determines which features and infrastructure choices are available. Head toConfluence Server and Data Center feature comparison for a full run down of the differences between a Server license and a Data Center license.

Home directories

To run Confluence in a cluster, you'll need an additional home directory, known as the shared home.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Each Confluence node has a local home that contains logs, caches, Lucene indexes and configuration files. Everything else is stored in the shared home, which is accessible to each Confluence node in the cluster. Marketplace apps can choose whether to store data in the local or shared home, depending on the needs of the app.

Here's a summary of what is found in the local home and shared home:

Local home	Shared home
• logs • caches • Lucene indexes • configuration files • plugins	• attachments • avatars / profile pictures • icons • export files • import files • plugins

If you are currently storing attachments in your database you can continue to do so, but this is not available for new installations.

Caching

When clustered, Confluence uses a combination of local caches, distributed caches, and hybrid caches that are managed using Hazelcast. This allows for better horizontal scalability, and requires less storage and processing power than using only fully replicated caches. SeeCache Statistics for more information.

Because of this caching solution, to minimize latency, your nodes should be located in the same physical location, or region (for AWS and Azure).

Indexes

Each individual Confluence application node stores its own full copy of the index. A journal service keeps each index in sync.

When you first set up your cluster, you will copy the local home directory, including the indexes, from the first node to each new node.

When adding a new Confluence node to an existing cluster, you will copy the local home directory of an existing node to the new node. When you start the new node, Confluence will check if the index is current, and if not, request a recovery snapshot of the index from either the shared home directory, or a running node (with a matching build number) and extract it into the index directory before continuing the start up process. If the snapshot can't be generated or is not received by the new node in time, existing index files will be removed, and Confluence will perform a full re-index.

If a Confluence node is disconnected from the cluster for a short amount of time (hours), it will be able to use the journal service to bring its copy of the index up-to-date when it rejoins the cluster. If a node is down for a significant amount of time (days) its Lucene index will have become stale, and it will request a recovery snapshot from an existing node as part of the node startup process.

If you suspect there is a problem with the index, you can rebuild the index on one node, and Confluence will propagate the new index files to each node in the cluster.

SeeContent Index Administration for more information on reindexing and index recovery.

Cluster safety mechanism

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

The ClusterSafetyJob scheduled task runs every 30 seconds in Confluence. In a cluster, this job is run on one Confluence node only. The scheduled task operates on asafety number a randomly generated number that is stored both in the database and in the distributed cache used across the cluster.The ClusterSafetyJob compares the value in the database with the one in the cache, and if the value differs, Confluence will shut the node down - this is known as cluster split-brain.This safety mechanism is used to ensure your cluster nodes cannot get into an inconsistent state.

If cluster split-brain does occur, you need to ensure proper network connectivity between the clustered nodes. Most likely multicast traffic is being blocked or not routed correctly.

Balancing uptime and data integrity

By changing how often the cluster safety scheduled job runs and the duration of the Hazelcast heartbeat (which controls how long a node can be out of communication before it's removed from the cluster) you can fine tune the balance between uptime and data integrity in your cluster.In most cases the default values will be appropriate, but there are some circumstances where you may decide to trade off data integrity for increased uptime for example.

Uptime over data integrity

Cluster safety job	Hazelcast heartbeat	Effect
1 minute	1 minute	You could have network interruptions or garbage collection pauses of up to 1 minute without triggering a cluster panic. However, if two nodes are no longer communicating, conflicting data could be being written to the database for up to 1 minute, affecting your data integrity.
10 minutes	30 seconds	You could have network interruptions or garbage collection pauses of up to 30 seconds without nodes being evicted from the cluster. Evicted nodes then have up to 10 minutes to rejoin the cluster before the Cluster Safety Job kicks in and shuts down the problem node. Although this may result in higher uptime for your site, conflicting data could be being written to the database for up to 10 minutes, affecting your data integrity.

Data integrity over uptime

Cluster safety job	Hazelcast heartbeat	Effect
15 seconds	15 seconds	Network interruptions or garbage collection pauses longer than 15 seconds will trigger a cluster panic. Although this may result in higher downtime for your site, nodes can only write to the database while out of communication with each other for a maximum of 15 seconds, ensuring greater data integrity.
15 seconds	1 minute	You could have network interruption or garbage collection pauses up to 1 minute without nodes being evicted from the cluster. Once a node is evicted, it can only write to the database for a maximum of 15 seconds, minimizing the impact on your data integrity.

To find out how to change the cluster safety scheduled job, see Scheduled Jobs.

You can change theHazelcast heartbeat default via the `confluence.cluster.hazelcast.max.no.heartbeat.seconds` system property. See Configuring System Properties.

Cluster locks and event handling

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Where an action must only run on one node, for example a scheduled job or sending daily email notifications, Confluence uses a cluster lock to ensure the action is only performed on one node.

Similarly, some actions need to be performed on one node, and then published to others. Event handling ensures that Confluence only publishes cluster events when the current transaction is committed and complete. This is to ensure that any data stored in the database will be available to other instances in the cluster when the event is received and processed.Event broadcasting is done only for certain events, like enabling or disabling an app.

Cluster node discovery

When configuring your cluster nodes you can either supply theIP address of each cluster node, or a multicast address.

If you're using multicast:

Confluence will broadcast a join request on the multicast network address. Confluence must be able to open a UDP port on this multicast address, or it won't be able to find the other cluster nodes.Once the nodes are discovered, each responds with a unicast (normal) IP address and port where it can be contacted for cache updates. Confluence must be able to open a UDP port for regular communication with the other nodes.

A multicast address can be auto-generated from the cluster name, or you can enter your own, during the set-up of the first node.

Infrastructure and hardware requirements

The choice of hardware and infrastructure is up to you. Below are some areas to think about when planning your hardware and infrastructure requirements.

AWS Quick Start deployment option

If you plan to run Confluence Data Center on AWS, a Quick Start is available to help you deploy Confluence Data Center in a new or existing Virtual Private Cloud (VPC). You'll get your Confluence and Synchrony nodes, Amazon RDS PostgreSQL database and application load balancer all configured and ready to use in minutes. If you're new to AWS, the step-by-step Quick Start Guide will assist you through the whole process.

Confluence can only be deployed in a region that supports Amazon Elastic File System (EFS). SeeRunning Confluence Data Center in AWS for more information.

It is worth noting that if you deploy Confluence using the Quick Start, it will use the Java Runtime Engine (JRE) that is bundled with Confluence (/opt/atlassian/confluence/jre/), and not the JRE that is installed on the EC2 instances (/usr/lib/jvm/jre/).

Server requirements

You should not run additional applications (other than core operating system services) on the same servers as Confluence. Running Confluence, Jira and Bamboo on a dedicated Atlassian software server works well for small installations but is discouraged when running at scale.

Confluence Data Center can be run successfully on virtual machines. If you plan to use multicast, you can't run Confluence Data Center in Amazon Web Services (AWS) environments as AWS doesn't support multicast traffic.

Cluster nodes

Each node does not need to be identical, but for consistent performance we recommend they are as close as possible. All cluster nodes must:

- be located in the same data center, or region (for AWS and Azure)
- run the same Confluence version on each Confluence node (exceptduring a rolling upgrade)
- run the same Synchrony version on each Synchrony node (if not using managed Synchrony)
- have the same OS, Java and application server version

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

- have the same memory configuration (both the JVM and the physical memory) (recommended)
- be configured with the same time zone (and keep the current time synchronized). Using ntpd or a similar service is a good way to ensure this.

⚠️You must ensure the clocks on your nodes don't diverge, as it can result in a range ofproblems with yourcluster.

How many nodes?

Your Data Center license does not restrict the number of nodes in your cluster. The right number of nodes depends on the size and shape of your Confluence site, and the size of your nodes. See our Confluence Data Center load profiles guide for help sizing your instance. In general, we recommend starting small and growing as you need.

Memory requirements

Confluence nodes

We recommend that each Confluence node has a minimum of 10GB of RAM. A high number of concurrent users means that a lot of RAM will be consumed.

Here's some examples of how memory may be allocated on different sized machines:

RAM	Breakdown for each Confluence node
10GB	• 2GB for operating system and utilities • 4GB for Confluence JVM (-Xmx 3GB) • 2GB for external process pool (2 sandboxes with -Xmx 512MB each) • 2GB for Synchrony
16GB	• 2GB for operating system and utilities • 10GB for Confluence JVM (-Xmx 8GB) • 2GB for external process pool (2 sandboxes with -Xmx 512MB each) • 2GB for Synchrony

The maximum heap (-Xmx) for the Confluence application is set in the `setenv.sh`or `setenv.bat`file. The default should be increased for Data Center. We recommend keeping the minimum (Xms) and maximum (Xmx) heap the same value.

The external process pool is used to externalise memory intensive tasks,to minimise the impact on individual Confluence nodes. The processes are managed by Confluence. The maximum heap for each process (sandbox) (-Xmx), and number of processes in the pool, is set using system properties.In most cases the default settings will be adequate, and you don't need to do anything.

Standalone Synchrony cluster nodes

Synchrony is required for collaborative editing. By default, it is managed by Confluence, but you can choose to run Synchrony in its own cluster. SeePossible Confluence and Synchrony Configurations for more information on the choices available.

If you do choose to run your own Synchrony cluster, we recommend allowing 2GB memory for standalone Synchrony.Here's an example of how memory could be allocated on a dedicated Synchrony node.

Physical RAM	Breakdown for each Synchrony node
4GB	• 2GB for operating system and utilities • 2GB for Synchrony JVM (-Xmx 1GB)

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Database

The most important requirement for the cluster database is that it have sufficient connections available to support the number of nodes.

For example, if:

- each Confluence node has a maximum pool size of 20 connections
- each Synchrony node has a maximum pool sizeof 15 connections (the default)
- you plan to run 3 Confluence nodes and 3 Synchrony nodes

your database server must allow at least 105 connections to the Confluence database. In practice, you may require more than the minimum for debugging or administrative purposes.

You should also ensure your intended database is listed in the currentSupported Platforms.The load on an average cluster solution is higher than on a standalone installation, so it is crucial to use the a supported database.

You must also use a supported database driver. Collaborative editing will fail with an error if you're using an unsupported or custom JDBC driver (or`driverClassName`in the case of a JNDI datasource connection). SeeDatabase JDBC Driversfor the list of drivers we support.

Additional requirements for database high availability

Running Confluence Data Center in a cluster removes the application server as a single point of failure. You can also do this for the database through the following supported configurations:

- Amazon RDS Multi-AZ: this database setup features a primary database that replicates to a standby in a different availability zone. If the primary goes down, the standby takes its place.
- Amazon PostgreSQL-Compatible Aurora: this is a cluster featuring a database node replicating to one or more readers (preferably in a different availability zone). If the writer goes down, Aurora will promote one of the writers to take its place.

The**AWS Quick Start deployment option**allows you to deploy Confluence Data Center with either one, from scratch. If you want to set up an Amazon Aurora cluster with an existing Confluence Data Center instance, refer toConfiguring Confluence Data Center to work with Amazon Aurora.

Shared home directory and storage requirements

All Confluence cluster nodes must have access to a shared directory in the same path.NFS and SMB/CIFS shares are supported as the locations of the shared directory. As this directory will contain large amount of data (including attachments and backups) it should be generously sized, and you should have a plan for how to increase the available disk space when required.

Remember me and session timeout

The 'remember me' option is enforced by default in a cluster. Users won't see the 'remember me' checkbox on the login page, and their session will be shared between nodes. See the following knowledge base articles if you need to change this, or change the session timeout.

- How to configure the 'Remember Me' feature in Confluence
- How to adjust the session timeout for Confluence

Load balancers

We suggest using the load balancer you are most familiar with. The loadbalancer needs to support session affinity and WebSockets. This is required for both Confluence and Synchrony. If you're deploying on AWS you'll need to use an Application Load Balancer (ALB).

Here are some recommendations when configuring your load balancer:

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

- Queue requests at the load balancer. By making sure the maximum number requests served to a node does not exceed the total number of http threads that Tomcat can accept, you can avoid overwhelming a node with more requests than it can handle. You can check the maxThreads in `<install-directory>/conf/server.xml`.
- Don't replay failed idempotent requests on other nodes, as this can propagate problems across all your nodes very quickly.
- Using *least connections* as the load balancing method, rather than *round robin*, can better balance the load when a node joins the cluster or rejoins after being removed.

Many load balancers require a URL to constantly check the health of their backends in order to automatically remove them from the pool. It's important to use a stable and fast URL for this, but lightweight enough to not consume unnecessary resources. The following URL returns Confluence's status and can be used for this purpose.

URL	Expected content	Expected HTTP Status
http://<confluenceurl>/status	`{"state":"RUNNING"}`	200 OK

HTTP Status Code	Response entity	Description
200	`{"state":"RUNNING"}`	Running normally
500	`{"state":"ERROR"}`	An error state
503	`{"state":"STARTING"}`	Application is starting
503	`{"state":"STOPPING"}`	Application is stopping
200	`{"state":"FIRST_RUN"}`	Application is running for the first time and has not yet been configured
404		Application failed to start up in an unexpected way (the web application failed to deploy)

Here are some recommendations, when setting up monitoring, that can help a node survive small problems, such as a long GC pause:

- Wait for two consecutive failures before removing a node.
- Allow existing connections to the node to finish, for say 30 seconds, before the node is removed from the pool.

Network adapters

Use separate network adapters for communication between servers.Cluster nodes should have a separate physical network (i.e. separate NICs) for inter-server communication.This is the best way to get the cluster to run fast and reliably. Performance problems are likely to occur if you connect cluster nodes via a network that has lots of other data streaming through it.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Additional requirements for collaborative editing

Collaborative editing in Confluence 6.0 and later is powered by Synchrony, which runs as a seperate process.

If you have a Confluence Data Center license, two methods are available for running Synchrony:

- **managed by Confluence**(recommended)
  Confluence will automatically launch a Synchrony process on the same node, and manage it for you. No manual setup is required.
- **Standalone Synchrony cluster (managed by you)**
  You deploy and manage Synchrony standalone in its own cluster with as many nodes as you need. Significant setup is required.During a rolling upgrade, you'll need to upgrade the Synchrony separately from the Confluence cluster.

If you want simple setup and maintenance, we recommend allowing Confluence to manage Synchrony for you. If you want full control, or if making sure the editor is highly available is essential, then managing Synchrony in its own cluster may be the right solution for your organisation.

App compatibility

The process for installing Marketplace apps (also known as add-ons or plugins) in a Confluence cluster is the same as for a standalone installation. You will not need to stop the cluster, or bring down any nodes to install or update an app.

The Atlassian Marketplace indicates apps that are compatible with Confluence Data Center.

If you have developed your own plugins (apps) for Confluence you should refer to our developer documentation onHow do I ensure my app works properly in a cluster?to find out how you can confirm your app is cluster compatible.

Ready to get started?

Head to Set up a Confluence Data Center cluster for a step-by-step guide to enabling and configuring your cluster.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

External Process Pool for Confluence Data Center

In Confluence Data Center we minimize the impact of particularly memory or CPU intensive actions byhandling them in an external process pool, which is a seperate pool of processes, managed by Confluence. These processes (also known as sandboxes) can crash or be terminated, and will be restarted automatically by Confluence, without affecting the Confluence application itself.

On this page:

- Memory requirements
- Configure the external process pool
- Monitor failed actions

The external process pool currently handles the following actions:

- Document conversion(thumbnail generation for file previews)
- Exporting a space to PDF

> The external process pool is **only available for Confluence Data Center**.
>
> In Confluence Server, these actions are handled by Confluence, so the information on this page does not apply.

Memory requirements

You will need to make sure that Confluence has enough memory for the external process pool. In a clustered Data Center installation, you'll need to do this for each cluster node.The pool contains two processes (sandboxes) by default, sowe recommend allowing an additional 2 GB on top of what is already required for Confluence (1 GB per sandbox).

If you increase the size of the external process pool, make sure each node has enough free memory to cater for the extra processes.

Configure the external process pool

In most cases the default values will be adequate, however system administrators can configure the external process pool usingsystem properties. For example you may want to increase the size of the pool (the number of processes available), or increase the amount of memory a process can consume. Here are the main properties you may need to change:

- **conversion.sandbox.pool.size**
  Use this property to increase the number of processes (sandboxes) in the pool. You'll need to allow additional memory on each node for each additional process.
- **conversion.sandbox.memory.limit.megabytes**
  Use this property to limit the amount of memory each process (sandbox) in the pool can consume.

SeeRecognized System Propertiesfor a full description of these properties, including additional properties that can be used to fine-tune, or disable sandboxes for particular actions.

Monitor failed actions

When an external process (sandbox) is terminated, we'll write the following to the application log on that node:

```
2018-04-09 17:35:35 WARN [sandbox-terminator]
[impl.util.sandbox.DefaultSandbox] lambda$startTerminator$0 Request
has taken 33384ms exceeds limit 30000ms terminating sandbox
```

This will be followed by an`Attempting to restart the sandbox`message, the next time someone performs an action that uses the external process pool.

Note that the process is not immediatley restarted after termination, as we don't re-attempt failed actions. We wait for the next request to spin up a new sandbox process.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Document conversion for Confluence Data Center

When you insert a file into a page (for example a Word document, or Excel spreadsheet), Confluence will convert the contents to a format that can be viewed inline in the page, in the preview, or in some macros. This can be quite memory and CPU intensive, and has been known to cause out of memory errors when processing very complex files.

In Confluence Data Center we minimize the impact byhandling the conversion in an external process pool, which is a seperate pool of processes, managed by Confluence. These processes (also known as sandboxes) can crash or be terminated, and will be restarted automatically by Confluence, without affecting the Confluence application itself.

For example, If you insert a very complex file, and the process crashes or is terminated, thumbnail generation will fail. When this happens, a placeholder thumbnail will be used on the page, and a download option will be provided in the file preview. Confluence Data Center doesn't re-attempt to generate thumbnails for failed files. A good example of a complex file, is a PowerPoint presentation that contains 50 embedded Excel charts. Most files will be processed without any problems.

The external process pool is used for the following conversions:

- thumbnail generation for images and documents inserted into a page, or viewed in the preview.
- HTML conversion for Word and Office documents viewed using the Office Word and Office Excel macros.

> The external process pool is **only available for Confluence Data Center**.
>
> In Confluence Server, thumbnail generation and HTML conversion is handled by Confluence, so the information on this page does not apply.

Configure the external process pool

In most cases the default values will be adequate, however system administrators can change the behaviour using system properties. For example you may want to increase the size of the pool (the number of processes available), or increase the time limit before a process is terminated. Here are the main properties you may need to change:

- **conversion.sandbox.pool.size**
  Use this property to increase the number of processes (sandboxes) in the pool. You'll need to allow additional memory on each node for each additional process.
- **conversion.sandbox.memory.limit.megabytes**
  Use this property to limit the amount of memory each thumbnail generation process in the pool can consume.
- **document.conversion.sandbox.memory.requirement.megabytes**
  Use this property to limit the amount of memory each HTML conversion process in the pool can consume.
- **document.conversion.sandbox.request.time.limit.secs**
  Use this property to change the amount of time (in seconds) that the sandbox will wait for the conversion process to complete, before terminating the process.

SeeRecognized System Propertiesfor a full description of these properties, plus a few additional properties that can be used to fine-tune, or disable the sandboxes completley.

Re-attempt thumbnail generation for failed files

Confluence does not re-attempt to generate thumbnails for a failed attachment, and re-inserting the attached file into the editor will not trigger the process.

If you do want to re-attempt thumbnail generation, for example after increasing the request time limit, you will need to re-upload the file, and then re-insert it into the page.

Other system properties that affect document conversion

The system properties listed on this page apply specifically to the external process pool. However there are some additional properties that apply in both Confluence Server and Confluence Data Center:

- **confluence.document.conversion.imaging.enabled.tif**
  Use this property to enable document conversion for TIFF files. This is disabled by default.
- **confluence.document.conversion.imaging.enabled.psd**
  Use this property to enable document conversion for Photoshop PSD files. This is disabled by default.
- **confluence.document.conversion.imaging.convert.timeout**
  Use this property to change the default 30 second time limit which applies when performing document conversion on complex image files (such as ICO, EMF, WMF).
- **confluence.document.conversion.slides.convert.timeout**
  Use this property to change the default 30 second time limit which applies when performing document conversion on presentation files (such as PPT, PPTX).

To override the default value of these properties, you'll need to use the `conversion.sandbox.java.options`system property to pass the property to the JVMs that make up the external process pool.

In this example we'll enable thumbnail generation for TIFF and PSD files.

1. Edit the `<install-directory>/bin/setenv.bat` file.
2. Add the following lines

```
set CATALINA_OPTS=-Dconversion.sandbox.java.options=-Dconfluence.document.conversion.imaging.
enabled.tif=true -Dconfluence.document.conversion.imaging.enabled.psd=true %CATALINA_OPTS%
```

You can pass multiple properties to the external process pool JVMs this way.

If you're running Confluence as a Windows Service or on AWS, see Configuring System Propertiesfor how to add this property.

In this example we'll enable thumbnail generation for TIFF and PSD files.

1. Edit the `<install-directory>/bin/setenv.sh` file.
2. Add the following lines. In this example we're enabling document conversion for TIFF and PSD files.

```
CATALINA_OPTS="-Dconversion.sandbox.java.options=-Dconfluence.document.conversion.imaging.enabled.
tif=true -Dconfluence.document.conversion.imaging.enabled.psd=true ${CATALINA_OPTS}"
```

You can pass multiple properties to the external process pool JVMs this way.

If you're running Confluence on AWS, see Configuring System Propertiesfor how to add this property.

If you decide to increase the timeout for generating thumbnails in the external process pool using the`document.conversion.sandbox.request.time.limit.secs`system property,you may also want to change the timeout for complex image files or presentations using the system properties above. Alternatively you could keep the default, and allow these types of files to fail sooner.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

PDF export in Confluence Data Center

When you export a space to PDF, Confluence exports the content of each page to HTML, converts that HTML to PDF, and then finally merges all the pages together into a single PDF file. This can be quite memory and CPU intensive, andhas been known to cause out of memory errors when processing spaces with very long or complex pages.

In Confluence Data Center we minimize the impact byhandling the export in an external process pool, which is a seperate pool of processes, managed by Confluence. These processes (also known as sandboxes) can crash or be terminated, and will be restarted automatically by Confluence, without affecting the Confluence application itself.

> The external process pool is **only available for Confluence Data Center**.
>
> In Confluence Server, PDF export is handled by Confluence, so the information on this page does not apply.

Troubleshooting failed exports

Exporting an entire space to PDF can sometimes fail, especially if the space is very large, or has very long or complex pages. If PDF export fails you'll see one of the following errors in your browser.

Page took too long to convert

This error occurs when the time it takes to convert the HTML of a page to PDF exceeds the set time limit. The page title will be included in the error message.

You should take a look at the page, and see if it can be simplified. It might have a lot of complex macros, or a lot of web images (images that are not attached to the page).If this error happens a lot, you can ask your admin toin crease the time limit.

Error converting page to HTML

This error occurs when Confluence runs out of memory, or hits another error while trying to convert the HTML of a page to PDF. The page title will be included in the error message.

As with the 'page took too long to convert' error above, you should take a look at the page, and see if it can be simplified.

Confluence admins can get more information about the cause of these errors from the Confluence application logs. If the failures are being caused by out of memory errors,your admin may be able to increase the amount of memory available to each sandbox in the external process pool. SeeExternal Process Pool for Confluence Data Center for more information.

Final PDF file wasn't merged in time

This error occurs at the last stage of the process, when the time it took to stitch together all the individual page PDFs into one PDF file, exceeds the set time limit.

If you hit this error you could try exporting the space again, or perhaps export the space in two sections (using the custom option on the PDF export screen). If this error happens a lot, you can ask your admin to increase the time limit.

Error merging the final PDF file

This error occurs when Confluence runs out of memory, or hits another error, when attempting to stitch together all the individual page PDFs into one file.

If you hit this error you could try exporting the space again, or perhaps export the space in two sections (using the custom option on the PDF export screen).

Confluence admins can get more information about the cause of these errors from the Confluenceapplication logs. If the failures are being caused by out of memory errors,they may be able to increase the amount of memory available to each sandbox in the external process pool. SeeExternal Process Pool for Confluence Data Centerfor more information.

Too many concurrent exports

This error occurs when multiple people are exporting to PDF at the same time. Confluence limits the number of PDF exports that can be processed concurrently.

If you hit this error, try exporting the space again later, after the other PDF exports have completed.

If this error happens a lot, your admin can increase the maximum number of concurrent PDF exports, or increase the time Confluence should wait when the maximum number of concurrent PDF exports has been reached using the following system properties:

confluence.pdfexport.permits.size
Use this property to set the maximum number of concurrent PDF exports that can be performed. This property applies per node, not per sandbox process.

confluence.pdfexport.timeout.seconds
Use this property to set the amount of time a new PDF export request should wait before failing, if the maximum number of concurrent PDF exportshas already been reached.

Change the time limit

Processes are automatically terminated once a time limit is exceeded. You can increase the time limit for PDF export using the following system property:

pdf.export.sandbox.request.time.limit.secs
Use this property to set the amount of time (in seconds) that a process should wait to complete, before being terminated. This time limit applies both to the time to convert the content from HTML to PDF, and the time to merge the final PDF file.

SeeRecognized System Propertiesfor a full list of properties, including a few additional properties that can be used to fine-tune, or disable the sandboxes for a particular action.

Don't use the external process pool for PDF export

If you don't want to use the external process pool for PDF export, you can revert back to the Confluence Server method of generating PDF exports using the following system property:

pdf.export.sandbox.disable
Set this property to true if you don't want to handle PDF exports in the external process pool.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Restricted Functions in Confluence Data Center

There are some features that are disabled or limited in clustered Confluence Data Center installations. This is to ensure the integrity and performance of your cluster.

The current restricted functions are:

Restricted function	Data Center Status	Explanation
Workbox plugins	Available from 5.7	The workbox provides notifications collected from Confluence page watches, shares, and mentions. This is disabled in Confluence Data Center 5.6 to ensure notifications are correctly handled across the cluster. Disabled plugins included Workbox common plugin, Workbox Jira provider plugin, Workbox confluence provider plugin, Workbox host plugin. You will not be able to enable these plugins in the universal plugin manager.
Confluence Quick Reload Plugin	Available from 5.6.3	The quick reload function notifies users when a new comment has been added to a page they are currently viewing. This is disabled in Confluence Data Center 5.6 and 5.6.1 for performance reasons. You will not be able to enable the Confluence Quick Reload Plugin in the universal plugin manager. See ~~CONFSERVER-34680~~ - Make quick reload plugin available in Confluence Data Center CLOSED for more info.
Application links 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 ConfluenceUsage Stats plugin providesspace 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.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Set up a Confluence Data Center cluster

Confluence Data Center allows you to run a cluster of multiple Confluence nodes, providing high availability, scalable capacity, and performance at scale.

This guides walks you through the process of configuring a Data Center cluster on your own infrastructure.

Youll need to be logged in as a System Administrator to do this.

Not sure if clustering is right for you? Check out Clustering with Confluence Data Center for a detailed overview.

On this page

- Clustering with AWS and Azure
- Before you begin
- Set up and configure your cluster
- Add more Confluence nodes
- Troubleshooting
- Were here to help

Clustering with AWS and Azure

You can also choose to deploy a Data Center cluster on public cloud providers, like AWS (Amazon Web Services) and Azure. We have specific guides and deployment templates to help you easily configure a cluster in AWS or Azure. Check them out to find out what's required.

Before you begin

Clustering requirements

To use Confluence Data Center you must:

- Have a Data Center license (you canpurchase a Data Center licenseo r create an evaluation license atmy.atlassian.com)
- Use asupportedexternal database, operating system and Java version
- Use OAuth authentication if you haveapplication linksto other Atlassian products (such as Jira)

To run Confluence in a cluster you must also:

- Use a load balancer with session affinity in front of the Confluence cluster. WebSockets support is also recommended for collaborative editing.
- Have a shared directory accessible to all cluster nodes in the same path (this will be your shared home directory). This must be a separate directory, and not located within the local home or install directory.

See Clustering with Confluence Data Center for a complete overview of hardware and infrastructure considerations.

Security

Ensure that only permittedcluster nodes are allowed to connect to the following portsthrough the use of a firewall and / or network segregation:

- 5801 -Hazelcastport for Confluence
- 5701 -Hazelcastport for Synchrony
- 25500 - Cluster base port for Synchrony

If you use multicast for cluster discovery:

- 54327- Multicast port for Synchrony (only required if running Synchrony standalone cluster)

Terminology

In this guide we'll use the following terminology:

- Installation directory The directory where you installed Confluence.
- Local home directory The home or data directory stored locally on each cluster node (if Confluence is not running in a cluster, this is simply known as the home directory).
- Shared home directory The directory you created that is accessible to all nodes in the cluster via the same path.

Set up and configure your cluster

We recommend completing this process in a staging environment, and testing your clustered installation, before moving to production.

1. Back up

We strongly recommend that you backup your existing Confluence local home and install directories and your database before proceeding.

You can find the location of your home directory in the `<installation-directory>/confluence/WEB-INF/classes/confluence-init.properties` file.

This is where your search indexes and attachments are stored. If you store attachments outside the Confluence Home directory, you should also backup your attachments directory.

2. Create a shared home directory

1. Create a directory that's accessible to all cluster nodes via the same path. This will be your **shared home** directory.
2. In your existing Confluence home directory, move the contents of `<local home directory>/shared-home` to the new shared home directory you just created. To prevent confusion, we recommend deleting the empty `<local home directory>/shared-home` directory once you've moved its contents.
3. Move your `<local home>/attachments` directory to the new `<shared home>/attachments` directory.

4. Enable cluster mode

Before you enable cluster mode, you should be ready to restart Confluence and configure your cluster. This will require some downtime.

1. Start Confluence.
2. Go to > **General Configuration**.
3. Choose **Clustering** from the sidebar.
4. Select **Enable cluster mode**.
5. Select **Enable** to confirm youre ready to proceed.

5. Restart Confluence

Restart Confluence to configure your cluster. Once you restart, Confluence will be unavailable until youve completed the set up process.

6. Configure your cluster

The setup wizard will prompt you to configure the cluster, by entering:

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

- A name for your cluster
- The path to the shared home directory you created earlier
- The network interface Confluence will use to communicate between nodes
- How you want Confluence to discover cluster nodes:
    - Multicast - enter your own multicast address or automatically generate one.
    - TCP/IP - enter the IP address of each cluster node
    - AWS - enter your IAM Role or secret key, and region.

    We recommend using our Quick Start or Cloud Formation Template to deploy Confluence Data Center in AWS, as it will automatically provision, configure and connect everything you need.

    If you do decide to do your own custom deployment, you can provide the following information to allow Confluence to auto-discover cluster nodes:

Field	Description
IAM Role or **Secret Key**	This is your authentication method. You can choose to authenticate by IAM Role or Secret Key.
Region	This is the region your cluster nodes (EC2 instances) will be running in.
Host header	Optional. This is the AWS endpoint for Confluence to use (the address where the EC2 API can be found, for example 'ec2.amazonaws.com'). Leave blank to use the default endpoint.
Security group name	Optional. Use to narrow the members of your cluster to only resources in a particular security group (specified in the EC2 console).
Tag key and **Tag value**	Optional. Use to narrow the members of your cluster to only resources with particular tags (specified in the EC2 console).

If Synchrony is managed by Confluence, the same network settings will be applied to Synchrony.

Follow the prompts to create the cluster.

When you restart, Confluence will start setting up the cluster. This can take a few minutes. Some core components of Confluence will also change to become cluster compatible. For example, Confluence will switch to a distributed caching layer, managed by Hazelcast.

⚠️Do not restart Confluence until your cluster is set up, and Confluence is back up and running.

Add more Confluence nodes

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Your Data Center license doesnt restrict the number of nodes in your cluster. To achieve the benefits of clustering, such as high availability, youll need to add at least one additional cluster node.

Weve found that typically between 2 and 4 nodes is sufficient for most organizations. In general we recommend starting small and growing as needed.

7. Copy Confluence to the second node

To copy Confluence to the second node:

1. Shut down Confluence on node 1.
2. Copy the installation directory from node 1 to node 2.
3. Copy the local home directory from node 1 to node 2.

Copying the local home directory ensures the Confluence search index, the database and cluster configuration, and any other settings are copied to node 2.

Copying the local home directory ensures the Confluence search index, the database and cluster configuration, and any other settings are copied to node 2.

Make sure your database has sufficient connections available to support the number of nodes.

8. Configure your load balancer

Configure your load balancer for Confluence. You can use the load balancer of your choice, but it needs to support session affinity and WebSockets.

You can verify that your load balancer is sending requests correctly to your existing Confluence server by accessing Confluence through the load balancer and creating a page, then checking that this page can be viewed /edited by another machine through the load balancer.

See Clustering with Confluence Data Center for further load balancer guidance.

9. Start Confluence one node at a time

You must only start Confluence **one node at a time**. The first node must be up and available before starting the next one.

1. Start Confluence on node 1
2. Wait for Confluence to become available on node 1
3. Start Confluence on node 2
4. Wait for Confluence to become available on node 2.

TheCluster monitoring console (> **General Configuration** > **Clustering**) shows information about the active cluster.

When the cluster is running properly, this page displays the details of each node, including system usage and uptime. Use the ••• menu to see more information about each node in the cluster.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

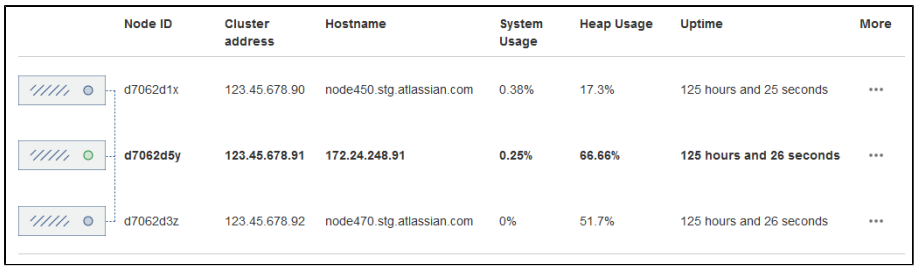

Node ID	Cluster address	Hostname	System Usage	Heap Usage	Uptime	More
d7062d1x	123.45.678.90	node450.stg.atlassian.com	0.38%	17.3%	125 hours and 25 seconds	•••
d7062d5y	**123.45.678.91**	**172.24.248.91**	**0.25%**	**66.66%**	**125 hours and 26 seconds**	•••
d7062d3z	123.45.678.92	node470.stg.atlassian.com	0%	51.7%	125 hours and 26 seconds	•••

10. Test your Confluence cluster

To test creating content you'll need to access Confluence via your load balancer URL. You can't create or edit pages when accessing a node directly.

A simple process to ensure your cluster is working correctly is:

1. Access a node via your load balancer URL, and create a new document on this node.
2. Ensure the new document is visible by accessing it directly on a different node.
3. Search for the new document on the original node, and ensure it appears.
4. Search for the new document on another node, and ensure it appears.

> If Confluence detects more than one instance accessing the database, but not in a working cluster, it will shut itself down in a *clu ster panic*. This can be fixed by troubleshooting the network connectivity of the cluster.

11. Set up a Synchrony cluster (optional)

Synchrony is required for collaborative editing. You have two options for running Synchrony with a Data Center license:

- **managed by Confluence**(recommended)
  This is the default setup. Confluence will automatically launch a Synchrony process on the same node, and manage it for you. No manual steps are required.
- **Standalone Synchrony cluster** (managed by you)
  You deploy and manage Synchrony standalone in its own cluster with as many nodes as you need. Significant setup is required.See S et up a Synchrony cluster for Confluence Data Center for a step-by-step guide.

Head to Administering Collaborative Editing to find out more about collaborative editing.

Troubleshooting

If you have problems with the above process, check our cluster troubleshooting guide.

Were here to help

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Need help setting up your cluster? There are a range of support services available to help you plan and implement a clustered Data Center installation.

- An Atlassian Technical Account Manager can provide strategic guidance. They work with you to develop best practices for configuring, deploying and managing Confluence in a cluster.
- The Atlassian Premier Support team can provide technical support. Premier Support also offers health check analyses to validate the readiness of your environment.
- Atlassian Enterprise Partners offers a wide array of services to help you get the most out of your Atlassian tools.
- You can also ask questions in the Atlassian Community.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Confluence Data Center Performance

This document describes the performance tests we conducted on clustered Confluence Data Center within Atlassian, and the results of those tests. You can compare these data points to your own implementation to predict the type of results you might expect from implementing Confluence Data Center in a cluster in your own organization.

We started our performance tests by taking a fixed load profile (read/write ratio), then tested different cluster set ups against multiples of that load profile.

On this page

- Testing results summary
- Testing methodology and specifications
  - How we tested
  - What we tested
  - Hardware
- Comparison to Confluence Server response times

Testing results summary

Performance gains - Under a high load, clustered Confluence has improved performance overall.

Request responses don't diminish under increased load -Adding more nodes increases throughput, handles higher load and decreases response times.

1547

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

You might observe a different trend/behavior based on your configuration and usage. For details, please see the What we tested section below.

Testing methodology and specifications

The following sections detail the testing environment and methodology we used in our performance tests.

How we tested

Our performance tests were all run on the same controlled isolated lab at Atlassian. For each test, the entire environment was reset and rebuilt. The testing environment included the following components and configuration:

- Apache proxy_balancer
- Postgres database and the required data
- G1GC garbage collector
- 8GB Xmx settings per node
- 6 CPUs per node
- Confluence Server on one machine or Confluence Data Center on two, or four machines as required for the specific test.

To run the test, we used a number of machines in the lab to generate load using scripted browsers and measuring the time taken to perform an action. An action here, means a complete user operation like creating a page or adding comment. Each browser was scripted to perform an action from a predefined list of actions and immediately move on the to next action (i.e. zero think time). Please note that this resulted in each browser performing more tasks than would be possible by a real user and you should not interpret the number of browsersto be equal to the number of real world users. Each test was run for 20 minutes, after which statistics were collected.

What we tested

- All tests used the same Postgres database containing the same number of spaces and pages.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

- The mix of actions we included in the tests represented a sample of the most common user actions* representing six typical types of users (personas). The table below showthe ratio of actionsperformed by each of these personas. These user-based actions were repeated until the test was completed.

Persona	Ratio of actions
PageReader	7
Searcher	1
Editor	1
Creator	1
Commenter	1
Liker	1

Tests were performed with differing load sizes, from 4 up to 96 browsers.For larger load sets, profiles were scaled up, that is, doubling each amount for the 24 browser load, tripled for the 36 browser load.

*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.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

- 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%

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

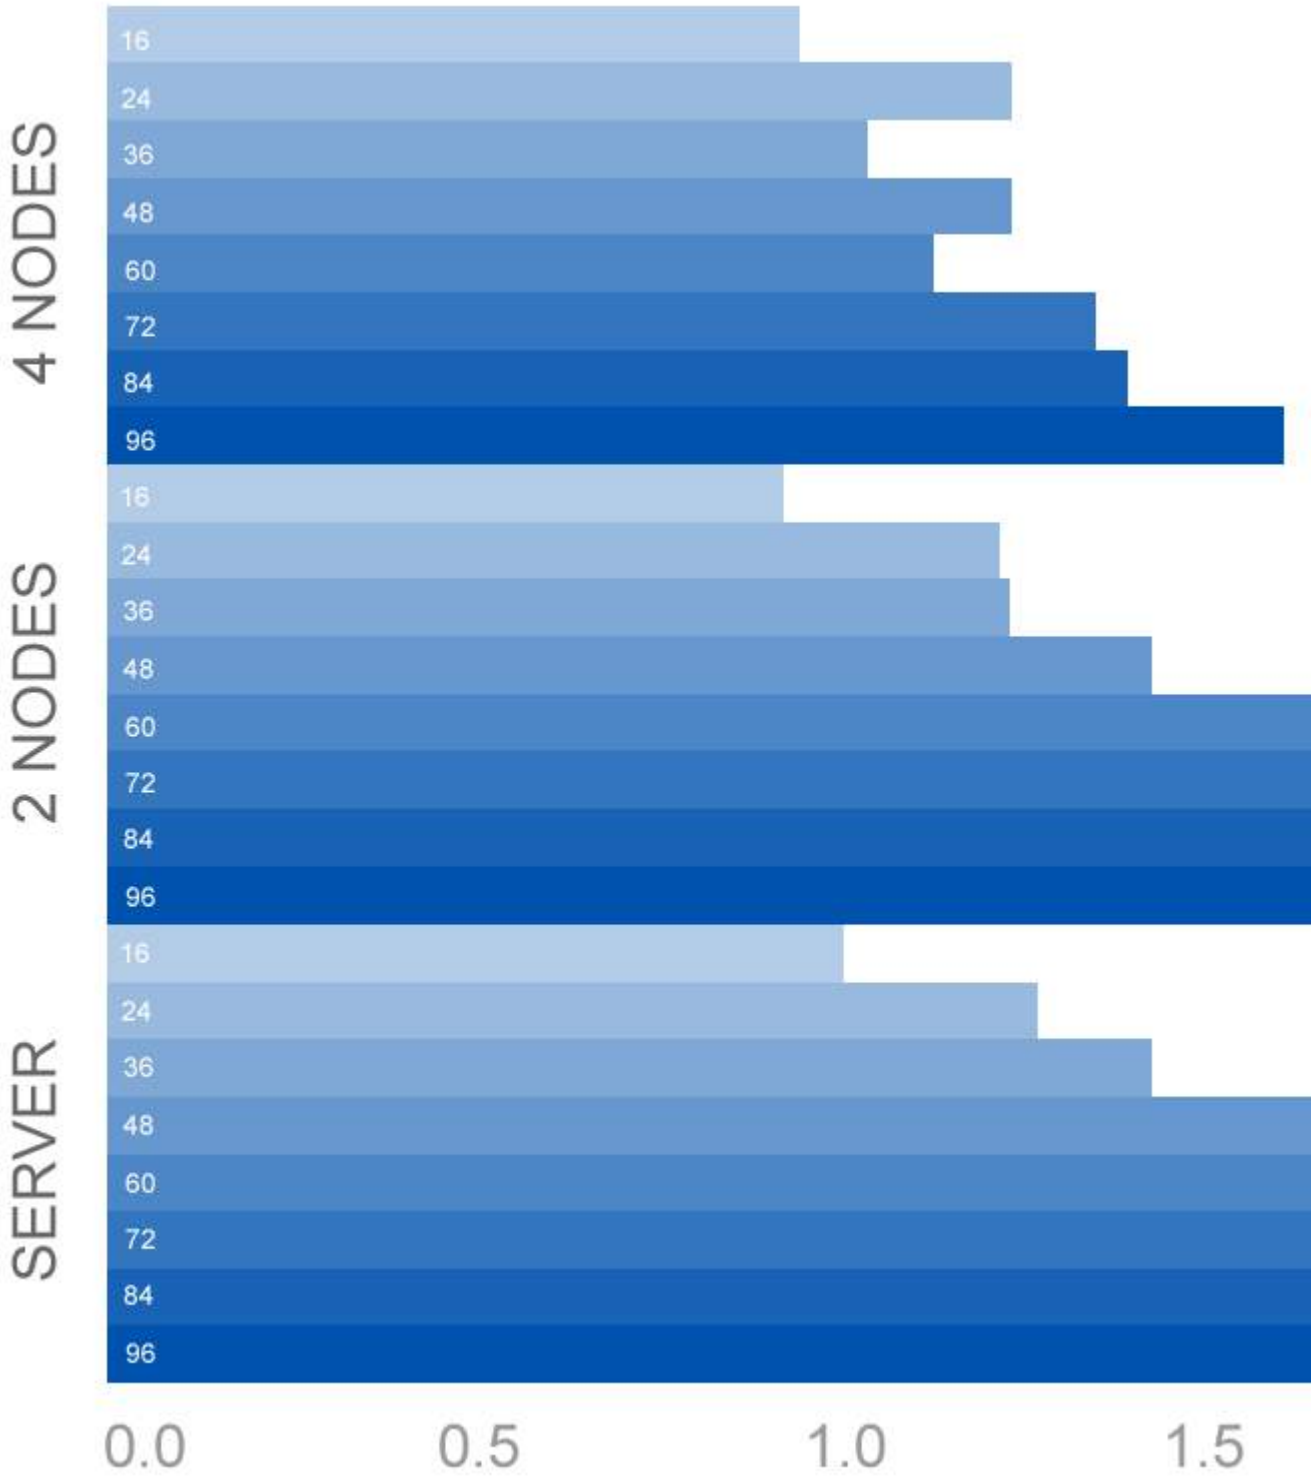

Ready to get started?

Contact usto speak with an Atlassian orget going with Data Centerstraight away.

For a detailed overview of Confluence's clustering solution see Clustering with Confluence Data Center. For help with installation, take a look atInstalling Confluence Data Center.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

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.ForConfluence, this means ensuringConfluence'savailability in the event that your primary site becomes unavailable.

Confluence Data Center documentationis the only Atlassian-supported high-availability solution for Confluence.

Not sure if you should upgrade from Confluence Server to Data Center? Learn more about thebenefits of Confluence Data Center.

This page demonstrates how you can useConfluence Data Center5.9 or laterin implementing and managing a disaster recovery strategy for Confluence. It doesn't, however, cover the broader business practices, like setting the key objectives (RTO, RPO & 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 needConfluence Data Center 5.9 or laterto implement the strategy described in this guide. We'll also assume you've already set up and configured your cluster. See Set up a Confluence Data Center cluster.

This page describes what is generally referred to as a 'cold standby'strategy, which means thestandby Confluence instance isn't continuously running and that you need to take some administrative steps to start the standby instance and ensure it's in a suitable state to service the business needs of your organization.

Maintaining a runbook

The detailed steps will vary from organization to organization and, as such, we recommend you keep a full runbook of steps on file, away from the production system it references.Make your runbook detailed enough such that anyone in the relevant team should be able to complete the steps and recover your service, regardless of prior knowledge or experience. We expect any runbook to contain steps that cover the following parts of the disaster recovery process:

1. Detection of the problem
2. Isolation of the current production environment and bringing it down gracefully
3. Synchronization of data between failed production and intended recovery point
4. Warm up instructions for the recovery instance
5. Documentation, communication, and escalation guidelines

The major components you need to consider in your disaster recovery plan are:

Confluence installation	Your standby site should have exactly the same version of Confluence installed as your production site.

Database	This is the primary source of truth for Confluence and contains most of the Confluence data (except for attachments, avatars, etc). You need to replicate your database and continuously keep it up to date to satisfy your RPO1
Attachments	All attachments are stored in the Confluence Data Center shared home directory, and you need to ensure it's replicated to the standby instance.
Search Index	The search index isn't a primary source of truth, and can always be recreated from the database. For large installations, though, this can be quite time consuming and the functionality of Confluence will be greatly reduced until the index is fully recovered. Confluence Data Center stores search index backups in the shared home directory, which are covered by the shared home directory replication.
Plugins	User installed plugins are stored in the database and are covered by the database replication.
Other data	A few other non-critical items are stored in the Confluence Data Center shared home. Ensure they're also replicated to your standby instance.

Set up a standby system

Step 1. InstallConfluence Data Center5.9or higher

Install the same versionof 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: • Oracle:http://www.oracle.com/technetwork/database/features/data-integration/index.html • PostgreSQL:https://wiki.postgresql.org/wiki/Binary_Replication_Tutorial • MySQL:http://dev.mysql.com/doc/refman/5.7/en/replication.html • Microsoft SQL Server:http://msdn.microsoft.com/en-us/library/ms151198.aspx You need to implement a database replication strategythat meets your RTO, RPO and RCO[1].
Files	You also need to implement a file server replication strategy for the Confluence shared home directory that meets your RTO, RPO and RCO[1].

Clustering considerations

For your clustered environment you need to be aware of the following, in addition to the information above:

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

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**initiallystart only one node**of the cluster, allow it to recover the search index, and check it's working correctly before starting additional nodes.

Disaster recovery testing

You should exercise extreme care when testing any disaster recovery plan. Simple mistakes may cause your live instance to be corrupted, for example, if testing updates are inserted into your production database. You may detrimentally impact your ability to recover from a real disaster, while testing your disaster recovery plan.

The key is to **keep the main data center as isolated as possible from the disaster recovery testing** .

This procedure will ensure that the standby environment will have all the right data, but as the testing environment is completely separate from the standby environment, possible configuration problems on the standby instance are not covered.

Prerequisites

Before you perform any testing, you need to isolate your production data.

Database	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>/confluence.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`

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

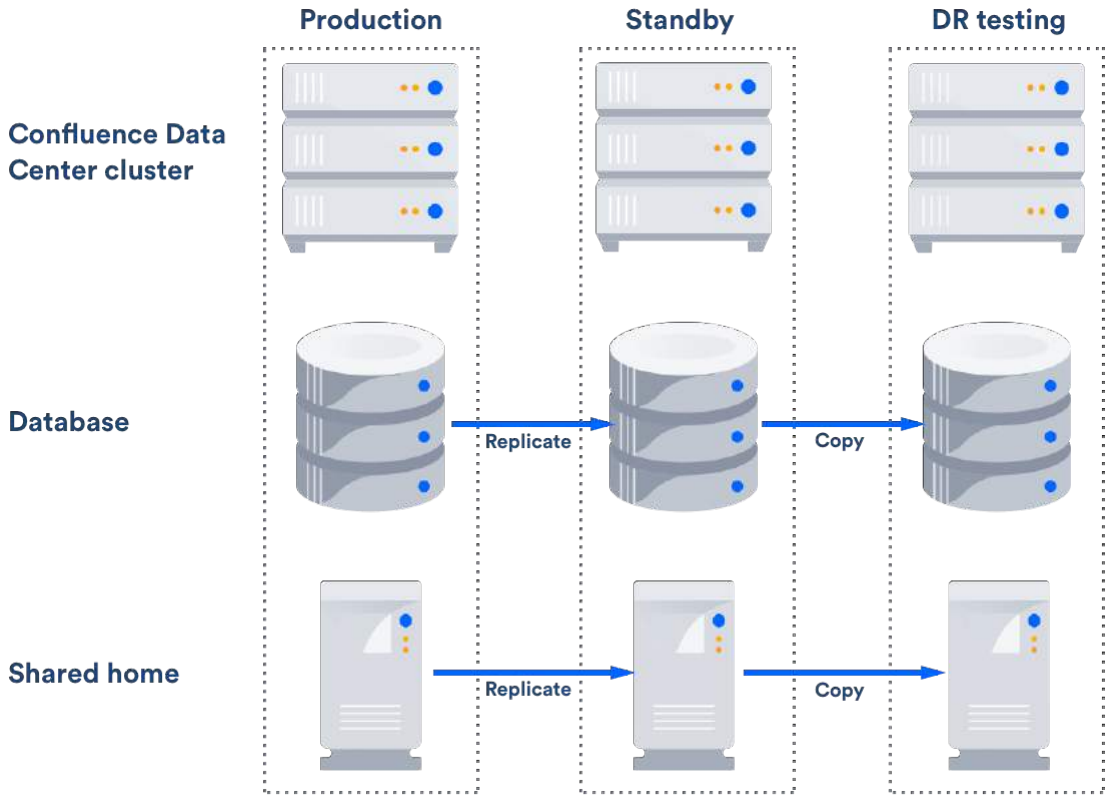

After this you can resume all replication to the standby instance, including the database.

Perform disaster recovery testing

Once you have isolated your production data, follow the steps below to test your disaster recovery plan:

1. Ensure that the new database is ready, with the latestsnapshot andno replication
2. Ensure that the new shared home directory is ready, with the latestsnapshot andno replication
3. Ensureyou have a copy of Confluence on a clean server with the right database and shared home directory settings in`<confluencelocalhome>/confluence.cfg.xml`
4. Ensure you haveconfluence.homemapped, as it was in the standby instance, in the test server
5. Disable email (See`atlassian.mail.senddisabled`inConfiguring System Properties)
6. StartConfluence

Handling a failover

In the event your primary site is unavailable, you'll need to fail over to your standby system. The steps are as follows:

1. Ensure your live system is shutdown and no longer updating the database
2. Ensure the contents of`<confluencesharedhome>` is synced to your standby instance
3. Perform whatever steps are required to activate your standby database
4. Start Confluenceon one node in the standby instance
5. Wait for Confluence to start andcheck 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.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

You need to:

- Synchronize your primary database with the state of the secondary
- Synchronize the primary shared home directory with the state of the secondary

Perform the cut over

1. Shutdown Confluence on the standby instance
2. Ensure the database is synchronized correctly and configured to as required
3. Use rsync or a similar uililty to synchronize the shared home directory to the primary server
4. Start Confluence
5. Check that Confluence is operating as expected
6. Update your DNS, HTTP Proxy, or other front end devices to route traffic to your primary server

Other resources

Troubleshooting

If you encounter problems after failing over to your standby instance, check these FAQs for guidance:

If your database doesn't have the data available that it should, then you'll need to restore the database from a backup.

Once you've restored your database, the search index will no longer by in sync with the database. You can ei ther do a full re-index, background or foreground, or recover from the latest index snapshot if you have one. This includes the journal id file for each index snapshot. The index snapshot can be older than your database backup; it'll synchronize itself as part of the recovery process.

If the search index is corrupt, you can either do a full re-index, background or foreground, or recover from an earlier index snapshot from the shared home directory if you have one.

You may be able to recover them from backups if you have them, or recover from the primary site if you have access to the hard drives.Tools such asrsync may be useful in these circumstances. Missing attachments won't stop Confluence performing normally; the missing attachments won't be available, but users may be able to upload them again.

Application links are stored in the database. If the database replica is up to date, then the application links will be preserved.

You do, however, also need to consider how each end of the link knows the address of the other:

- If you use host names to address the partners in the link and the backup Confluence server has the same hostname, via updates to the DNS or similar, then the links should remain intact and working.
- If the application links were built using IP addresses and these aren't the same, then the application links will need to be re-established.
- If you use IP addresses that are valid on the internal company network but your backup system is remote and outside the original firewall, you'll need to re-establish your application links.

Definitions

RPO	Recovery Point Objective	How up-to-date you require your Confluence instance to be after a failure.
RTO	Recovery Time Objective	How quickly you require your standby system to be available after a failure.
RCO	Recovery Cost Objective	How much you are willing to spend on your disaster recovery solution.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Data Center Troubleshooting

This page covers troubleshooting for a Data Center installation of Confluence.

If you're experiencing Cluster Panic messages in non-clustered installation of Confluence, visit the Knowledge Base article 'Database is being updated by an instance which is not part of the current cluster' Error Message.

⚠️You must ensure the clocks on your cluster nodes don't diverge, as it can result in a range ofproblems with yourcluster.

On this page:

- Symptoms
- Didn't find a solution?

Related pages:

- Troubleshooting a Data Center cluster outage

Symptoms

Below is a list of potential problems with Confluence Data Center, and their likely solutions.

Problem	Likely solutions
`Database is being updated by an instance which is not part of the current cluster` errors on a stand-alone	'Database is being updated by an instance which is not part of the current cluster' Error Message
`Database is being updated by an instance which is not part of the current cluster` errors on a cluster	Add multicast route, Check firewall, Cluster Panic due to Multiple Deployments
`Cannot assign requested address` on startup, featuring an IPv6 address	Prefer IPv4
Error in log: `The interface is not suitable for multicast communication`	Change multicast interface, Add multicast route
Multicast being sent, but not received	Check firewall, Check intermediate routers, Increase multicast TTL
App is unlicensed on some nodes after updating the license on one node.	Disable and re-enable the app in the Universal Plugin Manager.
After an app update, strings appear in the UI instead of buttons and icons on some nodes.	Restart the affected node.
`Hazelcast CANNOT start on this node. No matching network interface found.`	See Hazelcast CANNOT start on this node. No matching network interface found KB article
Any issue not covered here	Contact support

Multicast

- Which multicast address?

The multicast address and port used by Confluence can be found on the Cluster Configuration page, or in `confluence.cfg.xml` in the Confluence home directory.

- Multicast address generation.

Confluence uses a hashing algorithm to take the inputted name during setup and it is then turned into a multicast address stored in the config file. Thus, once the initial setup is completed, Confluence will use the address this is the reason why user can change the address if needed, without actually changing the name. Consequently the additional nodes using the same multicast address specified in the config file are able to join the cluster.

Each node has a multicast address configured in the `confluence.cfg.xml` file

```
name="confluence.cluster.address">xxx.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 interface`	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

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

There are known issues relating to IPv6. You should configure your JVM to try binding to an IPv4 address first.

Change multicast interface

Confluence might have selected the incorrect interface for multicast traffic, which means it cannot connect to other nodes in the cluster. To override the interface used for multicast traffic after initial setup, edit the`confluence.cluster.interface` property in`<local-home>/confluence.cfg.xml`and specify the network interface. For example to tell Confluence to use `eth1`:

```
<property name="confluence.cluster.interface">eth1</property>
```

Overriding Hazelcast Configuration

If the solution to your problem involves changes to the Hazelcast configuration, these changes should**not**be made to the Confluence configuration files. Instead, to ensure your configuration survives upgrades, make your changes by creating a Hazelcast override file.

Increase multicast TTL

The multicast time-to-live (TTL) specifies how many *hops* a multicast packet should be allowed to travel before it is discarded by a router. It should be set to the number of routers in between your clustered nodes: 0 if both are on the same machine, 1 if on two different machines linked by a switch or cable, 2 if on two different machines with one intermediate router, and so on.

To increase the multicast TTL by edit the`confluence.cluster.ttl` property in the`<local home>/confluence.cfg.xml` file on each node. For example to set the TTL to 3:

```
<property name="confluence.cluster.ttl">3</property>
```

Check intermediate routers

Advanced switches and routers have the ability to understand multicast traffic, and route it appropriately. Unfortunately sometimes this functionality doesn't work correctly with the multicast management information (IGMP) published by the operating system running Confluence.

If multicast traffic is problematic, try disabling advanced multicast features on switches and routers in between the clustered nodes. These features can prevent multicast traffic being transmitted by certain operating systems.

Didn't find a solution?

Check Related Articles from the Confluence Knowledge Base

- How to create a support zip via command line when Confluence is down
- Starting Confluence node fails with 'Port [5801] is already in use and auto-increment is disabled. Hazelcast cannot start' error
- Recovering from a Data Center cluster split-brain
- "Exception bootstrapping cluster:Shared home directory is not configured correctly" Error during Confluence Data Center startup
- Cannot find "external_id" column when trying to upgrade to a Confluence CDC license after upgrading from a pre-5.5 Confluence Clustered installation
- Multicast communication works only one-way
- Cluster Panic due to Multicast Traffic Communication Problem
- Hazelcast CANNOT start on this node. No matching network interface found.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

- List of REST APIs available to configure SSO on Confluence Data Center
- 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.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Troubleshooting a Data Center cluster outage

Confluence Data Center cluster outages can be difficult to troubleshoot as the environments are complex and logging can be very verbose.

This page provides a starting point for investigating outages in your cluster.

On this page:

- Establish the originating node
- Investigate common root causes
  - Garbage collection
  - Database connections
  - Network connectivity
- Still having trouble?

Establish the originating node

The most common outage scenario is when something, such as database connectivity issue, network outage or a long garbage collection (GC) process, causes a node to fail to communicate with the cluster for 30 seconds or more and is removed by Hazelcast. The affected node then continues to write to the database, causing a cluster panic.

To establish the originating node:

1. Gather the application log filefrom each node as soon as possible after the outage. Time is critical as the logs will roll over and you may lose the relevant time period.
2. Record identifying information about each node to help you interpret the log messages (IP address, node ID and name of each node).
3. Make a chronological timeline of the events:
    a. Record the time that users or monitoring systems started reporting problems.
    b. View the logs for each node side by side (Hint: we find opening three tabs in node number order helps you always know which logs you are viewing).
    c. Search the logs for '`removing member`'and '`panic`'. This will give you a good idea of which nodes caused the issue and when.
    d. Make a chronological timeline of events from errors to node removal to panics. You can essentially disregard all logging that happens post-panic because once a node panics it needs to be restarted to function effectively. There will be a lot of noise in the logs, but it won't be very useful. The time period we're most interested in will be the minute or so leading up to the first removal or panic event in the logs.

        For example:

        ```
        2:50:15 (approx) Node 3 stopped heartbeating to the cluster for 30s
        (we can estimate this from the time of node removal)
        02:50:45 Node 3 was removed by Node 2
        02:53:15 Node 4 panics
        02:54:15 Node 1, Node 3 and Node 4 receive the panic event and stop processing
        Node 2 remains serving requests
        ```

    e. When you've established when the first affected node was removed, or when the first cluster panic occurred, look back in time in the logs on that node, to look for root causes.

Investigate common root causes

Once you know when the first affected node was removed you can start investigating root causes. From this point on, you're only looking at events on the affected node around the time of removal (in our example above, this is Node 3 at around 2:50). The subsequent removals and panics are usually flow-on effects of the original node removal event, and aren't likely to provide useful root cause information.

Garbage collection

Check the GC logs for the node that was removed (Node 3 in our example). Were there any GC pauses longer than the Hazelcast heartbeat interval (30 seconds by default)? Nodes can't heartbeat during Garbage Collection, so they will be removed from the cluster by one of the other nodes.

If there was a cluster panic, but the node was not removed from the cluster first, check the GC logs for pauses around the time of the panic - pauses that are relatively short (less than 30 seconds) can sometimes still cause panics (due to a race condition) in Confluence 5.10.1 and earlier.

Database connections

Check any database monitoring tools you may have. How many connections to the database were there at the time of the outage?Heartbeats can fail to send if a node can get a connection from its connection pool but not from the database itself, which can lead to nodes being removed from the cluster.

You won't be able to diagnose this from the Confluence logs and will need to look at any external monitoring tools you have for your database. If the outage happens again, check the current number of connections at the db level during the outage.

Network connectivity

Check your network monitoring tools. If a node drops off the network for a short time and cannot communicate with the cluster, it can be removed by the other nodes. Your load balancer logs may be useful here.

Still having trouble?

ContactSupportfor help troubleshooting these outages. Provide them with as much of the information above as possible, to help their investigation.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Use a CDN with Atlassian Data Center applications

On this page:

- Get started with CDN

How it works
How to determine whether a CDN will help your users
What is cached?
Planning your implementation
- Infrastructure requirements
- Considerations for private instances
- Marketplace apps and third party customizations

If your users are distributed across the world and experience poor performancewhen using Data Center products, you may be able to improve their experience by using a Content Delivery Network (CDN). Common CDNs include AWS CloudFront, Cloudflare, Akamai, and others.

CDN support is available in**Data Center**editions of:

- Jira Software 8.3
- Jira Service Management (formerly Jira Service Desk)4.3
- Confluence 7.0
- Bitbucket 6.8.

Get started with CDN

Here's a quick summary of what's involved to enable your CDN in Confluence Data Center:

1. Useour templateto spin up an AWS CloudFront distribution, or create an account with the CDN vendor of your choice.
2. Update yourload balancer and firewallto allow the CDN to reach your site.
3. In Confluence Data Center, provide the CDN URL, and enable CDN support.

As end users access your site, static assets will be cached on the edge server closest to them, and served from there until they expire. This means it might take some time before you can start measuring the impact of the CDN, depending on when your users are online and accessing the site in each location. We don't provide the ability to preload the cache, so assets will be cached as they are served for the first time.

SeeConfigure your CDN for Confluence Data Centerfor the full step-by-step guide.

As always, we recommend testing this on your staging environment, before making any changes to your production site.

How it works

Static assets (such as JavaScript, , and fonts) are cached on edge servers provided by a vendor that are geographically closer to the user. This means when someone views a page, some of the assets needed to display the page are delivered by a server in their region, rather than from yourserver, known as the origin server.This can speed up page load times.

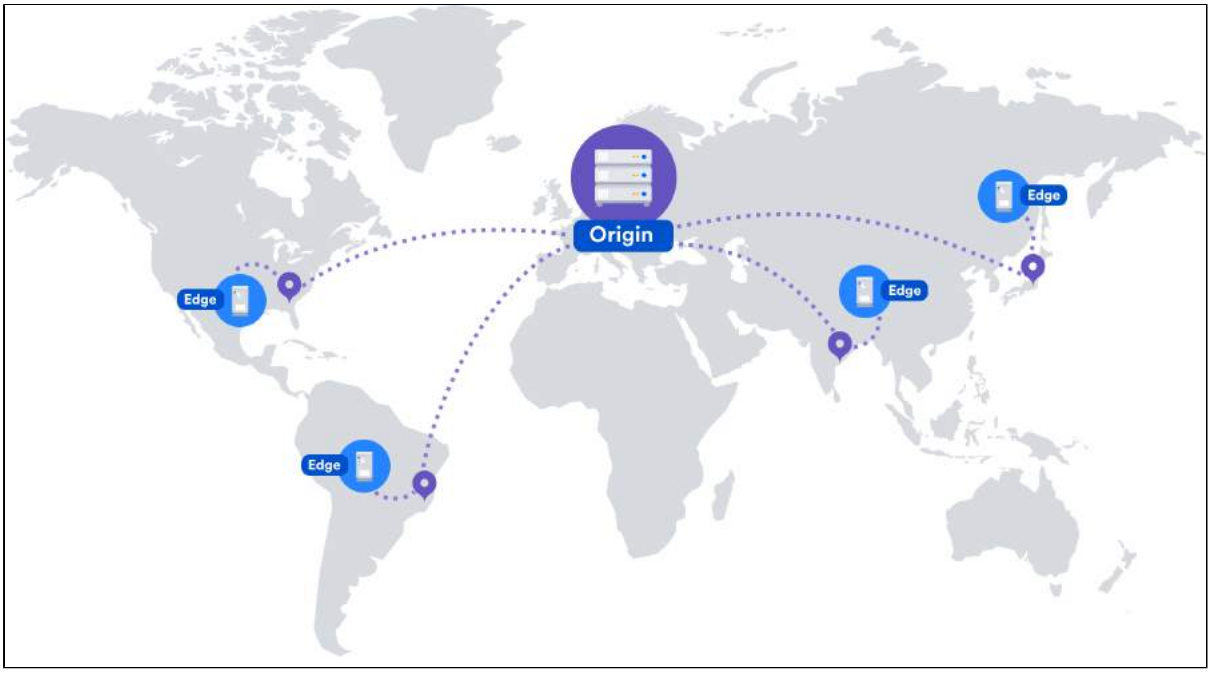

For example, if your server (known as the origin) is in Germany, a can improve page load time by as much as 50% for users located in Rio de Janeiro, as static assets can be served from an edge server in Brazil.If you're new to CDNs and would like to learn more about how they work, CloudFlare provides a great introduction, seehttps://www.cloudflare.com/learning/cdn/performance/.

It's important to note that using a will not make your application inherently faster, what it will do is reduce the load on your cluster,andreduce the latency experienced by some users, which should result in faster page load times for users.

Tests on our internal dogfooding instances located in Gdask, Poland have shown the response time for the View Issue action in Jira Data Center is ~50% faster for people accessing from US East, when is enabled.

How to determine whether a CDN will help your users

A good starting point when assessing whether a CDN will help your users, is to take a look at the network overhead experienced in your site.

Go to**Content Delivery Network**in the admin console of your Data Center application. On the **Performance** tab you'll see the percentage of requests that had a transfer cost of more than one second. Put simply, the higher the percentage, the more likely it is that your users requests are being affected by network conditions, such as latency and connection quality.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

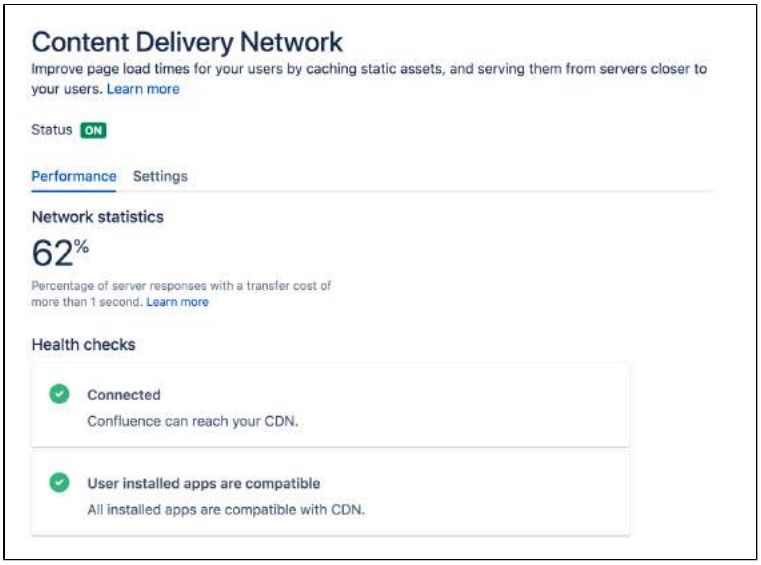

This network statistic is a useful indicator of the network conditions your users experience when using the product. If the percentage is high, it's likely that using a will benefit your users in these conditions.

> As users access pages in your site (for example a Confluence page, Jira issue, or Bitbucket pull request page), we measure the amount of time the browser has to wait to get the content of that page. We then subtract the time required to render the page on the server. This leaves us with the time it took to send the request and retrieve the response.
>
> This time is dependent mostly on the latency between the server and the browser, but also includes things like SSL connection setup time.
>
> This metric is collected on requests that don't use , so it will continue to provide consistent statistics on your network, even after you enable .

You should also consider where your users are geographically located.For example, if your servers are located in Frankfurt, and the majority of your teams are located in Germany and Austria, your team based in Malaysia may be suffering from high latency, resulting in slow page load times.

Network diagnostic tools such as `traceroute`, `ping`, and `mtr`can be helpful to determine the amount of latency being experienced.

> In these examples we'll use traceroute todisplay some basic network statistics, including latency information. Remember to replace yoursite.comwith your base URL.
>
> In Windows, open Command Prompt and enter the following:
>
> ```
> > tracert yoursite.com
> ```
>
> In Linux or Mac OS, open Terminal and enter the following:
>
> ```
> $ traceroute yoursite.com
> ```
>
> This will display the number of hops, and three latency times, in milliseconds, for each server. Average the three figures to get the latency for that server.
>
> The `mtr` command (my traceroute) is a useful combination of `ping` and `traceroute`. You will need to install `mtr` to be able to use it in MacOS or Windows.

What is cached?

We only cache static assets served by a Data Center application or Marketplace app. These are things that are only going to change when you upgrade your Data Center application or app.Dynamic content is not cached.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Here's a summary of what will be cached when you enable :

Cached	Not cached
• JavaScript • Fonts	• attached files • pages or issues • personal information, including avatars • assets that are part of a theme

You shouldn't need to ever manually invalidate the cache, as we handle this when you upgrade your Data Center product, or an app.

⚠ If youre performingZDU(Zero Downtime Upgrade), we highly recommend that you disableCDNbefore the upgrade and enable it after the cluster is in a stable state. Otherwise, you might experience some issues related to theCDNperformance.

Planning your implementation

Infrastructure requirements

You can use any origin pull. You're responsible for any costs associated with your CDN.

We've prepared a CloudFormation template that you can use to configure Amazon CloudFront with minimal effort. You can find all our deployment resources in this repositoryhttps://bitbucket.org/atlassian/atlassian-aws-deployment/src/master/templates/cdn/.

There are some other infrastructure requirements that you need to be aware of before you start:

- **HTTP/2 is highly recommended**
  Your load balancer, firewall, or proxy should allow HTTP/2 traffic. Using HTTP/2 will provide the best performance for your end users. Check the documentation for your particular provider to find out how to do this.
- **Firewall considerations**
  Your must be able to access and cache static assets. If your instance is not publicly accessible will you need to make some changes to your firewall to allow requests from the to pass through.We recommend using application firewalls instead of standard IP range filtering, as IP ranges can change without notice.

Considerations for private instances

If your site is publicly accessible on the internet, you should be able to enable without any problems.

If your site is not publicly accessible you can:

- configure your firewall to allow requests from your to pass through. More information on how to do this is provided in our step-by-step guides below.
- set up your own caching servers closer to your users which will not require opening any traffic to the internet, instead of using a vendor. SeeHow to configure Apache for caching and HTTP/2to learn more about this workaround.

Marketplace apps and third party customizations

Some marketplace apps or customizations may not be compatible with the feature. A health check, on the Content Delivery Network admin screen will let you know if any of your apps are not compatible.

SeeUser-installed apps health check fails in Data Center when configuring CDNto find out what to do if any of your apps are incompatible.

If you've developed your own plugin, seePreparing for Confluence 7.0for information about the APIs you can use to confirm your plugin is compatible.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Configure your CDN for Confluence Data Center

On this page:

- Configure an internet facing load balancer (optional)
  - Add an internet-facing load balancer
  - Update your firewall rules for the internet-facing load balancer
- Configure your CDN to cache assets
- Enable CDN in Confluence
  - Configure CDN in Confluence via REST API
- Troubleshooting
  - Frequently asked questions

If your users are distributed across the world and experience high latency when using Confluence Data Center, you may be able to improve their experience by using a Content Delivery Network (CDN). Common CDNs include AWS CloudFront, Cloudflare, Azure CDN, Akamai, and others.

Head toUse a CDN with Atlassian Data Center applicationsto learn about our CDN capabilities, and how to assess whether it will improve your users' experience.

Once you're ready to start using a CDN, there are three main steps:

1. Configure an internet-facing load balancer (optional)
2. Configure your CDN.
3. Enable the CDN feature in Confluence.

Configure an internet facing load balancer (optional)

If your site is not publicly accessible, you'll need to make sure that your CDN can reach it, but only to access and cache static assets. The way you do this depends on your particular load balancer and web application firewall. Refer to the documentation for your load balancer and firewall for detailed guidance.

Add an internet-facing load balancer

Add an internet-facing load balancer to your setup. This is in addition to your primary load balancer. Your CDN is the only entity that will interact with this load balancer. We recommend you:

- Enable HTTPS - the traffic from this load balancer will be sent over the public internet and should be encrypted.
- Enable HTTP/1.1 - currently, the caching proxies and CDNs do not handle HTTP/2 well (or at all) on the way to the origin.
- For AWS deployments, you would set up an internet-facing application load balancer.

Update your firewall rules for the internet-facing load balancer

Unlike your primary load balancer, this internet-facing load balancer must be locked down to ensure that your CDN can only pull data it is allowed to cache. When configuring your firewall rules we recommend:

1. The configuration should only allow requests for paths that start with "/s/". If your application is deployed with a context path (for example yoursite.com/wiki or yoursite.com/jira) you will need to include it in the path. All other requests must be blocked.
2. You can also choose to limit the allowedHTTP methods to GET, HEAD, OPTIONS.

For AWS deployments, you will configure a Web Access Control List (WebACL) in the Web Application Firewall attached to your application load balancer. The condition to use is a "string match condition" applied to "URI".

To check that your setup is secure, perform the following manual tests:

1. A GET on`https://internet-facing-proxy/` should return "403 FORBIDDEN".
2. A GET on`https://internet-facing-proxy/s` should return "403 FORBIDDEN".

3. A GET on `https://internet-facing-proxy/s/` should return "404 NOT FOUND".
4. A GET on `https://internet-facing-proxy/s/` should return "403 FORBIDDEN".
5. A GET on `https://internet-facing-proxy/s/../s/` should return "404 NOT FOUND".

Configure your CDN to cache assets

You'll need an account with a CDN provider. You're responsible for all costs associated with your CDN.We only support serving static assets from a CDN at this time. This means page content, attached files, and personally identifiable information, including things like user avatars, won't be cached by your CDN.

We've prepared a CloudFormation template that you can use to configure Amazon CloudFront with minimal effort. You can find all our AWS deployment resources in this repositoryhttps://bitbucket.org/atlassian/atlassian-aws-deployment/src/master/templates/cdn/.

If you choose not to use our template, define the following in your CDN configuration. This example is based on AWS CloudFront.

Origin domain name	This is your Atlassian application base URL, including the context path if you've configured one. For example:`mycompany.com/confluence`
Origin path	Leave blank. There is no need to specify a path.
Allowed HTTP methods	Optionally limit to: GET, HEAD, OPTIONS
Viewer protocol policy	redirect HTTP to HTTPS
Object caching	Use origin cache headers
Forward cookies	None This is important to make sure static assets are cached without the user context.
Query String Forwarding and Caching	Forward all, cache based on all
HTTP protocols	Must include HTTP/2
Error pages/Error Caching Minimum TTL (seconds)	The default error page caching time for CloudFront is 5 minutes. Consider lowering it to a value in the range of 10-30 seconds to decrease the time required to recover from an outage.
Compress Objects Automatically	Yes

Using the default should be fine for most of the other settings.

You will need to adapt this information for your particularCDN provider. You should refer to the documentation for your CDN for details, as we've found that terminology differs between CDNs.

Enable CDN inConfluence

Once you've configured your CDN, you can enable the CDN option in Confluence.

To turn on CDN:

1. Go to > **General Configuration** > **Content Delivery Network**.
2. Navigate to the**Settings**tab.
3. Set the status to**On**
4. Paste the URL generated by your CDNinto the URL field and hit**Validate**.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

5. If successful,save your changes.

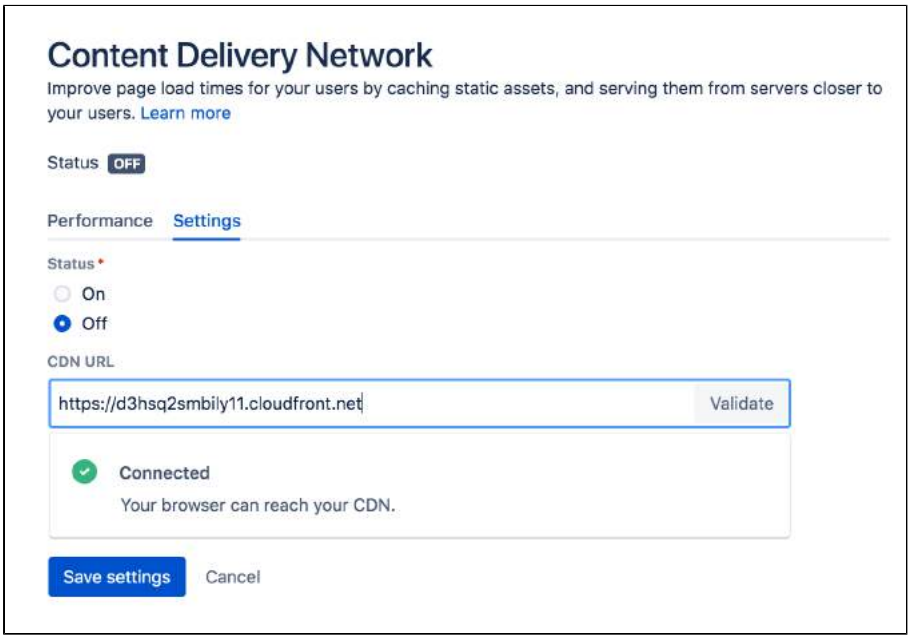

As end users access Confluence, static assets will be cached on the edge server closest to them, and served from there until they expire. This means it might take some time before you can start measuring the impact of the CDN, depending on when your users are online and accessing the site in each location.

Configure CDN in Confluence via REST API

You can also interact with the CDN feature using the following REST endpoint:`<base-url>/rest/static-asset-caching/configuration`

- **GET** -returns the current CDN status, and URL.
- **DELETE** -deletes the existing configuration and reverts to the default state (CDN disabled, no URL). This is useful if you can't access the UI because of a caching problem.
- **PUT** -sets the CDN URL and status to the values passed in the body of the request as follows:

```
{
  enabled: true,
  url: https://yourcdnurl.com
}
```

Troubleshooting

Here are some common problems that you may encounter.

- **We only accept HTTPS CDN URLs**
  This is particularly important if you're using Azure CDN,as Azure CDN will mirror the same protocol as the originating request, which means your Data Center application will need to be provisioned with HTTPS.
- **Data Center application UI is inaccessible or not functional**
  Although unlikely, a misconfiguration of your CDNor a CDN service outage may mean your application's UI is not accessible.If this happens, you will need to disable the CDN feature using the REST API, as follows.

```
curl -v -u <admin username>:<admin password> -X DELETE http://<your-base-url>/rest/static-asset-caching/configuration
```


Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

This example uses Curl, but you can use any language. Don't forget to replace the username, password, and base URL placeholders with your own details.

- **HTTP/2 disabled**
  Your load balancer, firewall, or reverse proxy shouldallow HTTP/2 traffic.Using HTTP/2 will provide the best performance for your end users. See HTTP/2 health check fails in Data Center when configuring CDNfor more information.
- **User-installed apps may not be compatible**
  This warning is displayed when we detect that a Marketplace or other user-installed app is using a deprecated method, which may result in assets being cached incorrectly. SeeUser-installed apps health check fails in Data Center when configuring CDNfor more information on what to do if you see this warning.

Frequently asked questions

Can I control what static assets are cached?

No, the application controls this. All requests for static assets are routed to the CDN. Requests for non-static assets are routed directly to your product.

Is personally identifiable information cached?

User created content, usernames, mentions, avatars etc are not static assets, so are not cached. Your CDN should also be configured to pull content from your product with cookies stripped to make sure it operates without user context.

Is dynamic content such as `batch.js`cached?

Although dynamically generated, `batch.js`is considered static content, so is cached.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Improving instance stability with rate limiting

When automated integrations or scripts send requests to Confluence in huge bursts, it can affect Confluences stability, leading to drops in performance or even downtime. With rate limiting, you can control how many external REST API requests automations and users can make and how often they can make them, making sure that your Confluence instance remains stable.

Rate limiting is available for **Confluence Data Center**.

On this page:

- How rate limiting works
- How to turn on rate limiting
- Limiting requests what its all about
- Adding exemptions
- Identifying users who have been rate limited
- Viewing limited requests in the Confluence log file
- Getting rate limited users perspective
- Other tasks

How rate limiting works

Heres some details about how rate limiting works in Confluence.

Rate limiting targets only external REST API requests, which means that requests made within Confluence arent limited in any way. When users move around Confluence, creating pages, commenting, and completing other actions, they wont be affected by rate limiting, as were seeing this as a regular user experience that shouldnt be limited.

Lets use an example to better illustrate this:

- When a user visits a space in Confluence, a number of requests are sent in the background these requests ask Confluence for the pages, blog posts, etc. Since this traffic is internal to Confluence, it wont be limited.
- When the same user opens up the terminal on their laptop and sends a request (like the one below) to get the contents of a space, it will be rate limited because its made outside of Confluence.

```
curl -u user:password http://localhost:8090/rest/api/space/SPACEKEY/content
```

Authentication mechanisms

To give you more details on how we recognize which requests should be limited, were targeting external HTTP requests with these authentication mechanisms:

- Basic auth
- OAuth
- JSESSIONID cookie

Out of the many available techniques for enforcing rate limits, weve chosen to use token bucket, which gives users a balance of tokens that can be exchanged for requests. Heres a summary of how it works:

Users are given tokens that are exchanged for requests. One token equals one request.

Users get new tokens at a constant rate so they can keep making new requests. This is their Requests allowed, and can be, for example, 10 every 1 minute.

Tokens are added to a users personal bucket until its full. This is their Max requests and allows them to adjust the usage of tokens to their own frequency, for example 20 every 2 minutes instead of 10 every 1 minute, as specified in their usual rate.

When a user tries to send more requests than the number of tokens they have, only requests that can draw tokens from the bucket will be successful. The remaining ones will end in a 429 error message (too many requests). The user can retry those requests once they get new tokens.

Confluence tastes best when used with our other products like Jira. Technically, products like these are external to Confluence, so they should be limited. In this case, however, were treating them as belonging to the same user experience and dont want to enforce any limits for requests coming from or to these products.

The way it is now:

- Server: Not limited in any way.
- Cloud: Theres a known issue that applies rate limits to requests coming from/to cloud products. Were working hard to disable rate limits for cloud products and should make that happen soon. For now, if youre integrating Confluence with Jira cloud, you should make rate limits higher than usual.

The general assumption is that Marketplace apps are installed on a Confluence instance, make internal requests from within Confluence, and shouldnt be limited. But, as always, **it depends on how an app works**.

- **Internal**: If an app in fact works internally, enhancing the user experience, it wont be limited. An example of such app would be a special banner thats displayed in a Confluence space. Lets say this banner checks all pages that were created and shows this spaces winner a user whos created the most pages in the last month. Traffic like that would be internal, not limited.
- **External**: Apps whose requests are external to Confluence are limited. Lets say we have an app that displays a wallboard on TV. It asks Jira for details about boards, issues, assignees, etc. and then reshuffles and displays them in its own way as the earlier mentioned wallboard. An app like that sends external requests and behaves just like a user sending requests over a terminal.

It really depends on the app, but were assuming most of them shouldnt be limited.

Rate limiting is available for Data Center, so you most likely have a cluster of nodes behind a load balancer. You should know that each of your users will have a separate limit on each node (rate limits are applied per node, not per cluster).

In other words, if they have used their Requests allowed on one node and were rate limited, they could theoretically send requests again if they started a new session on a different node. Switching between the nodes isnt something users can do, but keep in mind that this can happen.

Whatever limit youve chosen (e.g. 100 requests every 1 hour), the same limit will apply to each node, you dont have to set it separately. This means that each users ability to send requests will still be limited, and Confluence will remain stable regardless of which node their requests are routed to.

Setting the right limit depends on many factors, so we cant give you a simple answer. We have some suggestions, though.

Finding the right limit

The first step is to understand the size of traffic that your instance receives. You can do this by parsing the access log and finding a user than made the most REST requests over a day. Since UI traffic is not rate limited, this number will be higher than what you need as your rate limit. Now, thats a base number you need to modify it further based on the following questions:

1. Can you afford to interrupt your users work? If your users integrations are mission-critical, consider upgrading your hardware instead. The more critical the integrations, the higher the limit should be consider multiplying the number you found by two or three.
2. Is your instance already experiencing problems due to the amount of REST traffic? If yes, then choose a limit thats close to the base number you found on a day when the instance didnt struggle. And if youre not experiencing significant problems, consider adding an extra 50% to the base number this shouldnt interrupt your users and you still keep some capacity.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

In general, the limit you choose should keep your instance safe, not control individual users. Rate limiting is more about protecting Confluence from integrations and scripts going haywire, rather than stoping users from getting their work done.

How to turn on rate limiting

You need the System Administrator global permission to turn on rate limiting.

To turn on rate limiting:

1. In Confluence, go to > **General Configuration**> **Rate limiting**.
2. Change the status to **Enabled**.
3. Select one of the options: **Allow unlimited requests**, **Block all requests**, or **Limit requests**. The first and second are all about allowlisting and blocklisting. For the last option, youll need to enter actual limits. You can read more about them below.
4. **Save** your changes.

Make sure to add exemptions for users who really need those extra requests, especially if youve chosen allowlisting or blocklisting. See Adding exemptions.

Limiting requests what its all about

As much as allowlisting and blocklisting shouldnt require additional explanation, youll probably be using the **L imit requests** option quite often, either as a global setting or in exemptions.

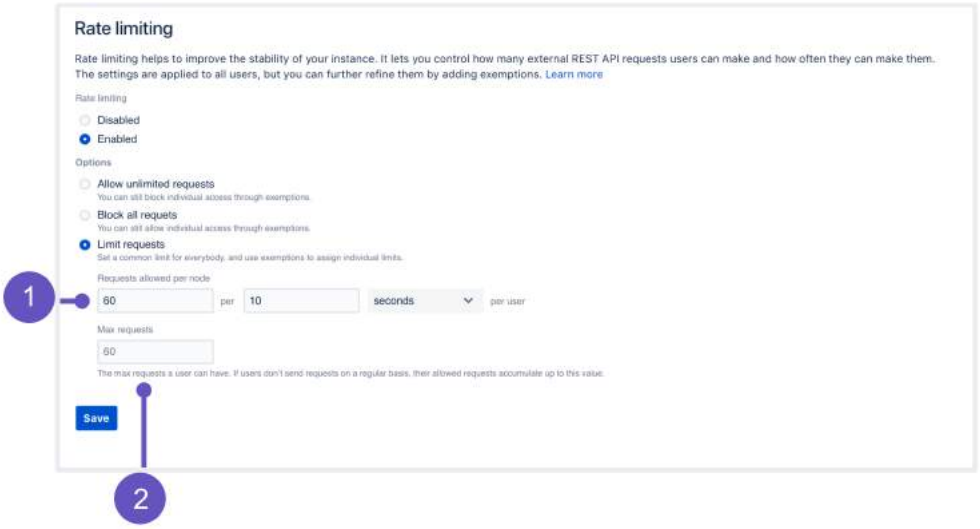

Lets have a closer look at this option and how it works:

1. **Requests allowed**: Every user is allowed a certain amount of requests in a chosen time interval. It can be 10 requests every second, 100 requests every hour, or any other configuration you choose.
2. **Max requests (advanced)**: Allowed requests, if not sent frequently, can be accumulated up to a set maximum per user. This option allows users to make requests at a different frequency than their usual rate (for example, 20 every 2 minutes instead of 10 every 1 minute, as specified in their rate), or accumulate more requests over time and send them in a single burst, if thats what they need. Too advanced? Just make it equal to Requests allowed, and forget about this field nothing more will be accumulated.

Examples

Requests allowed: 10/hour | Max requests: 100

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

One of the developers is sending requests on a regular basis, 10 per hour, throughout the day. If they try sending 20 requests in a single burst, only 10 of them will be successful. They could retry the remaining 10 in the next hour when theyre allowed new requests.

Another developer hasnt sent any requests for the past 10 hours, so their allowed requests kept accumulating until they reached 100, which is the max requests they can have. They can now send a burst of 100 requests and all of them will be successful. Once they used up all available requests, they have to wait for another hour, and theyll only get the allowed 10 requests.

If this same developer sent only 50 out of their 100 requests, they could send another 50 right away, or start accumulating again in the next hour.

Requests allowed: 1/second | Max requests: 60

A developer can choose to send 1 request every second or 60 requests every minute (at any frequency).

Since they can use the available 60 requests at any frequency, they can also send all of them at once or in very short intervals. In such a case, they would be exceeding their usual rate of 1 request per second.

Finding the right limit

Setting the right limit depends on many factors, so we cant give you a simple answer. We have some suggestions, though.

Finding the right limit

The first step is to understand the size of traffic that your instance receives. You can do this by parsing the access log and finding a user that made the most REST requests over a day. Since UI traffic is not rate limited, this number will be higher than what you need as your rate limit. Now, thats a base number you need to modify it further based on the following questions:

1. Can you afford to interrupt your users work? If your users integrations are mission-critical, consider upgrading your hardware instead. The more critical the integrations, the higher the limit should be consider multiplying the number you found by two or three.
2. Is your instance already experiencing problems due to the amount of REST traffic? If yes, then choose a limit thats close to the base number you found on a day when the instance didnt struggle. And if youre not experiencing significant problems, consider adding an extra 50% to the base number this shouldnt interrupt your users and you still keep some capacity.

In general, the limit you choose should aim at keeping your instance safe, not to control individual users. Rate limiting is more about protecting Jira from integrations and scripts going haywire, rather than stoping users from getting their work done.

Adding exemptions

Exemptions are, well, special limits for users who really need to make more requests than others. Any exemptions you choose will take precedence over global settings.

After adding or editing an exemption, youll see the changes right away, but it takes up to 1 minute to apply the new settings to a user.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

To add an exemption:

1. Go to the **Exemptions** tab.
2. Click **Add exemption**.
3. Find the user and choose their new settings.
   You cant choose groups, but you can select multiple users.
4. The options available here are just the same as in global settings: **Allow unlimited requests**, **Block all requests**, or **Assign custom limit**.
5. **Save** your changes.

If you want to edit an exemption later, just click **Edit** next to a users name in the **Exemptions** tab.

Recommended: Add an exemption for anonymous access

Confluence sees all anonymous traffic as made by one user: **Anonymous**. If your site is public, and your rate limits are not too high, a single person may drain the limit assigned to anonymous. Its a good idea to add an exemption for this account with a higher limit, and then observe whether you need to increase it further.

Identifying users who have been rate limited

When a user is rate limited, theyll know immediately as theyll receive an HTTP 429 error message (too many requests). You can identify users that have been rate limited by opening the **List of limited accounts** tab on the rate limiting settings page. The list shows all users from the whole cluster.

When a user is rate limited, it takes up to 5 minutes to show it in the table.

Unusual accounts

Youll recognize the users shown on the list by their name. It might happen, though, that the list will show some unusual accounts, so heres what they mean:

- **Unknown**: Thats a user that has been deleted in Confluence. They shouldnt appear on the list for more than 24 hours (as they cant be rate limited anymore), but you might see them in the list of exemptions. Just delete any settings for them, they dont need rate limiting anymore.
- **Anonymous**: This entry gathers all requests that werent made from an authenticated account. Since one user can easily use the limit for anonymous access, it might be a good idea to add an exemption for anonymous traffic and give it a higher limit.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Viewing limited requests in the Confluence log file

You can also view information about rate limited users and requests in the Confluence log file. This is useful if you want to get more details about the URLs that requests targeted or originated from.

When a request has been rate limited youll see a log entry similar to this one:

```
2019-12-24 10:18:23,265 WARN [http-nio-8090-exec-7] [ratelimiting.internal.filter.RateLimitFilter]
lambda$userHasBeenRateLimited$0 User [2c9d88986ee7cdaa016ee7d40bd20002] has been rate limited
 -- url: /rest/api/space/DS/content | traceId: 30c0edcb94620c83 | userName: exampleuser
```

Getting rate limited users perspective

When users make authenticated requests, theyll see rate limiting headers in the response. These headers are added to every response, not just when youre rate limited.

Header	Description
`X-RateLimit-Limit`	The max number of requests (tokens) you can have. New tokens wont be added to your bucket after reaching this limit. Your admin configures this as Max requests.
`X-RateLimit-Remaining`	The remaining number of tokens. This value is as accurate as it can be at the time of making a request, but it might not always be correct.
`X-RateLimit-Interval-Seconds`	The time interval in seconds. You get a batch of new tokens every time interval.
`X-RateLimit-FillRate`	The number of tokens you get every time interval. Your admin configures this as Requests allowed.
`retry-after`	How long you need to wait until you get new tokens. If you still have tokens left, it shows 0; this means you can make more requests right away.

When youre rate limited and your request doesnt go through, youll see the HTTP 429 error message (too many requests). You can use these headers to adjust scripts and automations to your limits, making them send requests at a reasonable frequency.

Other tasks

Allowlisting URLs and resources

Weve also added a way to allow whole URLs and resources on your Confluence instance using a system property. This should be used as quick fix for something that gets rate limited, but shouldnt.

> For example, a Marketplace app added some new API to Confluence. The app itself is used from the UI, so it shouldnt be limited, but it might happen that Confluence sees this traffic as external and applies the rate limit. In this case, you could disable the app or increase the rate limit, but this brings additional complications.
>
> To work around issues like this, you can allowlist the whole resource added by the app so it works without any limits.

To allow specific URLs to be excluded from rate limiting:

1. Stop Confluence.
2. Add the`com.atlassian.ratelimiting.whitelisted-url-patterns`system property, and set the value to a comma-separated list of URLs, for example:

```
-Dcom.atlassian.ratelimiting.whitelisted-url-patterns=/**/rest/applinks/**,/**/rest/capabilities,
/**/rest/someapi
```


Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

The way you add system properties depends on how you run Confluence. See Configuring System Properties for more information.

3. Restart Confluence.

For more info on how to create URL patterns, see AntPathMatcher: URL patterns.

Allowlisting external applications

You can also allowlist consumer keys, which lets you remove rate limits for external applications integrated through AppLinks.

> If you're integrating Confluence with other Atlassian products, you don't have to allowlist them as this traffic isn't limited.

1. Find the consumer key of your application.
    a. Go to > **General Configuration> Application Links**.
    b. Find your application, and click **Edit**.
    c. Copy the Consumer Key from**Incoming Authentication**.
2. Allowlist the consumer key.
    a. Stop Confluence.
    b. Add the`com.atlassian.ratelimiting.whitelisted-oauth-consumers` system property, and set the value to a comma-separated list of consumer keys, for example:

    ```
    -Dcom.atlassian.ratelimiting.whitelisted-oauth-consumers=app-connector-for-confluence-server
    ```

    The way you add system properties depends on how you run Confluence. SeeConfiguring System Propertiesfor more information.
    c. Restart Confluence.

After entering the consumer key, the traffic coming from the related application will no longer be limited.

Adjusting your code for rate limiting

Weve created a set of strategies you can apply in your code (scripts, integrations, apps) so it works with rate limits, whatever they are.

For more info, see Adjusting your code for rate limiting.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Adjusting your code for rate limiting

Whether its a script, integration, or app youre using if its making external REST API requests, it will be affected by rate limiting. Until now, you could send an unlimited number of REST API requests to retrieve data from Confluence, so were guessing you havent put any restrictions on your code. When admins enable rate limiting in Confluence, theres a chance your requests will get limited eventually, so we want to help you prepare for that.

Before you begin

To better understand the strategies weve described here, its good to have some some basic knowledge about rate limiting in Confluence. When in doubt, head to Improving instance stability with rate limiting and have a look at the first paragraph.

Quick reference

Success: When your request is successful, youll get a 2xx code.

Error: When your request fails, youll get a 4xx code. If youre rate limited, it will be 429 (too many requests).

The following HTTP headers are added to every **authenticated** request affected by rate limiting:

Header	Description
`X-RateLimit-Limit`	The max number of requests (tokens) you can have. New tokens wont be added to your bucket after reaching this limit. Your admin configures this as Max requests.
`X-RateLimit-Remaining`	The remaining number of tokens. This value is as accurate as it can be at the time of making a request, but it might not always be correct.
`X-RateLimit-Interval-Seconds`	The time interval in seconds. You get a batch of new tokens every time interval.
`X-RateLimit-FillRate`	The number of tokens you get every time interval. Your admin configures this as Requests allowed.
`retry-after`	How long you need to wait until you get new tokens. If you still have tokens left, it shows 0; this means you can make more requests right away.

Strategies

Weve created a set of strategies you can apply in your code so it works with rate limits. From very specific to more universal, these reference strategies will give you a base, which you can further refine to make an implementation that works best for you.

1. Exponential backoff

This strategy is the most universal and the least complex to implement. Its not expecting HTTP headers or any information specific to a rate limiting system, so the same code will work for the whole Atlassian suite, and most likely non-Atlassian products, too. The essence of using it is observing whether youre already limited (wait and retry, until requests go through again) or not (just keep sending requests until youre limited).

Universal, works with any rate limiting system.

Doesnt require too much knowledge about limits or a rate limiting system.

High impact on a Confluence instance because of concurrency. Were assuming most active users will send requests whenever theyre available. This window will be similar for all users, making spikes in Confluence performance. The same applies to threads most will either be busy at the same time or idle.

Unpredictable. If you need to make a few critical requests, you cant be sure all of them will be successful.

Summary of this strategy

Heres the high-level overview of how to adjust your code:

1. **Active**: Make requests until you encounter a 429. Keep concurrency to a minimum to know exactly when you reached your rate limit.
2. **Timeout**: After you receive a 429, start the timeout. Set it to 1 second for starters. Its a good idea to wait longer than your chosen timeout up to 50%.
3. **Retry**: After the timeout has passed, make requests again:
    a. **Success**: If you get a 2xx message, go back to step 1 and make more requests.
    b. **Limited**: If you get a 429 message, go back to step 2 and double the initial timeout. You can stop once you reach a certain threshold, like 20 minutes, if thats enough to make your requests work.

With this strategy, youll deplete tokens as quickly as possible, and then make subsequent requests to actively monitor the rate limiting status on the server side. It guarantees youll get a 429 if your rate is above the limits.

2. Specific timed backoff

This strategy is a bit more specific, as its using the `retry-after` header. Were considering this header an industry standard and plan to use it across the Atlassian suite, so you can still be sure the same code will work for Bitbucket and Confluence, Server and Cloud, etc. This strategy makes sure that you will not be limited, because youll know exactly how long you need to wait before youre allowed to make new requests.

Universal, works with any rate limiting system within the Atlassian suite (and other products using `retry-after`) Bitbucket and Confluence, Server and Cloud, etc.

Doesnt require too much knowledge about limits or a rate limiting system.

High impact on a Confluence instance because of concurrency. Were assuming most active users will send requests whenever theyre available. This window will be similar for all users, making spikes in Jira performance. The same applies to threads most will either be busy at the same time or idle.

Summary of this strategy

Heres the high-level overview of how to adjust your code:

1. **Active**: Make requests and observe the `retry-after` response header, which shows the number of seconds you need to wait to get new tokens. Keep concurrency level to a minimum to know exactly when the rate limit kicks in.
    a. **Success**: If the header says 0, you can make more requests right away.
    b. **Limited**:If the header has a number greater than 0, for example 5, you need to wait that number of seconds.
2. **Timeout**: If the header is anything above 0, start the timeout with the number of seconds specified in the header. Consider increasing the timeout by a random fraction, up to 20%.
3. **Retry**: After the timeout specified in the header has passed, go back to step 1 and make more requests.

With this strategy, youll deplete tokens as quickly as possible, and then pause until you get new tokens. You should never hit a 429 if your code is the only agent depleting tokens and sends requests synchronously.

3. Rate adjustment

This strategy is very specific and expects particular response headers, so its most likely to work for Confluence Data Center only. When making requests, youll observe headers returned by the server (number of tokens, fill rate, time interval) and adjust your code specifically to the number of tokens you have and can use.

It can have the least performance impact on a Confluence instance, if used optimally.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Highly recommended, especially for integrations that require high-volume traffic.

Safe, as you can easily predict that all requests that must go through will in fact go through. It also allows for a great deal of customization.

Very specific, depends on specific headers and rate limiting system.

Summary of this strategy

Heres the high-level overview of how to adjust your code:

1. **Active:** Make requests and observe all response headers.
2. **Adjust:** With every request, recalculate the rate based on the following headers:
    a. `x-ratelimit-interval-seconds`: The time interval in seconds. You get a batch of new tokens every time interval.
    b. `x-ratelimit-fillrate`: The number of tokens you get every time interval.
    c. `retry-after`: The number of seconds you need to wait for new tokens. Make sure that your rate assumes waiting longer than this value.
3. **Retry:** If you encounter a 429, which shouldnt happen if you used the headers correctly, you need to further adjust your code so it doesnt happen again. You can use the `retry-after` header to make sure that you only make requests when the tokens are available.

Customizing your code

Depending on your needs, this strategy helps you to:

By following the headers, you should know how many tokens you have, when you will get the new ones, and in what number. The most useful headers here are `x-ratelimit-interval-seconds` and `x-ratelimit-fillrate`, which show the number of tokens available every time interval. They help you choose the perfect frequency of making your requests.

You can wait to perform complex operations until youre sure you have enough tokens to make all the consecutive requests you need to make. This allows you to reduce the risk of leaving the system in an inconsistent state, for example when your task requires 4 requests, but it turns out you can only make 2. The most useful headers are `x-ratelimit-remaining` and `x-ratelimit-interval-seconds`, which show how many tokens you have right now and how long you need to wait for the new ones.

With all the information returned by the headers, you can create more strategies that work best for you, or mix the ones weve described here. For example:

If youre making requests once a day, you can focus on the max requests you can accumulate (`x-ratelimit-limit`), or lean towards the remaining number of tokens if a particular action in Confluence triggers your app to make requests (`x-ratelimit-remaining`).

If your script needs to work both for Confluence Data Center and some other application, use all headers for Confluence and focus on the universal retry-after or request codes if the app detects different software.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.

Running Confluence Data Center on a single node

Data Center allows you to run Confluence in a cluster with multiple nodes, or on a single server (also known as non-clustered, or standalone Data Center).

This page outlines the architecture and requirements of a non-clustered Confluence Data Center deployment, as well as some of the benefits and considerations.

Architecture

The deployment architecture of a non-clustered Data Center deployment is the same as a Server installation. Heres what a typical setup looks like:

As you can see, Confluence Data Center deployed on a single node looks just as a Server installation, and consists of:

- Confluence Data Center, running on a single node
- A database that Confluence reads and writes to

SeeGetting started as a Confluence administrator to learn more about single server Confluence installations.

Requirements

Non-clustered Confluence Data Center installations have the same minimum requirements as a Confluence Server installation. Check our Confluence System Requirements guide for a full overview of the supported platforms and hardware youll need.

Benefits of running a non-clustered Data Center deployment

There are a range of reasons you may choose a single node Data Center. Some of the benefits include:

- **Keeping your existing infrastructure**
  Running on a single node means that you can upgrade from Server to Data Center without adding to your infrastructure. In most cases, moving to Data Center will be as simple as updating your license.

- **Accessing Data Center-only features**
  Your Data Center license unlocks a suite of additional security, compliance, and administration features to help you easily manage enterprise-grade Confluence site like SAML single sign-on, advanced permission management, rate limiting, and more. See the complete list.

As non-clustered Confluence Data Center installations are cluster-compatible, you can still enable and configure clustering whenever youre ready to scale. Learn more about setting up a cluster.

Considerations

Non-clustered Data Center is the simplest setup, but it has some limitations. Just like a Server installation, youll still have the application server as a single point of failure, so it cant support high availability or disaster recovery strategies.

Some deployments start to experience performance or stability issues once their size profile hits Large or XLarge. Most clustered deployments provide you the flexibility to scale up your infrastructure to address heavy loads (or even scale down to save costs during light loads). On AWS or Azure, you can also quickly address most stability issues by replacing misbehaving nodes with fresh ones.

For more information about size profiles, see Data Center performance sizing. We also explain our own strategies for managing our clustered deployments in How Atlassians monitor their enterprise deployments.

Created in 2023 by Atlassian. Licensed under a Creative Commons Attribution 2.5 Australia License.