

1. Confluence Documentation Home . . . 10
1.1 Confluence 101 . . . 10
1.2 Confluence Administrator's Guide . . . 13
1.2.1 Administration . . . 16
1.2.1.1 Cache Statistics . . . 17
1.2.1.2 Changing time of Daily Backup . . . 19
1.2.1.3 Confluence Data Directory Configuration . . . 20
1.2.1.4 Confluence home directory contents . . . 20
1.2.1.5 Content Index Administration . . . 22
1.2.1.6 Finding Unused Spaces . . . 25
1.2.1.7 Important Directories and Files . . . 26
1.2.1.7.1 Confluence Home Directory . . . 28
1.2.1.7.2 Confluence Installation Directory . . . 28
1.2.1.8 Installing a Language Pack . . . 28
1.2.1.9 Manually Backing Up The Site . . . 30
1.2.1.9.1 Configuring Daily Backups . . . 31
1.2.1.10 Migrating Confluence Between Servers . . . 33
1.2.1.11 Rebuilding the Ancestor Table . . . 35
1.2.1.12 Restoring a Site . . . 36
1.2.1.13 Restoring a Space . . . 36
1.2.1.13.1 Changing the version of a space backup . . . 37
1.2.1.14 Restoring a Test Instance from Production . . . 38
1.2.1.15 Restoring Data from other Backups . . . 38
1.2.1.16 Restoring Data from the Administration Console . . . 39
1.2.1.17 Retrieve file attachments from a backup . . . 40
1.2.1.18 Troubleshooting failed XML site backups . . . 41
1.2.1.19 Troubleshooting XML backups that fail on restore . . . 43
1.2.1.19.1 Migrating from HSQLDB to MySQL . . . 45
1.2.1.20 Viewing and Editing License Details . . . 49
1.2.1.21 Viewing System Information . . . 51
1.2.1.21.1 Live Monitoring Using the JMX Interface . . . 51
1.2.1.21.2 Site Statistics . . . 53
1.2.1.22 Installing Patched Class Files . . . 54
1.2.2 Configuring Confluence . . . 55
1.2.2.1 Site Configuration . . . 55
1.2.2.1.1 Editing the Site Title . . . 56
1.2.2.1.2 Editing the Global Logo . . . 56
1.2.2.1.3 Configuring the Site Home Page . . . 57
1.2.2.1.4 Editing the Site Welcome Message . . . 58
1.2.2.1.5 Configuring the Server Base URL . . . 62
1.2.2.1.6 Configuring the Site Support Address . . . 63
1.2.2.1.7 Configuring the Destination of View Space Links . . . 63
1.2.2.1.8 Customising Default Space Content . . . 64
1.2.2.1.9 Showing Link Icons . . . 65
1.2.2.2 Configuring Encoding . . . 66

1.2.2.2.1 Character encodings in Confluence . . . 66
1.2.2.2.2 Troubleshooting Character Encodings . . . 67
1.2.2.3 Configuring Mail . . . 71
1.2.2.3.1 Configuring a Server for Outgoing Mail . . . 71
1.2.2.3.2 Enabling the 'Mail Page' plugin . . . 72
1.2.2.3.3 The Mail Queue . . . 73
1.2.2.4 Optional Settings . . . 73
1.2.2.4.1 Attachment Storage Configuration . . . 73
1.2.2.4.2 Configuring Quick Navigation . . . 79
1.2.2.4.3 Enabling CamelCase Linking . . . 80
1.2.2.4.4 Enabling OpenSearch . . . 81
1.2.2.4.5 Enabling Remote APIs . . . 81
1.2.2.4.6 Enabling Rich Text Editing Option . . . 81
1.2.2.4.7 Enabling the Did You Mean Feature . . . 82
1.2.2.4.8 Enabling Threaded Comments . . . 82
1.2.2.4.9 Enabling Trackback . . . 82
1.2.2.4.10 Making Rich Text Editing default . . . 83
1.2.2.4.11 WebDAV Configuration . . . 83
1.2.2.5 Other Settings . . . 86
1.2.2.5.1 Configuring Attachment Size . . . 86
1.2.2.5.2 Configuring Character Encoding . . . 87
1.2.2.5.3 Configuring HTTP Timeout Settings . . . 88
1.2.2.5.4 Configuring Indexing Language . . . 88
1.2.2.5.5 Configuring JIRA Issues Icon mappings . . . 89
1.2.2.5.6 Configuring Number Formats . . . 89
1.2.2.5.7 Configuring Shortcut Links . . . 90
1.2.2.5.8 Configuring Time and Date Formats . . . 91
1.2.2.5.9 Number of Ancestors to Show in Breadcrumbs . . . 92
1.2.2.5.10 Thumbnail Settings . . . 92
1.2.2.6 Configuring System Properties . . . 93
1.2.2.6.1 Recognised System Properties . . . 94
1.2.2.7 Configuring Logging . . . 96
1.2.2.8 External Gadgets . . . 98
1.2.3 Confluence and JIRA . . . 99
1.2.3.1 Add Confluence EAR-WAR to JIRA Standalone . . . 99
1.2.3.2 Integrating JIRA and Confluence . . . 100
1.2.3.3 Override properties in JIRA to Confluence Bridge . . . 102
1.2.3.4 Setting Up Trusted Communication between JIRA and Confluence . . . 103
1.2.4 Confluence Clustering Overview . . . 108
1.2.4.1 Technical Overview of Clustering in Confluence . . . 109
1.2.4.1.1 Cluster safety mechanism . . . 112
1.2.4.2 Changing Datasources Manually in a Cluster . . . 112
1.2.4.3 Cluster Troubleshooting . . . 113
1.2.4.3.1 Cluster Panic triggers . . . 118
1.2.4.3.2 Multicast Test . . . 118
1.2.4.4 Clustering for Scalability vs Clustering for High Availability (HA) . . . 118
1.2.4.5 Recommended network topology . . . 120
1.2.4.6 Cluster Administration page . . . 121
1.2.4.7 Cluster Checklist . . . 123
1.2.5 Confluence Security . . . 127
1.2.5.1 Confluence Community Security Advisory 2006-01-19 . . . 129
1.2.5.2 Confluence Security Advisory 2005-02-09 . . . 130
1.2.5.3 Confluence Security Advisory 2005-12-05 . . . 131
1.2.5.4 Confluence Security Advisory 2006-01-20 . . . 131
1.2.5.5 Confluence Security Advisory 2006-01-23 . . . 132
1.2.5.6 Confluence Security Advisory 2006-06-14 . . . 132
1.2.5.7 Confluence Security Advisory 2007-07-26 . . . 132
1.2.5.8 Confluence Security Advisory 2007-08-08 . . . 133
1.2.5.9 Confluence Security Advisory 2007-11-19 . . . 134
1.2.5.10 Confluence Security Advisory 2007-11-27 . . . 135
1.2.5.11 Confluence Security Advisory 2007-12-14 . . . 136
1.2.5.12 Confluence Security Advisory 2008-01-24 . . . 137
1.2.5.13 Confluence Security Advisory 2008-03-06 . . . 138
1.2.5.14 Confluence Security Advisory 2008-03-19 . . . 139
1.2.5.15 Confluence Security Advisory 2008-05-21 . . . 140
1.2.5.16 Confluence Security Advisory 2008-07-03 . . . 141
1.2.5.17 Confluence Security Advisory 2008-09-08 . . . 142
1.2.5.18 Confluence Security Advisory 2008-10-14 . . . 145
1.2.5.19 Confluence Security Advisory 2008-12-03 . . . 148
1.2.5.20 Confluence Security Advisory 2009-01-07 . . . 149
1.2.5.21 Confluence Security Advisory 2009-02-18 . . . 150
1.2.5.22 Confluence Security Advisory 2009-04-15 . . . 151
1.2.5.23 Confluence Security Advisory 2009-06-01 . . . 153
1.2.5.24 Confluence Security Advisory 2009-06-16 . . . 154
1.2.5.25 Confluence Security Advisory 2009-08-20 . . . 154
1.2.5.26 Confluence Security Advisory 2009-10-06 . . . 156
1.2.5.27 Confluence Security Advisory 2009-12-08 . . . 158
1.2.6 Design and Layout . . . 158
1.2.6.1 Choosing a Default Language . . . 159
1.2.6.2 Custom Decorator Templates . . . 159

1.2.6.3 Customising Look and Feel Overview . . . 162
1.2.6.3.1 Customising Colour Schemes . . . 162
1.2.6.3.2 Customising Layouts . . . 164
1.2.6.3.3 Global Templates . . . 171
1.2.6.3.4 Importing Templates . . . 172
1.2.6.3.5 Modify Confluence Interface Text . . . 173
1.2.6.3.6 Working With Decorator Macros . . . 174
1.2.6.3.7 Customising a Specific Page . . . 175
1.2.6.3.8 Customising PDF or HTML Content . . . 175
1.2.6.3.9 Customising the Dashboard . . . 175
1.2.6.3.10 Customising the eMail Templates . . . 176
1.2.6.3.11 Customising the Login Page . . . 176
1.2.6.4 Themes Overview . . . 177
1.2.6.4.1 Applying a Theme to a Site . . . 177
1.2.6.4.2 Customising the Left Navigation Theme . . . 179
1.2.6.4.3 Modifying Look and Feel (for themes) . . . 180
1.2.6.4.4 Creating a Theme . . . 186
1.2.7 Importing Data . . . 187
1.2.7.1 Importing Content from another Wiki . . . 187
1.2.7.2 Snip Snap Import . . . 188
1.2.8 Installing Plugins and Macros . . . 189
1.2.8.1 Installing and Configuring Plugins Manually . . . 190
1.2.8.2 Installing and Configuring Plugins using the Plugin Repository Client . . . 192
1.2.8.3 Plugin loading strategies in Confluence . . . 194
1.2.8.4 Removing Malfunctioning Plugins . . . 195
1.2.8.5 Enabling and Configuring Macros . . . 196
1.2.8.5.1 Adding, Editing and Removing User Macros . . . 196
1.2.8.5.2 Configuring a URL Whitelist . . . 197
1.2.8.5.3 Configuring the userlister Macro . . . 199
1.2.8.5.4 Enabling HTML macros . . . 199
1.2.8.5.5 Troubleshooting the Gallery Macro . . . 201
1.2.8.6 Configuring the Office Connector . . . 201
1.2.9 Operating Large or Mission-Critical Confluence Installations . . . 205
1.2.10 Performance Tuning . . . 212
1.2.10.1 Cache Performance Tuning . . . 215
1.2.10.1.1 Cache Performance Tuning for Specific Problems . . . 218
1.2.10.1.2 Confluence Cache Schemes . . . 219
1.2.10.2 Configuring a Large Confluence Installation . . . 220
1.2.10.3 How Adaptavist Runs Confluence . . . 221
1.2.10.4 Memory usage and requirements . . . 222
1.2.10.5 Requesting Performance Support . . . 224
1.2.10.5.1 Access Log Scripts . . . 226
1.2.10.5.2 Obtaining Confluence Instance Metrics . . . 227
1.2.10.6 Troubleshooting Slow Performance Using Page Request Profiling . . . 230
1.2.10.7 Compressing an HTTP Response within Confluence . . . 233
1.2.10.8 Performance Testing Scripts . . . 234
1.2.10.9 Garbage Collector Performance Issues . . . 238
1.2.11 Scheduled Jobs . . . 243
1.2.12 Search . . . 243
1.2.12.1 Setup Confluence To Index External Sites . . . 243
1.2.12.2 Setup External Search Tool To Index Confluence . . . 244
1.2.13 Security . . . 245
1.2.13.1 Adding SSL for Secure Logins and Page Security . . . 245
1.2.13.2 Anonymous Access to Remote API . . . 248
1.2.13.3 Confluence Cookies . . . 248
1.2.13.4 Enabling or Disabling Public Signup . . . 250
1.2.13.5 Hiding External Links From Search Engines . . . 250
1.2.13.6 Hiding the People Directory . . . 251
1.2.13.7 Managing External Referrers . . . 252
1.2.13.7.1 Excluding external referrers . . . 253
1.2.13.7.2 Hiding external referrers . . . 254
1.2.13.7.3 Ignoring External Referrers . . . 255
1.2.13.8 User Email Visibility . . . 256
1.2.13.9 Spam Prevention via Captcha . . . 257
1.2.14 Working with Confluence Logs . . . 257
1.2.14.1 log4j Logging Levels . . . 259
1.2.15 Configuring OAuth . . . 259
1.2.16 User Management . . . 261
1.2.16.1 Confluence User Management . . . 261
1.2.16.1.1 Searching For and Managing Users . . . 261
1.2.16.1.2 Adding a Group . . . 265
1.2.16.1.3 Adding a New User . . . 265
1.2.16.1.4 Adding or Removing Users in Groups . . . 266
1.2.16.1.5 Changing Usernames . . . 269
1.2.16.1.6 Editing User Details . . . 272
1.2.16.1.7 Global Groups Overview . . . 273
1.2.16.1.8 Global Permissions Overview . . . 274
1.2.16.1.9 Migrating to new User Management . . . 277
1.2.16.1.10 Removing a Group . . . 278
1.2.16.1.11 Removing a User . . . 279

1.2.16.1.12 Setting up Anonymous Access . . . 280
1.2.16.1.13 Viewing members of a group . . . 281
1.2.16.1.14 How to Improve User Search Performance . . . 281
1.2.16.1.15 Restoring Passwords To Recover Admin User Rights . . . 282
1.2.16.2 Integrating with Crowd . . . 283
1.2.16.3 JIRA User Management . . . 283
1.2.16.3.1 Delegate user management to use JIRA logins . . . 283
1.2.16.4 LDAP User Management . . . 290
1.2.16.4.1 Add LDAP Integration . . . 290
1.2.16.4.2 Add LDAP Integration For User Authentication Only . . . 297
1.2.16.4.3 atlassian-user.xml reference . . . 298
1.2.16.4.4 Changes in osuser.xml from 1.0.3a to 1.1.x . . . 300
1.2.16.4.5 Configuring multiple LDAP repositories . . . 300
1.2.16.4.6 Connect to LDAP, JIRA or Other Services Via SSL . . . 303
1.2.16.4.7 Disabling the Built-In User Management . . . 304
1.2.16.4.8 Legacy User Management Documentation . . . 305
1.2.16.4.9 Troubleshooting LDAP User Management . . . 308
1.2.16.4.10 Troubleshooting the "Not Permitted" Screen under LDAP Integration . . .
1.2.16.5 Migrating users from Confluence to JIRA . . . 314
1.2.16.6 Requesting External User Management Support . . . 315
1.2.16.6.1 Paddle . . . 316
1.2.16.7 Understanding User Management in Confluence . . . 316
1.2.16.8 User Management Frequently Asked Questions . . . 319
1.3 Confluence Installation and Upgrade Guide . . . 319
1.3.1 System Requirements . . . 319
1.3.1.1 Server Hardware Requirements Guide . . . 322
1.3.1.2 Example Size & Hardware Specifications From Customer Survey . . . 325
1.3.2 Supported Platforms . . . 326
1.3.2.1 End of Support Announcements for Confluence . . . 327
1.3.2.1.1 Java 1.4 Support Timeline . . . 329
1.3.2.2 Supported Platforms FAQ . . . 330
1.3.3 Confluence Installation Guide . . . 331
1.3.3.1 Installing Confluence Standalone . . . 332
1.3.3.1.1 Change listen port for Confluence Standalone . . . 332
1.3.3.1.2 Installing Confluence Standalone on Mac OS X from Zip File . . . 333
1.3.3.1.3 Installing Confluence Standalone on Unix or Linux . . . 336
1.3.3.1.4 Installing Confluence Standalone on Windows from Zip File . . . 340
1.3.3.1.5 Installing Confluence Standalone Using the Windows Evaluation Installer . . . 342
1.3.3.1.6 Installing Confluence Standalone Using the Mac OS X Evaluation Installer . . . 346
1.3.3.2 Installing the Confluence EAR-WAR Edition . . . 348
1.3.3.2.1 Installing Confluence EAR-WAR on Tomcat . . . 349
1.3.3.2.2 Installing Confluence EAR-WAR on Weblogic . . . 351
1.3.3.2.3 Installing Confluence EAR-WAR on Websphere . . . 355
1.3.3.2.4 Installing Confluence EAR-WAR on Resin . . . 358
1.3.3.3 Confluence Cluster Installation . . . 359
1.3.3.3.1 Apache and Tomcat load balancing . . . 361
1.3.3.4 Confluence Cluster Installation with Existing Data . . . 364
1.3.3.4.1 Upgrading a Confluence Cluster . . . 365
1.3.3.5 Installing Sun JDK for Confluence . . . 367
1.3.3.5.1 Setting the JAVA_HOME Variable in Windows . . . 367
1.3.3.6 Confluence Unix and X11 Dependencies . . . 368
1.3.3.7 Get A Confluence Licence . . . 369
1.3.3.8 Running Confluence in a Virtualised Environment . . . 369
1.3.3.9 Uninstalling Confluence Standalone . . . 373
1.3.4 Confluence Setup Guide . . . 373
1.3.4.1 External Database . . . 379
1.3.4.2 Load Content for the Site . . . 380
1.3.4.3 Restoring from Backup During Setup . . . 381
1.3.5 Upgrading Confluence . . . 382
1.3.5.1 Upgrading Confluence Standalone Distribution . . . 383
1.3.5.2 Upgrading Confluence EAR-WAR Distribution . . . 386
1.3.5.3 Upgrading Beyond Current Licensed Period . . . 389
1.3.6 Confluence Release Cycle . . . 390
1.3.6.1 Release Notes . . . 391
1.3.6.1.1 Confluence Release Summary . . . 393
1.3.6.1.2 Confluence 3.2 Release Notes . . . 396
1.3.6.1.3 Confluence 3.1.2 Release Notes . . . 410
1.3.6.1.4 Confluence 3.1.1 Release Notes . . . 411
1.3.6.1.5 Confluence 3.1 Release Notes . . . 414
1.3.6.1.6 Confluence 3.0.2 Release Notes . . . 438
1.3.6.1.7 Confluence 3.0.1 Release Notes . . . 441
1.3.6.1.8 Confluence 3.0 Release Notes . . . 446
1.3.6.1.9 Confluence 2.10.4 Release Notes . . . 485
1.3.6.1.10 Confluence 2.10.3 Release Notes . . . 485
1.3.6.1.11 Confluence 2.10.2 Release Notes . . . 488
1.3.6.1.12 Confluence 2.10.1 Release Notes . . . 491
1.3.6.1.13 Confluence 2.10 Release Notes . . . 492
1.3.6.1.14 Confluence 2.9.3 Release Notes . . . 506
1.3.6.1.15 Confluence 2.9.2 Release Notes . . . 506
1.3.6.1.16 Confluence 2.9.1 Release Notes . . . 508

1.3.6.1.17 Confluence 2.9 Release Notes . . . 511
1.3.6.1.18 Confluence 2.8.3 Release Notes . . . 528
1.3.6.1.19 Confluence 2.8.2 Release Notes . . . 528
1.3.6.1.20 Confluence 2.8.1 Release Notes . . . 530
1.3.6.1.21 Confluence 2.8 Release Notes . . . 532
1.3.6.1.22 Confluence 2.8 Beta Release Notes . . . 547
1.3.6.1.23 Confluence 2.7.4 Release Notes . . . 548
1.3.6.1.24 Confluence 2.7.3 Release Notes . . . 548
1.3.6.1.25 Confluence 2.7.2 Release Notes . . . 550
1.3.6.1.26 Confluence 2.7.1 Release Notes . . . 551
1.3.6.1.27 Confluence 2.7 Release Notes . . . 553
1.3.6.1.28 Confluence 2.6.3 Release Notes . . . 565
1.3.6.1.29 Confluence 2.6.2 Release Notes . . . 566
1.3.6.1.30 Confluence 2.6.1 Release Notes . . . 568
1.3.6.1.31 Confluence 2.6 Release Notes . . . 570
1.3.6.1.32 Release Notes 1.0 . . . 584
1.3.6.1.33 Release Notes 1.0.1 . . . 588
1.3.6.1.34 Release Notes 1.0.3 . . . 589
1.3.6.1.35 Release Notes 1.0a2 . . . 591
1.3.6.1.36 Release Notes 1.0a3 . . . 591
1.3.6.1.37 Release Notes 1.0b1 . . . 591
1.3.6.1.38 Release Notes 1.0b2 . . . 592
1.3.6.1.39 Release Notes 1.0b3 . . . 593
1.3.6.1.40 Release Notes 1.0b4 . . . 596
1.3.6.1.41 Release Notes 1.0rc1 . . . 599
1.3.6.1.42 Release Notes 1.0rc2 . . . 602
1.3.6.1.43 Release Notes 1.0rc5 . . . 606
1.3.6.1.44 Release Notes 1.0rc6 . . . 610
1.3.6.1.45 Release Notes 1.1 . . . 614
1.3.6.1.46 Release Notes 1.1.1 . . . 618
1.3.6.1.47 Release Notes 1.1.2 . . . 620
1.3.6.1.48 Release Notes 1.2 . . . 620
1.3.6.1.49 Release Notes 1.2.1 . . . 627
1.3.6.1.50 Release Notes 1.2.2 . . . 629
1.3.6.1.51 Release Notes 1.2.3 . . . 629
1.3.6.1.52 Release Notes 1.3 . . . 631
1.3.6.1.53 Release Notes 1.3.1 . . . 645
1.3.6.1.54 Release Notes 1.3.2 . . . 646
1.3.6.1.55 Release Notes 1.3.4 . . . 647
1.3.6.1.56 Release Notes 1.3.5 . . . 648
1.3.6.1.57 Release Notes 1.3.6 . . . 650
1.3.6.1.58 Release Notes 1.4 . . . 650
1.3.6.1.59 Release Notes 1.4.1 . . . 652
1.3.6.1.60 Release Notes 1.4.2 . . . 653
1.3.6.1.61 Release Notes 1.4.3 . . . 655
1.3.6.1.62 Release Notes 1.4.4 . . . 659
1.3.6.1.63 Release Notes 2.0 . . . 661
1.3.6.1.64 Release Notes 2.0.1 . . . 668
1.3.6.1.65 Release Notes 2.0.2 . . . 670
1.3.6.1.66 Release Notes 2.0.3 . . . 672
1.3.6.1.67 Release Notes 2.1 . . . 673
1.3.6.1.68 Release Notes 2.1.1 . . . 677
1.3.6.1.69 Release Notes 2.1.2 . . . 678
1.3.6.1.70 Release Notes 2.1.3 . . . 680
1.3.6.1.71 Release Notes 2.1.4 . . . 682
1.3.6.1.72 Release Notes 2.1.5 . . . 686
1.3.6.1.73 Release Notes 2.2 . . . 689
1.3.6.1.74 Release Notes 2.2.1 . . . 697
1.3.6.1.75 Release Notes 2.2.2 . . . 700
1.3.6.1.76 Release Notes 2.2.3 . . . 702
1.3.6.1.77 Release Notes 2.2.4 . . . 703
1.3.6.1.78 Release Notes 2.2.5 . . . 704
1.3.6.1.79 Release Notes 2.2.6a . . . 705
1.3.6.1.80 Release Notes 2.2.7 . . . 707
1.3.6.1.81 Release Notes 2.2.8 . . . 709
1.3.6.1.82 Release Notes 2.2.9 . . . 711
1.3.6.1.83 Release Notes 2.2.10 . . . 712
1.3.6.1.84 Release Notes 2.3 . . . 715
1.3.6.1.85 Release Notes 2.3.1 . . . 727
1.3.6.1.86 Release Notes 2.3.2 . . . 729
1.3.6.1.87 Release Notes 2.3.3 . . . 730
1.3.6.1.88 Release Notes 2.4 . . . 730
1.3.6.1.89 Release Notes 2.4.1 . . . 733
1.3.6.1.90 Release Notes 2.4.2 . . . 735
1.3.6.1.91 Release Notes 2.4.3 . . . 739
1.3.6.1.92 Release Notes 2.4.4 . . . 740
1.3.6.1.93 Release Notes 2.4.5 . . . 741
1.3.6.1.94 Release Notes 2.5 . . . 742
1.3.6.1.95 Release Notes 2.5.1 . . . 745
1.3.6.1.96 Release Notes 2.5.2 . . . 747

1.3.6.1.97 Release Notes 2.5.3 . . . 748
1.3.6.1.98 Release Notes 2.5.4 . . . 749
1.3.6.1.99 Release Notes 2.5.5 . . . 751
1.3.6.1.100 Release Notes 2.5.6 . . . 753
1.3.6.1.101 Release Notes 2.5.7 . . . 754
1.3.6.1.102 Release Notes 2.5.8 . . . 755
1.3.6.1.103 Release Notes_1.0.1 . . . 756
1.3.6.2 Development Releases . . . 757
1.3.6.2.1 Confluence 3.1 Newly Deprecated Code . . . 759
1.3.6.2.2 Development Release Warnings . . . 759
1.3.6.2.3 Milestone release advisory . . . 759
1.3.6.2.4 Beta Release Advisory . . . 760
1.3.6.2.5 Release Notes 1.3-DR1 . . . 761
1.3.6.2.6 Release Notes 1.3-DR2 . . . 761
1.3.6.2.7 Release Notes 1.3-DR3 . . . 763
1.3.6.2.8 Release Notes 1.3-DR4 . . . 764
1.3.6.2.9 Release Notes 1.3-final . . . 766
1.3.6.2.10 Release Notes 1.4-DR1 . . . 767
1.3.6.2.11 Release Notes 1.4-DR2 . . . 769
1.3.6.2.12 Release Notes 1.4-DR3 . . . 775
1.3.6.2.13 Release Notes 1.4-DR4 . . . 778
1.3.6.2.14 Release Notes 1.4-DR6 . . . 779
1.3.6.2.15 Release Notes 1.4-DR7 . . . 781
1.3.6.2.16 Release Notes 1.4-RC2 . . . 782
1.3.6.2.17 Release Notes 1.5-DR1 . . . 782
1.3.6.2.18 Release Notes 1.5-DR2 . . . 785
1.3.6.2.19 Release Notes 2.0-RC1 . . . 790
1.3.6.2.20 Release Notes 2.0-RC2 . . . 791
1.3.6.2.21 Release Notes 2.3-DR1 . . . 792
1.3.6.2.22 Release Notes 2.3-DR2 . . . 793
1.3.6.2.23 Release Notes 2.6-dr1 . . . 794
1.3.6.2.24 Release Notes 2.7-m2 ("Milestone 2") . . .
1.3.6.2.25 Release Notes 2.7-m4 ("Milestone 4") . . .
1.3.6.2.26 Release Notes 2.7-m5 ("Milestone 5") . . .
1.3.6.2.27 Release Notes 2.7-rc1 ("Release Candidate 1") . . .
1.3.6.2.28 Release Notes 2.8-m2 ("Milestone 2") . . .
1.3.6.2.29 Release Notes 2.8-m3 ("Milestone 3") . . .
1.3.6.2.30 Release Notes 2.8-m4 ("Milestone 4") . . .
1.3.6.2.31 Release Notes 2.8-m6 ("Milestone 6") . . .
1.3.6.2.32 Release Notes 2.8-m7 ("Milestone 7") . . .
1.3.6.2.33 Release Notes 2.8-m9 ("Milestone 9") . . .
1.3.6.2.34 Release Notes 2.9-m2 ("Milestone 2") . . .
1.3.6.2.35 Release Notes 2.9-m3 ("Milestone 3") . . .
1.3.6.2.36 Release Notes 2.9-m5 ("Milestone 5") . . .
1.3.6.2.37 Release Notes 2.9-rc1 ("Release Candidate 1") . . .
1.3.6.2.38 Release Notes 2.10-m1 ("Milestone 1") . . .
1.3.6.2.39 Release Notes 2.10-m2 ("Milestone 2") . . .
1.3.6.2.40 Release Notes 2.10-m3 ("Milestone 3") . . .
1.3.6.2.41 Release Notes 2.10-m4 ("Milestone 4") . . .
1.3.6.2.42 Release Notes 2.10-m5 ("Milestone 5") . . .
1.3.6.2.43 Release Notes 2.10-m7 ("Milestone 7") . . .
1.3.6.2.44 Release Notes 2.10-m8 ("Milestone 8") . . .
1.3.6.2.45 Release Notes 2.10-rc1 ("Release Candidate 1") . . .
1.3.6.2.46 Release Notes 3.0-beta2 ("Beta 2") . . .
1.3.6.2.47 Release Notes 3.0-m3 ("Milestone 3") . . .
1.3.6.2.48 Release Notes 3.0-m4 ("Milestone 4") . . .
1.3.6.2.49 Release Notes 3.0-m5 ("Milestone 5") . . .
1.3.6.2.50 Release Notes 3.0-m6 ("Milestone 6") . . .
1.3.6.2.51 Release Notes 3.0-m7 ("Milestone 7") . . .
1.3.6.2.52 Release Notes 3.0-m8 ("Milestone 8") . . .
1.3.6.2.53 Release Notes 3.0-m9 ("Milestone 9") . . .
1.3.6.2.54 Release Notes 3.0-rc1 ("Release Candidate 1") . . .
1.3.6.2.55 Release Notes 3.1-m1 ("Milestone 1") . . .
1.3.6.2.56 Release Notes 3.1-m3 ("Milestone 3") . . .
1.3.6.2.57 Release Notes 3.1-m4 ("Milestone 4") . . .
1.3.6.2.58 Release Notes 3.1-m5 ("Milestone 5") . . .
1.3.6.2.59 Release Notes 3.1-m6 ("Milestone 6") . . .
1.3.6.2.60 Release Notes 3.1-m7 ("Milestone 7") . . .
1.3.6.2.61 Confluence 3.1 Beta 1 Release Notes . . . 872
1.3.6.2.62 Confluence 3.1 Beta 2 Release Notes . . . 879
1.3.6.2.63 Confluence 3.1 Release Candidate 1 Release Notes . . . 888
1.3.6.2.64 Release Notes 3.2-m3 ("Milestone 3") . . .
1.3.6.2.65 Release Notes 3.2-m4 ("Milestone 4") . . .
1.3.6.2.66 Confluence 3.2 Beta Release Notes . . . 900
1.3.6.3 Coherence license changes SEPT 2009 - new Standard and Clustered Confluence Editions . . . 910
1.3.6.4 Upgrade Notes Overview . . . 911
1.4 Confluence Configuration Guide . . . 912
1.4.1 Application Server URL encoding . . . 913
1.4.1.1 Configuring Tomcat's URI encoding . . . 913
1.4.2 Database Configuration . . . 913

1.4.2.1 Migrate to Another Database . . . 914
1.4.2.1.1 The Embedded HSQLDB Database . . . 916
1.4.2.2 Database Setup Guides . . . 917
1.4.2.2.1 Database Setup For Any External Database . . . 917
1.4.2.2.2 Database Setup for PostgreSQL . . . 918
1.4.2.2.3 Database Setup For MySQL . . . 920
1.4.2.2.4 Database Setup for Oracle . . . 924
1.4.2.2.5 Database Setup for SQL Server . . . 927
1.4.2.2.6 Database JDBC drivers . . . 927
1.4.2.3 Creating Database Schema Manually . . . 928
1.4.2.4 Known Issues For Supported Databases . . . 928
1.4.2.4.1 Configuring Database Character Encoding . . . 929
1.4.2.4.2 Known Issues for DB2 . . . 931
1.4.2.4.3 Known Issues for MySQL . . . 933
1.4.2.4.4 Known Issues For Oracle . . . 936
1.4.2.4.5 Known Issues for PostgreSQL . . . 938
1.4.2.4.6 Known Issues For SQL Server . . . 938
1.4.2.4.7 Known Issues for Sybase Database . . . 942
1.4.2.5 Improving Database Performance . . . 942
1.4.2.5.1 Creating a Lowercase Page Title Index . . . 943
1.4.2.6 Troubleshooting External Database Connections . . . 945
1.4.2.6.1 Configuring database query timeout . . . 946
1.4.2.7 Troubleshooting the Embedded Database (hSQL DB) . . . 946
1.4.2.7.1 Connecting to HSQLDB using DBVisualizer . . . 946
1.4.2.8 Upgrading From HSQL 1.7.1 to 1.8 . . . 949
1.4.3 Guide to using Apache Tomcat's Virtual Hosts . . . 949
1.4.4 How to dump Active Directory data to a file . . . 950
1.4.5 Known Issues with Enterprise or Webhosting environments . . . 950
1.4.6 Managing Application Server Memory Settings . . . 950
1.4.7 Pull down RSS Feeds or use the Repository plugin through a web proxy . . . 951
1.4.8 Running Confluence behind Apache . . . 951
1.4.8.1 General Apache Configuration Notes . . . 952
1.4.8.2 Using Apache with mod_jk . . . 953
1.4.8.3 Using Apache with mod_proxy . . . 955
1.4.8.4 Using Apache with virtual hosts and mod_proxy . . . 957
1.4.9 Setting Up Public Access . . . 959
1.4.10 Setup a mail session in standalone version . . . 959
1.4.11 Start Confluence automatically on system startup . . . 959
1.4.11.1 Start Confluence automatically on Linux & Unix . . . 960
1.4.11.2 Start Confluence automatically on OS X using launchd . . . 963
1.4.11.3 Start Confluence automatically on Windows as a Service . . . 964
1.4.12 The Confluence data model . . . 967
1.4.13 Troubleshooting SQL Exceptions . . . 977
1.4.14 Webserver Configuration . . . 977
1.4.14.1 Apache and Apache Connector Tips . . . 978
1.4.14.2 Configure Web Proxy Support for Confluence . . . 978
1.5 Confluence FAQ . . . 979
1.5.1 Administration FAQ . . . 982
1.5.1.1 Add Spell Checking To Confluence . . . 983
1.5.1.2 Can I run multiple instances of Confluence & connect them to a central database? . . . 983
1.5.1.3 Capturing HTTP traffic using Wireshark or Fiddler . . . 983
1.5.1.4 Configuring Confluence to send email notifications . . . 984
1.5.1.5 Copy Or Rename A Space . . . 984
1.5.1.6 Customising Confluence Icons . . . 987
1.5.1.7 Disable public account signups . . . 987
1.5.1.8 Displaying System Properties . . . 987
1.5.1.9 Editing your database password . . . 988
1.5.1.10 Enable public anonymous access . . . 988
1.5.1.11 Fix Out of Memory errors by Increasing Available Memory . . . 988
1.5.1.11.1 Editing the Windows Registry . . . 990
1.5.1.12 Getting a License for a Staging Environment . . . 992
1.5.1.13 How can I retrieve a recently deleted space or page? . . . 992
1.5.1.14 How do I adjust the session timeout . . . 993
1.5.1.15 How do I change the space key? . . . 993
1.5.1.16 How do I check which spaces have email accounts . . . 993
1.5.1.17 How do I configure the Plugin Repository to update its plugins information offline? . . . 993
1.5.1.18 How do I disable indexing of attachments . . . 994
1.5.1.19 How Do I Find My License from the File System? . . . 995
1.5.1.20 How Do I Get More Statistics From Confluence? . . . 995
1.5.1.21 How Do I Identify Inactive Users in Confluence . . . 996
1.5.1.22 How do I know what Confluence version I am running? . . . 997
1.5.1.23 How do I make commenters added as page watchers automatically? . . . 997
1.5.1.24 How do I prevent personal spaces from being shown on the dashboard . . . 997
1.5.1.25 How do I Remove a User who has Content Created . . . 998
1.5.1.26 How do I Remove the Last Updated and Created By Text? . . . 998
1.5.1.27 How do I supress cluster warning message in confluence? . . . 998
1.5.1.28 How to Disable Emoticons . . . 998
1.5.1.29 How to display a banner like the Confluence Documentation space . . . 999
1.5.1.30 How to Force External Links to Open in a New Window . . . 999
1.5.1.31 How to get a Java Heap Dump . . . 100

1.5.1.32 How to Hide the Referrer 100
1.5.1.33 How to Restore Deactivated Users 100
1.5.1.34 How to run a SQL script on your database 100
1.5.1.35 How to Search Confluence for Uses of a Macro 100
1.5.1.36 How to turn on Debugging for indexing 100
1.5.1.37 List page- and space-related details for an attachment using the attachment's name 100
1.5.1.38 Migrate Confluence from one database to another 100
1.5.1.39 Need to Change Server ID for Test Install 100
1.5.1.40 Page Restrictions Performance Considerations 100
1.5.1.41 Passing username and password via URL 100
1.5.1.42 Preventing and Cleaning up SPAM 100
1.5.1.43 Rebuild the Content Indices from scratch 100
1.5.1.44 Redirect users to a page on login 100
1.5.1.45 Redirect users to a site-wide home page after a successful login 100
1.5.1.46 Restrict Attachments Based On File Type 100
1.5.1.47 Search for User Properties in the Database 100
1.5.1.48 Using Firebug Lite in Internet Explorer when browsing a Confluence page 100
1.5.1.49 What are the IP Address Ranges for Atlassian's Servers? 100
1.5.1.50 Where are the files that used to be in my Confluence installation directory? 100
1.5.1.51 Where are user macros stored? 100
1.5.2 Backup FAQ 100
1.5.2.1 Are there any scripts for backup creation and restore? 100
1.5.2.2 Backup will not import 100
1.5.2.3 Can Confluence be restored from a backup minus attachments? 100
1.5.2.4 Can XML backups be deleted automatically? 100
1.5.2.5 Does running a daily XML backup slow performance? 100
1.5.2.6 How can I reduce the space taken up by automatic backups? 100
1.5.2.7 How to Change the Version of a Space Backup 100
1.5.2.8 How to Find Attachments in Attachments Folder 100
1.5.2.9 Is it Possible to Store the Confluence Home Directory on a Network Share? 101
1.5.2.10 Providing MySQL database with Content Anonymised 101
1.5.3 Configuration FAQ 101
1.5.3.1 Running Tomcat on a Different Port 101
1.5.3.2 Change default font, color, or spacing in Confluence 101
1.5.3.3 How do I Change the Time of Daily Report Updates 101
1.5.3.4 How to audit Confluence - enabling user access logging 101
1.5.3.5 Share users and groups between Confluence and JIRA 101
1.5.3.6 How do I Modify the Frequency of Content Indexing 101
1.5.3.7 How do I change the default polling time for email in Confluence? 101
1.5.3.8 How do I Configure an Automatic Refresh of the Recently Updated List 101
1.5.3.9 How to Revert from Clustering to Single Node 101
1.5.3.10 Adding a Site-Wide Banner 101
1.5.3.11 Customise Confluence Page Exports 101
1.5.3.11.1 Available Velocity Context Objects in Exporters 101
1.5.3.11.2 Customise MS Word Exports 101
1.5.3.11.3 Customise PDF Exports 101
1.5.3.12 Editing the Footer 101
1.5.3.13 How do i completely remove the "Space Details" page from Confluence exports?
1.5.3.14 Where does Confluence store all its data? 101
1.5.3.15 Running Confluence Behind a Caching Proxy Server 101
1.5.3.16 I am trying to install Confluence but the demo-site.zip file is missing 101
1.5.3.17 How do I Disable Automatic Mail Polling? 101
1.5.3.18 Disabling Profile Pictures on the Recently Updated Dashboard 102
1.5.3.19 How to Disable Profile Pictures from the Recently Updated Section of the Dashboard 102
1.5.3.20 Remove Version from Footer 102
1.5.3.21 Configuring a MySQL Datasource in Apache Tomcat 102
1.5.4 Installation FAQ 102
1.5.4.1 The Confluence window closes immediately when started 102
1.5.4.2 How do I re-trigger the setup wizard 102
1.5.4.3 Confluence starts but a problem prevents me from accessing the dashboard 102
1.5.4.4 I receive a BUILD FAILED message when trying to create an EAR file in Confluence 2.6 or 2.7 102
1.5.4.5 How much disk space does Confluence need? 102
1.5.4.6 How Do I Make Confluence Accessible from the Root Context with a Tomcat EAR WAR configuration 102
1.5.4.7 How To Run Confluence Standalone and Apache on Port 80 (Different IP Addresses) 102
1.5.5 JIRA Integration FAQ 102
1.5.5.1 The JIRA Issues Macro behaves problematically 102
1.5.5.2 The JIRA Issues Macro generates an error 102
1.5.5.3 When setting up JIRA user management in Confluence, the Confluence login page loads but login fails 102
1.5.5.4 When setting up JIRA user management in Confluence, the Confluence login page loads with an error 102
1.5.6 LDAP FAQ 102
1.5.6.1 Are all users in LDAP visible in Confluence administration and can they can be assigned permissions or to groups? 102
1.5.6.2 Can we use LDAP and Confluence groups simultaneously, as a 'mixed mode', where some groups are kept in Confluence and others in LDAP? 102
1.5.6.3 Confluence Domino LDAP FAQs 102
1.5.6.4 Confluence integration with LDAP and Active Directory FAQs 102
1.5.6.5 How are LDAP or Active Directory users counted toward my license limit? 102
1.5.6.6 How can I assign an LDAP user a Confluence account? 102
1.5.6.7 How can I enable LDAP? 102
1.5.6.8 How does Confluence handle user deletions from LDAP? Is the user's assignment to one or more groups still

visible? . . . 102
1.5.6.9 I am having a problem with Confluence LDAP integration . . . 102
1.5.6.10 I enabled LDAP and some users are now returned twice under the user browser . . . 102
1.5.6.11 If a Confluence user had a lowercase username, but an LDAP user has the same username in UPPERCASE, does it matter which one I use? . . . 102
1.5.6.12 If a user already exists in Confluence and an LDAP user with the same username is added, which account's password gets used? . . . 102
1.5.7 Mail Archiving FAQ . . . 102
1.5.7.1 Can Confluence replace my regular mail client? . . . 103
1.5.7.2 How do I get mail into Confluence? . . . 103
1.5.7.3 How do I use the mail archive? . . . 103
1.5.7.4 Okay, I've imported the mail, but where is it? . . . 103
1.5.8 New User FAQ . . . 103
1.5.8.1 Can I use CamelCaseLinks like they do on WardsWiki? . . . 103
1.5.8.2 Can Users Edit Individual Sections Within a Page? . . . 103
1.5.8.3 How does Confluence differ from a wiki? . . . 103
1.5.9 RSS Feeds FAQ . . . 103
1.5.9.1 Create an RSS feed for mail from only specified mail accounts . . . 103
1.5.9.2 How do I fix a "Could not download (Feed URL) - Connection timed out (errno238)" error? . . .
1.5.9.3 How do I fix a "Could not retrieve (Feed URL) - Not Permitted" error? . . .
1.5.9.4 How do I fix an "Error formatting 'macro rss java.lang.NullPointerException'" error? . . .
1.5.9.5 How do I fix an "Unable to retrieve (Feed URL) - Connection refused - connect" error? . . .
1.5.9.6 How do I force authentication for public feeds? . . . 103
1.5.9.7 Is it possible to delete a feed? . . . 103
1.5.9.8 I want to remove RSS Feeds completely . . . 103
1.5.10 Upgrade FAQ . . . 103
1.5.10.1 I cannot find the "Rich Text" editor. Is the editor part of Confluence 1.4.3? . . .
1.5.10.2 Server ID FAQ . . . 103
1.5.10.3 Upgrade My Trial To A Commercial Version . . . 103
1.5.11 Usage FAQ . . . 103
1.5.11.1 Add many files to a page at once . . . 103
1.5.11.2 Create a page by passing parameters to a template . . . 103
1.5.11.3 How do I obtain content that hasn't been modified in a certain period of time . . . 103
1.5.11.4 How to Add a Quick Search for Firefox . . . 103
1.5.11.5 How to disable PDF Export . . . 103
1.5.11.6 How to Find Pages with no Label . . . 103
1.5.11.7 How to Make Confluence Open a New Tab when Clicking on the Attachments Link . . . 103
1.5.11.8 How to Reset a Custom Layout . . . 103
1.5.11.9 Redirect to a specific page (home page) within the site after logging in . . . 103
1.5.11.10 Setup email notifications of page updates . . . 103
1.5.12 Unsupported Platform Information . . . 103
1.5.12.1 Setting up Confluence with IIS . . . 103
1.5.12.1.1 Connecting Confluence with IIS 5.1 or 6 . . . 103
1.5.12.1.2 Connecting Confluence with IIS 7 . . . 104
1.5.12.2 Using the IBM 64bit J9 JDK . . . 104
1.5.13 Support Policies . . . 104
1.5.13.1 Bug Fixing Policy . . . 104
1.5.13.2 How to Report a Security Issue . . . 104
1.5.13.3 New Features Policy . . . 104
1.5.13.4 Patch Policy . . . 104
1.5.13.5 Security Advisory Publishing Policy . . . 104
1.5.13.6 Security Patch Policy . . . 104
1.5.13.7 Severity Levels for Security Issues . . . 104
1.6 Troubleshooting Problems & Requesting Technical Support . . . 104
1.6.1 Content Anonymiser for Data Backups . . . 105
1.6.2 Editing or Deleting a Page That Won't Render . . . 105
1.6.3 Enabling detailed SQL logging . . . 105
1.6.4 General Support Enquiries . . . 105
1.6.5 Generating a Thread Dump . . . 105
1.6.5.1 Generating a Thread Dump Externally . . . 105
1.6.6 Plugin Support Mode . . . 105
1.6.7 Profiling using the YourKit Plugin . . . 105
1.7 Confluence Resources . . . 105
1.7.1 Tips of the Trade . . . 105
2. Information For Evaluators . . . 106
2.1 Confluence Presentation Material . . . 106
3. TreeNavigation . . . 106
4. TreeNavigationVersions . . . 106
5. .bookmarks . . . 106

Confluence Documentation Home

Confluence 3.2.x

User's Guide

The Confluence User's Guide is for project managers, developers, testers – anyone who uses Confluence. New to Confluence? Start by exploring the Confluence dashboard and learning about spaces, pages and blog posts. Try creating a new space, then add a page to that space, add some content and a comment to that page and then export the page to PDF.

Administrator's Guide

The Confluence Administrator's Guide is for people with Confluence administration rights. It will help you set up users and groups, security for users, groups and spaces, and keep track of any changes and updates made within your Confluence site. You may want to customise the look and feel of Confluence, by applying a theme to a space or modifying colour schemes and layouts. Admin tasks such as backup are also covered. You may also find the Knowledge Base, FAQ and the Confluence Announcements and User forums useful.

Installation Guide

The Confluence Installation Guide is for people who are installing Confluence for the first time. Check the requirements and supported platforms, then download and install Confluence. Where to next? The Confluence 101 will help you get started. When setting up Confluence, load the Example Site ('Demonstration Space'), which contains a tutorial and additional content to help you get familiar with using Confluence. If you are using other Atlassian products, take a look at the Integration Guide.

Upgrade Guide

The Confluence Upgrade Guide is for people who are upgrading their instance of Confluence. Start by reading the latest Release Notes, the Upgrade Notes Overview and version-specific Upgrade Notes for the version to which you are upgrading. Then, download Confluence and follow the main Upgrade Guide.

Developer Resources

These resources are for software developers who want to create their own plugins for Confluence. Take a look at the Confluence developer documentation and the API documentation. You may also find the Confluence developers forum useful. (Click here to subscribe.)

Confluence 101

Confluence 101

Thank you for choosing Confluence. To help you get up and running quickly, we've compiled some easy instructions for configuring and using Confluence.

Are you using Confluence Hosted? If so, please visit our Confluence Hosted documentation instead.

Getting Started

1. Installing Confluence

First things first. If you haven't already got Confluence up and running, carry out the following steps:

For Windows: (click to expand)

You may want to watch the video showing how to do this.

1. Go to the Atlassian Download Center.
2. Click 'Show all' and download the Confluence Standalone Windows Installer (.EXE) file.
   If you are upgrading Confluence, please download the ZIP archive instead of the Standalone Installer.
3. Install Java and set JAVA_HOME.
4. Run the Windows Installer .EXE file, choose an installation directory, a home directory, and a port ('8080' will do). We recommend that you choose to 'Run Confluence as a service'.
5. Confluence will start automatically when the Installer finishes, if you selected the option to launch Confluence at the end of the Installer.
6. To access Confluence, go to your web browser and type this address: http://localhost:8080. Windows 'Start' menu shortcuts will also be added which you can also use to start and stop Confluence.
7. Follow the Setup Wizard. This will guide you through the process of setting up your Confluence server and creating an Admin user.

For more help on the technical procedures in this section, see the Confluence Installation Guide.

If you need assistance, please create a support ticket.

Before using Confluence as a production system, you need to switch from the default HSQL database, which is provided for evaluation purposes only. Please see the documentation for details.

For Mac: (click to expand)

You may want to watch the video showing how to do this.

Download the Confluence Standalone for Mac (TAR.GZ) file from the Atlassian Download Center, and follow the installation instructions. The Setup Wizard will guide you through the process of setting up your Confluence server and creating an Admin user.

If you need assistance, please create a support ticket.

Before using Confluence as a production system, you need to switch from the default HSQL database, which is provided for evaluation purposes only. Please see the documentation for details.

For Unix or Linux: (click to expand)

Download the Confluence Standalone for Linux (TAR.GZ) file from the Atlassian Download Center, and follow the installation instructions. The Setup Wizard will guide you through the process of setting up your Confluence server and creating an Admin user.

If you need assistance, please create a support ticket.

Before using Confluence as a production system, you need to switch from the default HSQL database, which is provided for evaluation purposes only. Please see the documentation for details.

2. Adding Users

Confluence was designed to improve team communication, so you will want to configure it such that your colleagues can participate on it:

For each of your users, you will need to do the following: (click to expand)

1. Create users — add user accounts for the people who will create most of the pages and blog posts in your Confluence site.
   - Your users will be automatically added to the 'confluence-users' group. For more about groups, please see the documentation.
2. (Optional) Enable 'anonymous' mode — allow non-logged-in users to access your Confluence site. Note that they will only be able to view content where the space permissions allow 'Anonymous' access.

3. Configuring Mail

Confluence can send notifications based on user preferences, such as watching the contributions to a particular space or page. To take advantage of this feature, you will first need to tell Confluence about your SMTP server.

Mastering the Basics

4. Creating Spaces, Pages and Blog Posts

Your new Confluence site contains a 'Demonstration' space, which contains a Tutorial. This will guide you through the basics of using the Dashboard, creating a space, adding pages and publishing blog posts ('news items').

Your users will need to access Confluence at **`http://<Confluence machine name>:8080`** (not `http://localhost:8080`).

5. Searching Confluence

- Search Confluence
- Advanced search
- Search the content of attachments

6. Subscribing to Updates

- Create an RSS feed
- 'Watch' a page (or space) to be emailed whenever it's edited
- Mark your 'favourite' places so you can quickly navigate to them later

Customising Confluence

(Note that you need to be an Administrator to do the tasks in this section.)

7. Adding Your Own Logo

Why would I do this? (click to expand)

By replacing the default Confluence logo with your organisation's logo, your Confluence site will immediately appear more familiar to visitors.

How do I do this? (click to expand)

Please see the following documentation:

- Changing the Confluence-wide 'global' logo
- Changing a space's logo

8. Changing Confluence's Colours

Why would I do this? (click to expand)

By replacing the default Confluence colours with your organisation's standard colours, your Confluence site will immediately appear more familiar to visitors.

How do I do this? (click to expand)

Please see the following documentation:

- Customising Confluence's colour scheme
- Styling Confluence via CSS

9. Choosing a Different 'Theme'

Why would I do this? (click to expand)

Depending on what you are using Confluence for, you may want to alter the 'theme' (i.e. screen layout) for an individual space or for the entire site. For example, for a space that contains documentation, you may prefer a theme that displays a table of contents on the left and allows you to configure a page header and footer.

How do I do this? (click to expand)

Please see the following documentation:

- Applying a theme to a site
- Applying a theme to a space

10. Using a Page Template

Why would I do this? (click to expand)

A page template is a form that is displayed when a user creates a page, so that the user can enter content in particular fields. A page template can be made available to your entire Confluence site or just to a particular space, e.g. a Human Resources space might have a template called 'Leave Request', or an IT space might have a template called 'Asset Configuration'.

How do I do this? (click to expand)

Please see the following documentation:

- Adding a template to a site
- Adding a template to a space

11. Installing Plugins

Why would I do this? (click to expand)

There are hundreds of new features and customisations available for Confluence as plugins. For example, many additions to Confluence's macros are available as plugins. Or, if you want to populate Confluence with content created elsewhere, take a look at the WebDAV plugin and the Universal Wiki Converter.

How do I do this? (click to expand)

Confluence comes with a bunch of bundled plugins and others can be installed in one click via our plugin repository, so why not why not browse our most popular plugins.

12. Import Existing Content

Why would I do this? (click to expand)

The best way to spur adoption is to get some content in Confluence. It is likely that you have existing documentation stored in another wiki or in static documents. Confluence makes it easy to import your existing content so that you do not have to re-invent the wheel or start from scratch.

How do I do this? (click to expand)

Refer to our documentation - Importing Content Into Confluence - for details on:

- Importing Content from Other Confluence Sites
- Importing Content from Other Wikis - e.g. Mediawiki, SharePoint, Socialtext, MindTouch etc.
- Importing Content from an Office Document
- Importing Web Content
- Importing Other Non-Wiki Content

Important Next Steps

(Note that you need to be an Administrator to do the tasks in this section.)

13. Connecting to an External Database

Before using Confluence as a production system, you need to switch from the default HSQL database, which is provided for evaluation purposes only. Please see the documentation for details.

14. Backing up Data

To back up your Confluence data, and establish processes for regular backups, please see the documentation.

Thank you for using Confluence.

We're always happy to help. Feel free to contact us with any questions you have.

Confluence Administrator's Guide

Confluence Administrator's Guide
Configuring Confluence Data and Backups System Administration Importing Data Mail Configuration Security User Management Design and Layout Integrating Confluence and JIRA Plugins and Macros Performance Tuning Character Encoding Support

Additional Resources

Visit the Configuration Guide for documentation on configuring databases and application servers.
The Confluence User's Guide has information on how to use Confluence as a collaborative tool.
Go to Documentation Home for links to more resources.

Download



You can download the Confluence Admin Guide in PDF, HTML or XML formats.

Site Administrator?

The **Confluence Administrators Guide** provides information to site administrators on how to manage their Confluence instances.

If you still have a question that hasn't been answered, write and tell us about it.

Configuring Confluence

Site Configuration
Configuring the Site Home Page
Editing the Site Title
Editing the Site Welcome Message
Configuring the Destination of View Space Links
Editing the Global Logo
Configuring the Server Base URL
Configuring the Site Support Address
Configuring HTTP Timeout Settings
Configuring System Properties
Customising Default Space Content

Optional Settings
Enabling Remote APIs
Enabling Rich Text Editing Option
Making Rich Text Editing default
Enabling CamelCase Linking
Enabling Trackback
Enabling Threaded Comments
Attachment Storage Configuration
Configuring Attachment Size

Display Settings
Configuring Indexing Language
Number of Ancestors to Show in Breadcrumbs
Configuring Character Encoding
Configuring Time and Date Formats
Configuring Number Formats
Thumbnail Settings
Configuring Shortcut Links
Showing Link Icons

Data and Backups

Backups Configuration
Configuring Daily Backups
Changing time of Daily Backup
Manually Backing Up The Site

Restoring Data
Restoring a Site
Restoring a Space
Restoring Data During Setup
Restoring Data from the Administration Console
Retrieve file attachments from a backup
Restoring a Test Instance from Production

Scheduled Jobs
Scheduled Jobs

Data Model
The Confluence data model

System Administration

System Configuration
Viewing System Information
Viewing and Editing License Details
Cache Statistics
Site Statistics

System Administration
Content Index Administration
Upgrading Confluence
Migrating Confluence Between Servers
Migrate to Another Database
Important Directories and Files
Rebuilding the Ancestor Table
Finding Unused Spaces

Large Confluence Installations
Operating Large or Mission-Critical Confluence Installations
Configuring a Large Confluence Installation
Overview of Confluence Clusters
Cluster Administration page

Importing Data

Universal Wiki Converter
Snip Snap Import
Importing Pages from Disk

Mail Configuration

Configuring a Server for Outgoing Mail
The Mail Queue
Enabling the 'Mail Page' plugin

Security

Overview and Advisories
Security Overview and Advisories

Security Options

Spam Prevention via Captcha
Managing External Referrers
Hiding external referrers
Hiding External Links From Search Engines
Excluding external referrers
User Email Visibility
Anonymous Access to Remote API
Adding SSL for Secure Logins and Page Security

User Management

Confluence User Management
Global Groups Overview
Global Permissions Overview
Setting up Anonymous Access
Adding a New User
Editing User Details
Removing a User
Enabling or Disabling Public Signup
Adding or Removing Users in Groups
Adding a Group
Removing a Group
Viewing members of a group
Restoring Passwords To Recover Admin User Rights

External User Management
Understanding User Management in Confluence
Disabling the Built-In User Management
LDAP Integration Overview

Troubleshooting LDAP User Management
Requesting External User Management Support
Troubleshooting the "Not Permitted" Screen under LDAP Integration

Crowd User Management
Integrating Confluence with Crowd

JIRA User Management
Delegating User Management in Confluence to JIRA
Override properties in JIRA to Confluence Bridge

Design and Layout

Configuring Layouts
Customising Look and Feel Overview
Customising Colour Schemes
Customising Layouts
Working With Decorator Macros
Global Templates
Customising the eMail Templates

Configuring Themes
Themes Overview
Creating a Theme
Applying a Theme to a Site
Including Cascading Stylesheets in Themes

Integrating Confluence and JIRA

Integrating JIRA and Confluence
Configuring JIRA Issues Icon mappings
Adding Confluence EAR-WAR to JIRA Standalone
Setting Up Trusted Communication between JIRA and Confluence

Plugin Management

Confluence Plugin Guide
Installing and Configuring Plugins Manually
Installing and Configuring Plugins using the Plugin Repository Client
Configuring the Office Connector
Setting Up Trusted Communication between JIRA and Confluence
Configuring a URL Whitelist
Enabling the 'Mail Page' plugin

Macros
Adding, Editing and Removing User Macros
Enabling HTML macros
Enabling the html-include Macro
Flowchart Macro
Troubleshooting the Gallery Macro
Configuring the userlister Macro
Working With Decorator Macros
Writing Macros

Performance Tuning

Memory usage and requirements
Configuring a Large Confluence Installation
Performance tuning
Enabling HTTP Compression
Working with Confluence Logs
Page Request Profiling
Profiling using the YourKit Plugin

Character Encoding

Configuring Encoding
Troubleshooting Character Encodings

Support

How to Get Support

Administration

- Cache Statistics
- Changing time of Daily Backup
- Confluence Data Directory Configuration
- Confluence home directory contents
- Content Index Administration
- Finding Unused Spaces
- Important Directories and Files
  - Confluence Home Directory
  - Confluence Installation Directory
- Installing a Language Pack
- Manually Backing Up The Site
  - Configuring Daily Backups
    - User Submitted Backup & Restore Scripts
- Migrating Confluence Between Servers
- Rebuilding the Ancestor Table
- Restoring a Site
- Restoring a Space
  - Changing the version of a space backup
- Restoring a Test Instance from Production
- Restoring Data from other Backups
- Restoring Data from the Administration Console
- Retrieve file attachments from a backup
- Troubleshooting failed XML site backups
- Troubleshooting XML backups that fail on restore
  - Migrating from HSQLDB to MySQL
- Viewing and Editing License Details
- Viewing System Information
  - Live Monitoring Using the JMX Interface
  - Site Statistics
- Installing Patched Class Files

Administrators Guide Home Confluence Documentation Home

Cache Statistics

Confluence provides statistics about its internal caches that allow you to track the size and hit ratio of each cache and tune it for better performance (if necessary). See Performance Tuning for more information.

Configurable Caches

System administrators can change the sizes of Confluence's internal caches through the Administration Console and these changes will take effect without the need to first shut down and then restart Confluence. The maximum number of units for any of the defined cache regions can be adjusted individually.

Note that larger cache sizes will require more memory at runtime, so you should review the memory allocation of the Confluence Java process and the physical memory available on your server.

Viewing Cache Statistics and Modifying Cache Sizes

To view the cache statistics:

1. Go to the Confluence '**Administration Console**'. To do this:
   - Open the '**Browse**' menu and select '**Confluence Admin**'. The 'Administration Console' view will open.
2. Click '**Cache Statistics**' in the left-hand panel. There you will find a list of all objects cached within Confluence.
3. Click the '**Advanced**' tab for more detail. Below is an example for one of the most frequently used caches, the 'Content Object' cache.

Name	Percent Used	Effectiveness	Objects / Size	Hit / Miss / Expiry	Adjust Size	Flush
Content Object	80%	73%	4023 / 5000	374550 / 140460 / 55044	*Adjust Size*	*Flush*

About the generated numbers:

Percent Used :	=(Objects)/(Size)
Effectiveness :	=(Hits)/(Hits + Misses)

Objects / Size:	The number of entries in the cache / the number of total possible entries allowed (configurable).
Hit / Miss / Expiry:	The number of reads accessing cache where required content was found / the number of reads accessing cache where required content was not found / the number of objects evicted from the cache.
Adjust Size	Use this option to specify a different maximum cache size. Enter a new cache size and click the '**Adjust Size**' button to set it.
Flush:	Flushes the cache.

For instance, to calculate **Percent Used**:

```
Percent Used = Objects / Size

Percent Used = 4023/5000 = 80%
```

To calculate **Effectiveness**:

```
Effectiveness = (Hits)/(Hits + Misses)

Effectiveness = 374550 / (374550 + 140460) = 73%
```



The clustered versions of Confluence use distributed cache called Tangosol Coherence.

Watching the Cache Contents

To see the specific items in the caches, view the cache statistics at `<baseUrl>/admin/cachecontents.jsp`.

Additional Notes about Configurable Caches

Changes to cache size configurations persist across confluence restarts as they are saved in the `<confluence-home>/config/confluence-coherence-cache-config.xml` file (or `<confluence-home>/config/confluence-coherence-cache-config-clustered.xml` for a clustered instance). In most cases, a Confluence administrator will never need to know about these files. However, if it is necessary to tune cache options other than the maximum cache size, this can be done by manually editing these files. See Cache Performance Tuning for details.



Important note about clustered Confluence installations
The cache configuration file is stored in a home directory of each cluster node. When a Confluence administrator changes a cache size, all running cluster nodes will automatically update their own configuration files in their respective home directories. However, if a cluster node is not running when an administrator adjusts a cache size, the `/config/confluence-coherence-cache-config-clustered.xml` file in its home directory will not be updated. Since cluster caches are configured by the first node to start, if a node with an outdated cache configuration is the first to start up, the whole cluster would end up using the configuration of that node. However, copying this file from one node to another would resolve this issue.

Performance Tuning

If you need to tune your application when under high usage, you may like to review this document for suggestions.

RELATED TOPICS


Cache Performance Tuning


Cache Performance Tuning for Specific Problems


Confluence Cache Schemes


Cache Statistics


Viewing and Editing License Details


Viewing System Information


Administrators Guide Home 
Confluence Documentation Home

Changing time of Daily Backup



Atlassian recommends disabling the XML backup both for performance and reliability. 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. Upgrading is better done without the XML backup. This page can also help with troubleshooting XML Space versions.

By default, Confluence runs its daily backup at 2.00 AM. You can configure Confluence to perform the backup at a time that is best suited to you or your organisational needs.



Time is derived from the Confluence server
The time zone is taken from the server on which Confluence is running. To check the time according to the server, do the following:

1. Go to the Confluence '**Administration Console**'. To do this:
    - Open the '**Browse**' menu and select '**Confluence Admin**'. The 'Administration Console' view will open.
2. Click '**System Information**' in the left-hand panel and look at the '**System Time**'.

Confluence uses Quartz for scheduling periodic jobs. To change the time of your daily backup, you will need to edit the Quartz configuration.

To change the time of your daily backup

1. Open the Quartz configuration file `schedulingSubsystemContext.xml` located under `confluence/WEB-INF/lib/confluence-x.x.x.jar`. Where x.x.x is your Confluence version number.
2. Find the following section of the file:

```
1.<bean id="backupTrigger" class=
"org.springframework.scheduling.quartz.CronTriggerBean">
2.      <property name="jobDetail">
3.           <ref bean="backupJob"/>
4.      </property>
5.      <property name="cronExpression">
6.           <value>0 0 2 * * ?</value>
7.      </property>
8.</bean>
```

3. The string `'0 0 2 * * ?'` sets up a Cron Trigger for the job to run at the zeroth second of the zeroth minute of the 2nd hour, every day of every month, every day of the week.
4. Re-jar the file, either with a zip utility (change the title of .zip back to .jar) or a java command.
5. You can set a new time by editing this string. Note that the date and time format in this configuration file is in this order:

   Second minute hour day

6. Restart Confluence.



For example, to set the new time to twenty past ten PM, change the string to `'0 20 22 * * ?'`.

If you wanted to back up only once a week, for example, at midnight on Sundays, you would change the string to `'0 0 0 ? * SUN'`.

For complete details on the formatting of the cron string, please see http://www.opensymphony.com/quartz/api/org/quartz/CronTrigger.html.

RELATED TOPICS

 Site Backup and Restore

 Configuring Daily Backups

 User Submitted Backup & Restore Scripts

 Backup FAQ

 Changing time of Daily Backup

Manually Backing Up The Site

Production Backup Strategy

Administrators Guide Home Confluence Documentation Home

Confluence Data Directory Configuration

Here is a link listing important Confluence files.

The home directory defines the location of the directory where Confluence will store it's data, including attachments, indexes and backups. Administrators can set this location by defining a value for the file `<MY-INSTALL>/confluence/WEB-INF/classes/confluence-init.properties`. To find what your home directory is currently set to, open this file and check the `confluence.home` property. It is unset on new installations.

Windows Configuration

On Windows, this path:

```
C:\confluence\data
```

will be written like so:

```
confluence.home=C:/confluence/data
```

Note that all backslashes (\) are written as forward slashes (/).

UNIX/Linux/Mac Configuration

On any UNIX-based system, the property is defined using the normal directory syntax:

```
confluence.home=/var/confluence/
```

Symbolic links

If your `confluence.home` directory contains a symbolic link, you must define the absolute path.

Please note that there can be no symbolic links within the `confluence.home` directory. If disk space is an issue, place the entire `confluence.home` directory on a disk partition where there is enough space.

The absolute path of generated files (such as exports) is compared with the absolute path of the `confluence.home` directory when constructing URLs. When a sub-directory has a different path, the URL will be incorrect, and you may receive "Page not found" errors. These measures are in place to prevent "directory traversal" attacks.

Fixing the Confluence Configuration

The Confluence configuration file: `confluence-cfg.xml` inside the home directory may contain references to the original location of your Confluence home. You will need to edit this file to update these references to also point to the new location. The two properties in this file that need to change are:

- `daily.backup.dir` if you have not configured your backups to be placed elsewhere already
- `hibernate.connection.url` if you are using the embedded HSQL database.

Confluence home directory contents

The Confluence home directory contains data that work in concert with the Confluence database to provide the wiki experience. This document outlines the purpose of the various files and directories in the Confluence home directory.

Tip: Another term for 'Home directory' would be 'data directory'.

Files and directories

`confluence.cfg.xml`

This file is the most critical file in the Confluence home directory. It contains all of the information necessary for Confluence to start up such

as:

- Product license
- Context path
- Database details such as location and connection pool settings
- Paths to important directories

attachments

This directory contains every version of each attachment stored in Confluence. This directory is not used when Confluence is configured to store attachments in the database. Attachments are always stored in the database in clustered instances of Confluence.

Paths within this directory have the following structure:

```
/attachments/PAGE_ID/ATTACHMENT_ID/VERSION
```

An alternative directory may be specified for attachment storage by setting the `attachments.dir` property in `confluence.cfg.xml`

backups

Confluence will place its daily backup archives in this directory, as well as any manually generated backups. Backup files in this directory take the following form:

```
daily-backup-YYYY_MM_DD.zip
```

An alternative directory may be specified for backups by setting the `daily.backup.dir` property in `confluence.cfg.xml`.

bundled-plugins



This directory exists for Confluence 2.3 and above

Recent versions of Confluence ship with a set of *bundled* plugins. These are plugins written by the Atlassian and the Confluence community that we think provide useful and broadly applicable functionality in Confluence. The {{bundled-plugins)) directory is where Confluence will unpack its bundled plugins when it starts up. This directory is refreshed on every restart, so removing a plugin from this directory will not uninstall the plugin; it will simply be replaced the next time Confluence starts up.

database

This is where Confluence stores its database when configured to run with the HSQL embedded database and as such contains all Confluence runtime data. Instances configured to run using an external database such as MySQL will not use this directory.

index

This is where Confluence stores its indexes for rapid retrieval of often used data. The Confluence index is used heavily by the application for content searching and recently updated lists and as such is critical for a running Confluence instance. It is important to note however that should the data in this directory be lost or corrupted, it can be restored by running a full reindex from within Confluence. This can take a long time depending on how much data is stored Confluence's database.

An alternative directory may be specified for the index by setting the `lucene.index.dir` property in `confluence.cfg.xml`. As this is the most heavily accessed directory in the Confluence home directory you might want to consider hosting it on the fastest disk available. It would also be useful if the disk holding the Confluence index was not heavily used by any other application to reduce access contention.

plugin-cache



This directory exists for Confluence 2.3 and above

As of Confluence 2.3, all Confluence plugins are now stored in the Confluence database. To allow for quicker access to classes contained within the plugin JARs, Confluence will cache these plugins in the `plugin-cache` directory. This directory is updated as plugins are installed and uninstalled from the system and is completely repopulated from the database every time Confluence is restarted. As such, removing plugins from this directory does not uninstall them.

resources

The `resources` directory stores any space logos used in your Confluence instance. For each space with a space logo, there is a directory within `resources` named after the space's key. That directory contains the space's logo.

temp

The `temp` directory is used for various runtime functions such as exporting, importing, file upload and indexing. As the name suggests, and file in this directory is of temporary importance and is only used during runtime. This directory can be safely emptied when Confluence is offline.

An alternative directory may be specified for temporary data by setting the `webwork.multipart.saveDir` property in `confluence.cfg.xml`.

thumbnails

When Confluence generates a thumbnail of an image (for example when the `gallery` macro is used), the resulting thumbnail is stored in this directory for quicker retrieval on subsequent accesses. This directory is essentially a thumbnail cache, and deleting files from this directory simply means the thumbnail will have to be regenerated on the next access.

Confluence 2.2 and older

The following files and directories were used by versions of Confluence older than 2.3

config

The `config` directory is used to store data used by Confluence's bandana data persistence framework. This system is used by Confluence to store the global instance settings and is used by various plugins for their own configuration and data persistence needs. Confluence versions 2.3 and later store these data in the Confluence database and do not use this directory.

The most important file in this directory is the `confluence-global.bandana.xml` file. This file is used to store all of the settings from the `Administration` console in Confluence.

plugins

The `plugins` directory is where Confluence stores all installed plugin JARs. It is possible to install and remove plugins by placing and deleting plugin JARs from this directory.

default-formatting.properties

This properties file contains various formatting information such as the formats for decimal numbers and dates used in the Confluence user interface. These configuration data where relocated to the Confluence database from Confluence 2.3 and onwards.

Content Index Administration

The Content Indexes power Confluence's search functionality and they are also used for a number of related functions such as building email threads in the mail archive, the Space Activity feature and lists of recently-updated content. The Gliffy Plugin also uses them for some of its functionality.

For reasons of efficiency, content is not immediately added to the index. New and modified Confluence content is first placed in a queue and the queue is processed once every minute (by default).

On this page:

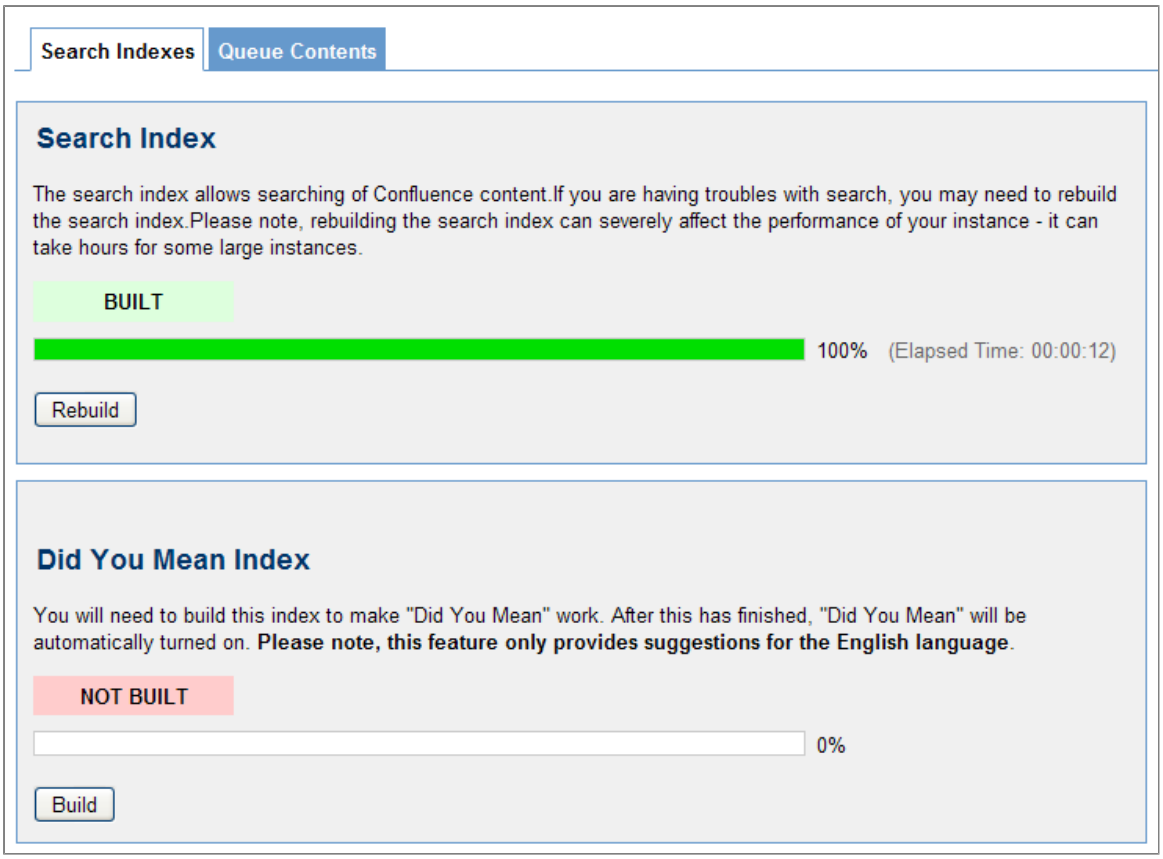
- Viewing the Content Index Summary
- Rebuilding the Content Indexes
- Slow Reindexing
- Viewing the Index Browser
- More Hints and Tips

Viewing the Content Index Summary

To see information about your Confluence instance's content indexing,

1. Go to the Confluence '**Administration Console**'. To do this:
    - Open the '**Browse**' menu and select '**Confluence Admin**'. The 'Administration Console' view will open.
2. Click '**Content Indexing**' under the heading '**Administration**' in the left-hand panel.

Screenshot: Index summary



Rebuilding the Content Indexes

The content indexes are maintained automatically, but you may need to rebuild one or both of them manually under circumstances such as these:

- Your searching and mail threading are malfunctioning. (Rebuild the Search Index.)
- Your 'Did You Mean' feature is malfunctioning. (Rebuild the Did You Mean Index.)
- After an upgrade. If a content re-index is required after an upgrade, it will be noted in an upgrade subsection of the relevant release notes.

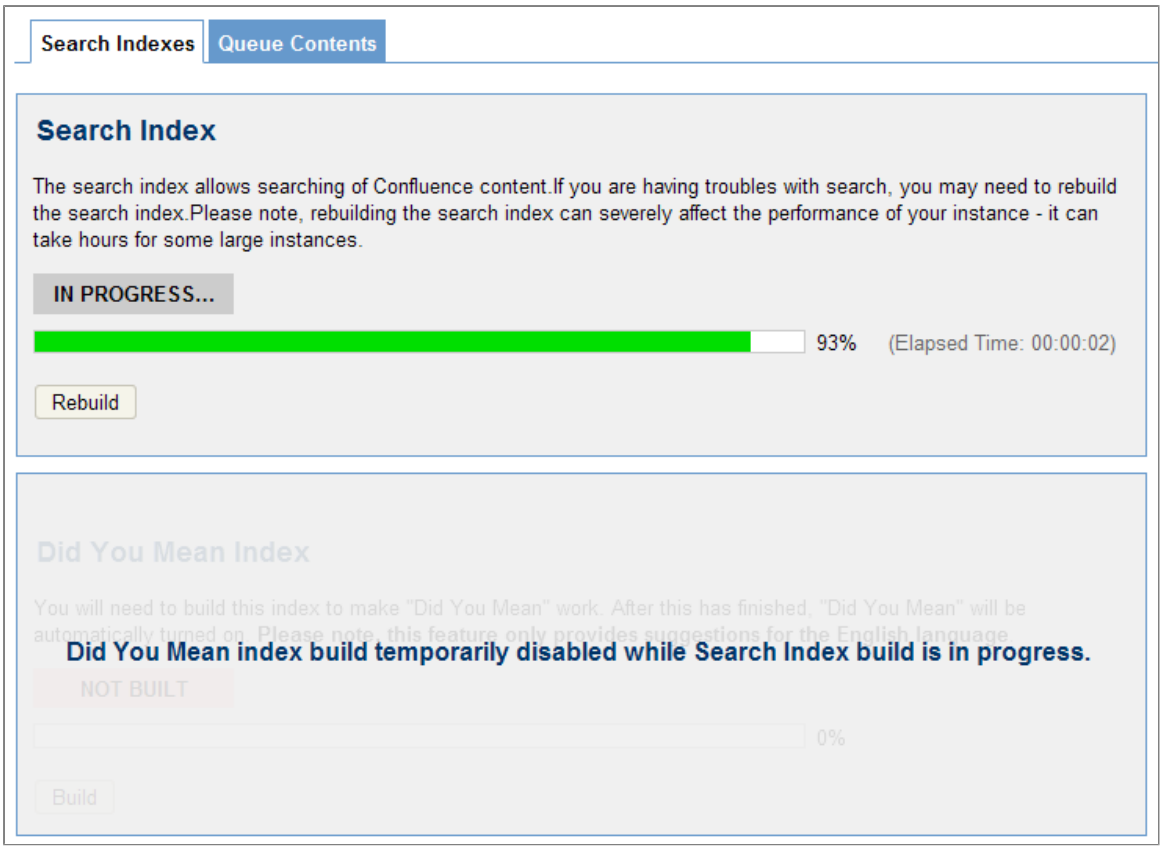
In new Confluence installations, the 'Did You Mean' feature is not initially activated. To activate it, you first need to build its index by clicking its '**Build**' button on this page.

To rebuild either of the content indexes,

1. Go to the Confluence '**Administration Console**'. To do this:
    - Open the '**Browse**' menu and select '**Confluence Admin**'. The 'Administration Console' view will open.
2. Click '**Content Indexing**' under the heading '**Administration**' in the left-hand panel.
3. Click the '**Rebuild**' button in either the 'Search Index' or 'Did You Mean Index' sections on this page, depending on the particular index you want to rebuild.

- If one of these indexes has not yet been built, its button will indicate 'Build' (instead of 'Rebuild').
- As shown in the image below, only one index can be (re)built at a time.

Screenshot: Content Indexing



Slow Reindexing

Does the reindexing take a long time to complete? The length of time depends on the following factors:

- Number of pages in your Confluence instance.
- Number, type and size of attachments.
- Amount of memory allocated to Confluence.

It may help to increase the heap memory allocation of Confluence by following the instructions in the JIRA documentation.

If you are running an older version of Confluence and find that the index rebuild is not progressing, you may need to shut down Confluence, and restart it with the following *Java system property* set: `bucket.indexing.threads.fixed=1`. This will cause the re-indexing to happen in a single thread and be much more stable (but slower).

Viewing the Index Browser

You can access the Lucene index browser used with Confluence, to see if the required documents have been added to the index.

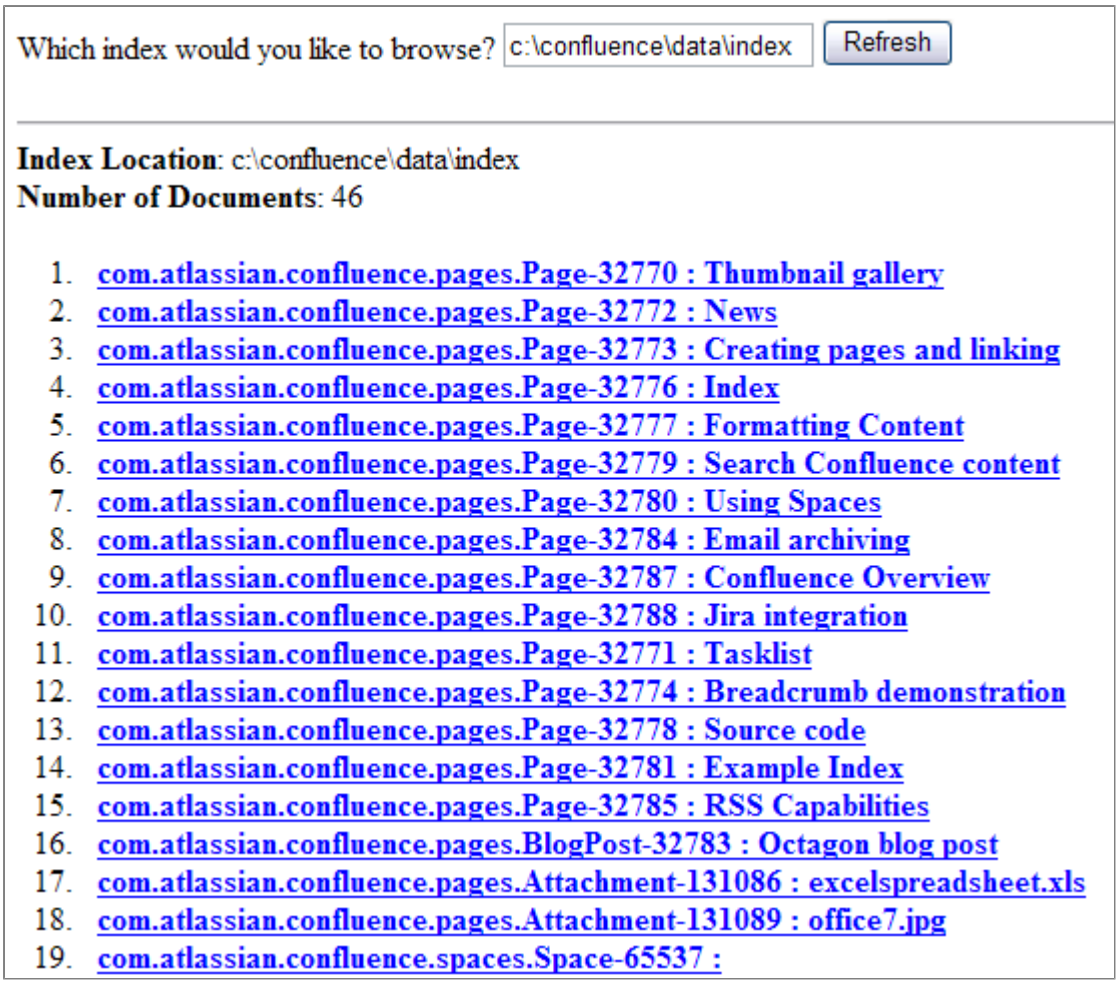
To view the index browser,

1. Open a browser window.
2. Log in to Confluence.
3. Type the following into the browser address bar: `<Confluence location>/admin/indexbrowser.jsp`.
   For example:

   ```
   http://localhost:8080/admin/indexbrowser.jsp
   ```

4. You will be asked "Which index would you like to browse?"
5. Type in the path of your index directory. By default, this is located at `<Confluence Home directory/index>`.
   For example, `c:\confluence\data\index`.
6. You should see a list of indexed file names.

Screenshot: Index Browser



More Hints and Tips

- If you are still experiencing problems after performing the above rebuild, the next step might be to remove the index and rebuild it from scratch.
  The Space Activity feature uses the index to store data. If you remove the index file, the existing activity data will disappear.
- A tip for the development community: If you have the Confluence source, you can look for references to the SmartListManager to find the screens and lists that rely on the content index.

RELATED TOPICS

Working with Macros

Creating a Lowercase Page Title Index

Configuring Indexing Language

Rebuild the Content Indices from scratch

Content Index Administration

[!Administration Guide Attachments directory^adminhome.gif!] Confluence Documentation Home

Finding Unused Spaces

Sometimes, you want to know what is *not* being used. It's great to know what's getting most attention, but what about stagnant pages, or even entire spaces that are no longer active?

While viewing space activity and the Global Statistics plugin can provide hints, they still don't always provide enough detail. The simple way is to go directly to the database. We recommend DbVisualizer, and have basic instructions for connecting it to HSQLDB.

The following query identifies the last date on which content was modified in each space within a single Confluence instance:

```
1.SELECT spaces.spacename, MAX(content.lastmoddate)
2.FROM content, spaces
3.WHERE content.spaceid = spaces.spaceid
4.GROUP BY spaces.spacename;
```

It returns a list of spacenames, and the last date and time at which any content was added or changed.

Alternatively, this one simply identifies spaces whose content hasn't changed since a specified date:

```
1.SELECT spaces.spacename
2.FROM content, spaces
3.WHERE content.spaceid = spaces.spaceid
4.GROUP BY spaces.spacename
5.HAVING MAX(content.lastmoddate) < '2006-10-10';
```

The result is a simple list of space names.

It's also possible to present the information in a wiki page, using the SQL plugin, which can be installed via the Plugin Repository. You'll also need to define a database resource in `conf/server.xml` and `confluence/WEB-INF/web.xml`, as described here. Having done so, you can use wiki markup code like the following, replacing `confluenceDS` with the name of your own local datasource:

```
1.h3. Space activity
2.{sql:dataSource=confluenceDS|output=wiki}
3.SELECT spaces.spacename AS Space, MAX(content.lastmoddate) AS LastModified
4.FROM content, spaces
5.WHERE content.spaceid = spaces.spaceid
6.GROUP BY Space;
7.{sql}
```

The result will be something like this:

Space activity:

space	lastmodified
Private Space	2007-10-11 11:34:04.914
Another space	2007-10-11 11:39:39.716
More space	2007-10-11 11:40:11.688

You can try the Chart plugin in combination with the SQL plugin to give more visually attractive results.

Important Directories and Files

The Installation Directory

The 'Confluence Installation directory' is the directory into which the Confluence application files and libraries have been unpacked (unzipped) when Confluence was installed. Confluence does not modify or store any data in this directory. This directory is also sometimes called the 'Confluence Install directory'.

Important Files and Directories

- `confluence/WEB-INF/classes/confluence-init.properties` : This file tells Confluence where to find the Confluence Home Directory. This file is modified by the administrator when installing Confluence.
- `confluence/WEB-INF/classes/osuser.xml` : This file is modified when connecting Confluence to an external user management system such as an LDAP server or JIRA instance in Confluence 2.0 and earlier. For more information, refer to Understanding User Management in Confluence.
- `confluence/WEB-INF/classes/atlassian-user.xml` : This file is modified when connecting Confluence to an external user management system such as an LDAP server or Crowd. For more information, refer to Understanding User Management in Confluence.
- `confluence/WEB-INF/lib/` : This directory is used when deploying plugins, especially those plugins that cannot automatically be loaded through the Administration Console.
- `confluence/WEB-INF/classes/log4j.properties` : Confluence's logging configuration file. See Working with Confluence Logs.
- `confluence/WEB-INF/classes/ehcache.xml` : This is where you can configure the size of Confluence's internal caches
- `confluence/WEB-INF/classes/styles/site-css.vm` : Confluence's main stylesheet, modify at your own risk
- `conf/server.xml` : SSL configuration.

Memory Settings

The file used to edit JAVA_OPTS memory settings will depend on the method used to install Confluence, as well as the operating system used for your installation.

- Windows Users
  - **Confluence Standalone** — `bin/setenv.bat`
  - **Confluence Installer** — `wrapperwin32.conf`
- Mac/Linux Users
  - **Confluence Standalone** — `bin/setenv.sh`
  - **Confluence Installer** — `wrapperosx.conf`

The Temp Directory

The temp directory is configured in the Java runtime and some Confluence components write temporary files or lockfiles into this directory.

Typically, this directory is `/tmp` on Unix systems, or `C:\Temp` on Windows.

To change the location of this directory, you should start the Java Virtual Machine in which confluence is running with the argument:

`-Djava.io.tmpdir=/path/to/your/own/temp/directory`.

The Confluence Home Directory

The Confluence Home directory is the folder where Confluence stores its configuration information, search indexes and page attachments. If you're using the embedded HSQLDB database supplied for evaluation purposes, the database files are also stored in this directory.

 Tip: Another term for 'Home directory' would be 'data directory'.
Administrators can expect the Confluence Home Directory to grow quite large in a busy site.

The location of this directory is configured by the system administrator during installation (see `confluence-init.properties` above).

Important Files and Directories

- `confluence.cfg.xml` : Confluence's core configuration file; includes the configuration for connecting to its database.
- `default-formatting.properties` : Some auxiliary configuration data concerning default number and date formats.
- `attachments/` : All file attachments in the Confluence site are stored under this directory. This is the only place Confluence keeps attachment files.
- `backups/` : If Confluence is configured to produce daily backups, these are kept in this directory. Administrators should occasionally delete old or unwanted backups from this directory to prevent it from growing too large.
- `config/` : Miscellaneous global and per-space configuration files are kept in this directory.
- `database/` : If Confluence is being run from the embedded HSQL database, the database files will be kept in this directory.
- `index/` : The full-text search index is kept in this directory. Removing or modifying files in this directory may cause search to no longer function. Rebuilding the search index from Confluence's global administration screen will completely regenerate the contents of this directory.
- `plugins/` : Dynamically uploaded plugins are stored in this directory. Administrators can install new plugins by copying them into this directory and triggering a scan from the plugin management page.
- `temp/` : Confluence stores temporary files in this directory, especially during backups and exports. A daily job within Confluence deletes files that are no longer needed.
- `thumbnails/` : Stores temporary files for image thumbnails. The contents of this directory can be safely deleted, as Confluence will regenerate thumbnails as required.
- `velocity/` : Storage for customised page layouts, globally and per-space.

Database

All other data — page contents, links, archived mail and so on — is kept in the database. If you have configured Confluence to use the embedded HSQL database, the database will store its files under `database/` in the Confluence Home Directory. Otherwise, the database management system you are connecting to is responsible for where and how your remaining data is stored.



Tip
All of Confluence's persistent data is stored either in the Confluence Home Directory, or the database. If you have backup copies of both of these, taken at the same time, you will be able to restore Confluence from them (see Restoring Data from other Backups).

RELATED TOPICS

Confluence Home Directory
Confluence Installation Directory
The Embedded HSQLDB Database
Database Configuration

Administrators Guide Home [!Administration Guide Attachments directory^dochome.gif!]

Confluence Home Directory

Often in the documentation, you'll see a reference to the 'Confluence Home directory'.

What is the Confluence Home Directory?

The Confluence Home directory is the folder where Confluence stores its configuration information, search indexes and page attachments. If you're using the embedded HSQLDB database supplied for evaluation purposes, the database files are also stored in this directory.

Tip: Another term for 'Home directory' would be 'data directory'.

You can also read about the contents of the Home directory.

Finding the Confluence Home Directory

The location of the Confluence Home directory is defined when you install Confluence. This location is stored in a configuration file called `confluence-init.properties`, which is located inside the `confluence/WEB-INF/classes` directory in your Confluence Installation directory.

When Confluence first starts up, it reads the `confluence-init.properties` file to determine where to look for the Home directory.

RELATED TOPICS

Confluence Installation Directory
Important Directories and Files
The Embedded HSQLDB Database

Confluence Installation Directory

The 'Confluence Installation directory' is the directory into which the Confluence application files and libraries have been unpacked (unzipped) when Confluence was installed. Confluence does not modify or store any data in this directory. This directory is also sometimes called the 'Confluence Install directory'.

RELATED TOPICS

Confluence Home Directory
Important Directories and Files

Installing a Language Pack

Confluence ships with a number of bundled language packs. These languages appear as options on the 'Language Configuration' screen in the Administration Console when choosing a default language and as 'Language' options for users in their user settings. You can make additional languages available for selection by installing language packs. Please note, you must be a Confluence administrator to install a language pack.

Language packs are essentially plugins. Hence, the process of installing a language pack is the same as installing a new plugin:

- Install a language pack using the Plugin Repository Client
- Install a language pack manually

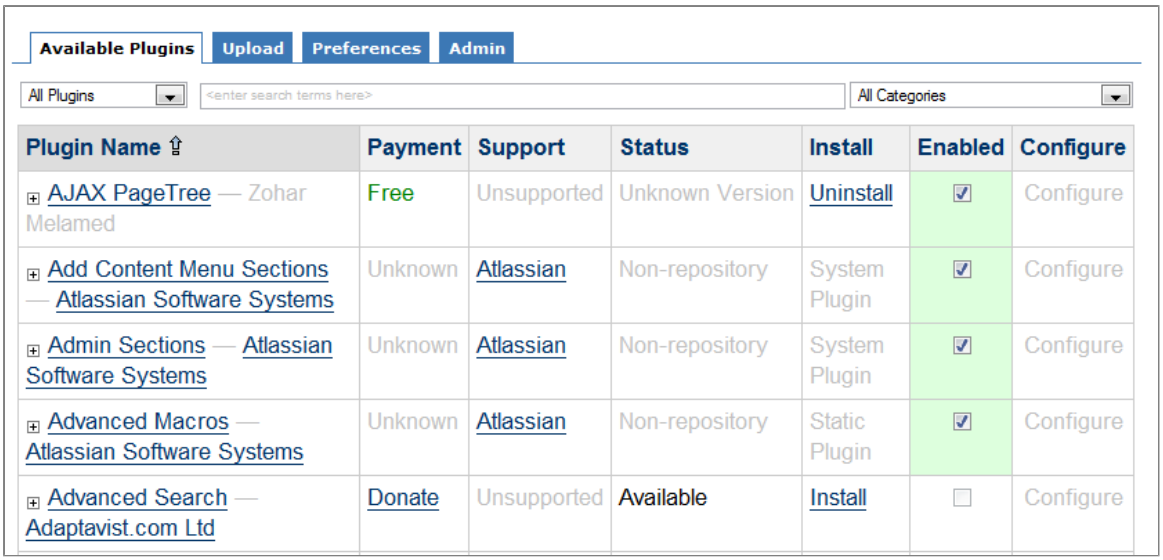
Quick guide to installing a language pack
To install a language pack using the Plugin Repository, you simply need to open the Plugin Repository, locate the language pack and install it via the Plugin Repository interface.

Installing a Language Pack using the Plugin Repository Client

To install a language pack using the Plugin Repository, you simply need to open the Plugin Repository, locate the language pack and install it via the Plugin Repository interface:

Go to the '**Administration Console**' and click '**Plugin Repository**' in the left-hand panel. The following will be displayed:

Screenshot: Plugin Repository



Filtering the List of Plugins

Along the top of the page, you'll see three items which allow you to choose the plugins you want displayed:

- **Status filter** — Defaults to 'All Plugins'. Choose one of the following to limit the list of plugins displayed:
  - 'Installed Plugins' – Plugins which have been installed on your Confluence instance.
  - 'Outdated Plugins' — Plugins for which updates are available
  - 'Available Plugins' — Plugins which are available, but have not yet been installed.
  - 'Supported Plugins' — All plugins which are supported by Atlassian or a third-party.
- **Search** — Enter text in the middle textbox to quickly find what you're looking for:
  - Quickly locate plugins by searching on their title, description, vendor and other details. Just type what you are looking for and let Confluence do the rest
  - The search results are filtered by the status filter (as discussed above), so if you want to see all installed plugins from Adaptavist.com, for example, set the filter to 'Installed Plugins' and type 'Adaptavist.com' into the search box.
- **Categories list** — Filter by category:
  - Simply choose the desired category from the list and only plugins relating to that category will be displayed.
  - Note that plugins can exist in more than one category.

Using the List of Plugins

Under the filter options, the list of plugins matching the current filter settings is shown in a table.

Each column in the table shows information about a particular plugin, and allows you further configuration options:

- **Plugin Name** — Displays the name of the plugin (linked to the detailed information page)and the plugin vendor (linked to their website if applicable)
  - Click the '**+**' icon to expand the information display showing plugin description.
  - Click the '**-**' icon to hide the description again.
- **Payment** — Can be one of the following:
  - Free (self explanatory)
  - Donate (it's free, but you should consider donating to keep it that way)
  - Buy (it's commercial - click the link to show a price list and purchase online)
- **Status** — Shows the current status of this plugin in respect to your Confluence installation:
  - Installed - installed and up-to-date
  - Outdated - installed, but there are new versions available
  - Available - not installed yet
  - Non Repository - a version is installed which is not in the repository
- **Support** — Tells you who supports the plugin:
  - 'Atlassian' — The plugin is supported by Atlassian. If you have any problems, please raise a ticket at the Atlassian Support System.
  - 'Unsupported' — The plugin has been developed by a third party, not by Atlassian, and is not currently supported by Atlassian. In addition, the third party has not yet given detailed information about support arrangements. This does **not** necessarily mean that the plugin is not supported. Please refer to the plugin's home page in the Confluence Extension space or the Atlassian Plugins Exchange site.
  - The 'Support' column can also contain a link to the third-party plugin support site.
  - Read more information about supported plugins.
- **Install** — Install, upgrade or uninstall a plugin:
  - When installing or upgrading, everything is automatic (i.e. it downloads and installs for you, etc). Although the client (since 1.0.2) warns you of dependencies and (since 1.0.3) will do its best to check what has been downloaded is what you asked for - Confluence may break as if you had uploaded the plugins to the Plugin Manger yourself. Where it can, the client will error constructively allowing you to choose the best course of action for yourself. In general, **things usually work - and if they don't its a bug with the client or the metadata**.
  - If the Confluence Repository Client encounters a password request when downloading the plugin (usually case with Commercial plugins), you will be prompted for a username and password.

    - If the plugin is installed into WEB-INF (or otherwise uninstallable) it will display 'Manually Installed. where the actions would be.
- **Enabled** — If the box is ticked, the plugin is enabled, otherwise it's either disabled or partially disabled. You can enable or disable individual modules within the plugin from the plugin details screen (see later).
- **Configure** — If the plugin offers further configuration options, you can click the 'Configure' link. A new screen will open, showing the specific options offered by the plugin.

You can click the table headings to sort the table. Click a second time to reverse the sort.

Installing a Language Pack Manually

To install a language pack manually, you will need to upload the language pack plugin as described below. The language pack plugin will be enabled by default once you have installed it.

> **Finding more language packs**
> You can download official language packs from the Atlassian Plugin Exchange. You can also download language packs developed by the Confluence user community from the Language Pack Translations page.

Plugins are distributed as a jar file. To install a plugin:

1. In the '**Administration**' section of Confluence, click the '**Plugins**' link.
2. Use the '**Choose file**' button to find the plugin jar you wish to install from your hard drive or network location, and select it.
3. Click '**Upload**'.
4. The plugin will be uploaded to Confluence and will be automatically installed.
5. Check the '**Plugin Administration**' screen to ensure if the plugin is available.
6. Enable the plugin if necessary. (Some plugins will be enabled by default when they are installed. Others will have to be manually enabled from the Plugin Administration screen.)

Manually Backing Up The Site

Confluence is configured to make a daily backup of your data and store it as a zipped XML file in the '**backups**' folder under the Confluence Home Directory. A System Administrator can also manually back up the data from the **Administration Console**.

You need to have System Administrator permissions in order to perform this function.

> Consider an Production backup strategy if your Confluence site is large or you are encountering problems with your automated backup.

To manually back up your site,

1. Go to the Confluence '**Administration Console**'. To do this:
    - Open the '**Browse**' menu and select '**Confluence Admin**'. The 'Administration Console' view will open.
2. Click '**Backup & Restore**' in the 'Administration' section of the left-hand panel.
3. Select '**Archive to backups folder**' to store a copy of the backup in the same folder as Confluence's daily backups. (If you do not archive the backup it will be made available for you to download, and then deleted from the server after 24 hours).
4. Select '**Backup attachments**' to include attachments in your backup.
5. Click '**Backup**'.
   Please note that this process will take a few minutes.
6. Once the backup is completed, you will be prompted to download the zipped backup file.

> If you are running Confluence behind Apache and are facing timeout errors, please consider creating the export directly from Tomcat, instead of going through Apache. This will speed up the process and prevent timeouts.

RELATED TOPICS

Site Backup and Restore

Configuring Daily Backups

User Submitted Backup & Restore Scripts

Backup FAQ

Changing time of Daily Backup

Manually Backing Up The Site

Production Backup Strategy

Administrators Guide Home Confluence Documentation Home

Configuring Daily Backups

Confluence backs up your data on a daily basis into a zipped XML file. By default, the backup is performed at 2.00 a.m. and the backup files are stored in the `backups` folder under the Confluence Home Directory.

Time is derived from the Confluence server
The time zone is taken from the server on which Confluence is running. To check the time according to the server, do the following:

1. Go to the Confluence '**Administration Console**'. To do this:
    - Open the '**Browse**' menu and select '**Confluence Admin**'. The 'Administration Console' view will open.
2. Click '**System Information**' in the left-hand panel and look at the '**System Time**'. You can change the time of the daily backup.

Consider an [alternative backup strategy] if your Confluence site is large or you are encountering problems with your automated backup.

The default naming convention for the backup files is 'daily-backup-yyyy_MM_dd'. Confluence can write backups to both local and mapped network drives.

From the **Administration Console**, you can:

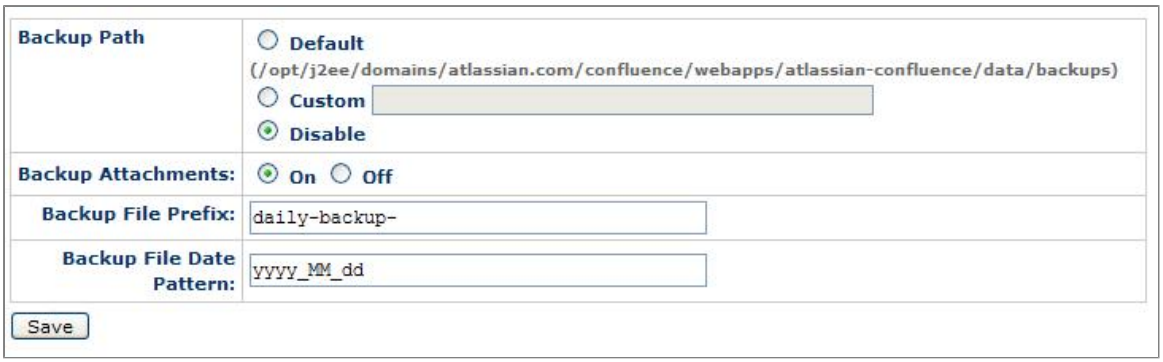
- Enable or disable backups.
- Include or exclude attachments in backups.
- Configure a different path to store backup files.
- Change the naming format used for the files.

You need to have System Administrator permissions in order to perform this function.

To configure your daily backups,

1. Go to the Confluence '**Administration Console**'. To do this:
    - Open the '**Browse**' menu and select '**Confluence Admin**'. The 'Administration Console' view will open.
2. Click '**Daily Backup Admin**' in the '**Configuration**' section.
3. Click the '**Edit**' button on the '**Daily Backup Administration**' screen.
4. Now you can do the following:
    - To disable backups — Select '**Disable**'.
    - To specify an alternate path to store backup files — Select '**Custom**' and then enter the path. The directory must be on either a local drive or a mounted network drive.
      Please ensure the mapped drive is on a physical server, not a Virtual Machine image.
    - To exclude attachments from backups — Select '**Off**' beside '**Backup Attachments**'. By default, this is 'On'.
    - To use a different naming prefix format — Enter the new format in the '**Backup File Prefix**' input field.
    - To use a different date format — Enter the date format in the '**Backup File Date Pattern**' input field using the syntax described in this document from Sun.
5. '**Save**' your changes.

Below is an example of daily backup being *disabled*.



RELATED TOPICS

- Site Backup and Restore
- Configuring Daily Backups
- User Submitted Backup & Restore Scripts
- Backup FAQ
- Changing time of Daily Backup
- Manually Backing Up The Site
- Production Backup Strategy

Administrators Guide Home Confluence Documentation Home

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 the Confluence user forum. Feel free to submit new scripts or post updates by logging in and adding them to the page as a comment.

Delete Old Backups - Wscript Script On Windows

This script examines backup filename and deletes them if necessary, it may need to be edited.

```
01.'If you want 3 day old files to be deleted then insert 3 next to Date - "your number here"
02.'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.
03.'You can always do a wscript.echo strYesterday or strFileName to see what the script thinks you
are searching for.
04.
05.dtmYesterday = Date - 3
06.
07.strYear = Year(dtmYesterday)
08.
09.strMonth = Month(dtmYesterday)
10.If Len(strMonth) = 1 Then
11.          strMonth = "0" & strMonth
12.End If
13.
14.strDay = Day(dtmYesterday)
15.If Len(strDay) = 1 Then
16.          strDay = "0" & strDay
17.End If
18.
19.strYesterday = strYear & "-" & strMonth & "-" & strDay
20.
21.strFileName = "C:\test*." & strYesterday &"-*"
22.
23.Set objFSO = CreateObject("Scripting.FileSystemObject")
24.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:

```
1.ls -t <path to your backup dir>/* | tail -n +6 | xargs -i rm {}
```

Or, using the older form of the `tail` command if your system does not support the standard form:

```
ls -t <path to your backup dir>/* | tail +6 | xargs -i rm {}
```

Delete Old Backups - Advanced Bash Script For Linux

Old XML backups can be deleted automatically by inserting a nightly or weekly automation script or cron similar to the following. Set the BACKUP_DIR and DAYS_TO_RETAIN variables to appropriate values for your site. Between runs, more files than DAYS_TO_RETAIN builds up.

```
#!/bin/sh

# Script to remove the older Confluence backup files.
# Currently we retain at least the last two weeks worth
# of backup files in order to restore if needed.

BACKUP_DIR="/data/web/confluence/backups"
DAYS_TO_RETAIN=14

find $BACKUP_DIR -maxdepth 1 -type f -ctime +$DAYS_TO_RETAIN -delete
```

Manual Database & Home Backup - Bash Script For Linux

This backs up a mySQL database and the Confluence home directory.

```
#!/bin/bash
CNFL=/var/confluence
CNFL_BACKUP=/backup/cnflBackup/`date +%Y%m%d-%H%M%S`

rm -rf $CNFL/temp/*
mkdir $CNFL_BACKUP
mysqldump -uroot -p<password> confluence|gzip > $CNFL_BACKUP/confluence.mysql.data.gz
tar -cjvf $CNFL_BACKUP/data.bzip $CNFL > $CNFL_BACKUP/homedir.status
```

Backup by Date - Postgres

```
export d=`date +%u`
mkdir -p /home/backup/postgres/$d

sudo -u postgres pg_dumpall | bzip2 > /home/backup/postgres/$d/sql.bz2
```

Related Topics

- Site Backup and Restore
- Backup FAQ

Migrating Confluence Between Servers

Some customers have experienced problems with Confluence's search functions after performing a migration, or that the content of their {`recently-updated`} macro is not being updated correctly. Errors in the `atlassian-confluence.log` file corroborate such problems. Hence, to avoid these issues, it is strongly recommended that you perform a rebuild of your content indices after performing a migration.

How to Create a Test or Development Instance



Development licenses are available for any Commercial or Academic license. Create one or contact us for help.

Adminstrators may need to move a Confluence instance from one server to another for upgrades or downtime. This page tells you how to copy a Confluence instance from one server to another. For example, you may want to transfer your current production snapshot to a test server as permitted in the licence agreement.



Avoid upgrades while transferring
If you are planning to switch databases, application servers or Confluence versions, perform the transfer and test that it is successful **separately** to any other changes.

Transferring Confluence To Another Server Using The Same Operating System

If the operating systems on both servers are the same, then the home and install folders can be copied straight into an identical external database and user management setup.

1. On the original server, create zips of the Confluence install and home directories. Copy the zips to the new server.
2. On the new server, unzip the install and home directories. Windows users should avoid unzipping with the Windows built-in extractor, instead use Winzip or the free 7Zip.

   
   If you are changing the location of the home directory, open the Confluence install\confluence\WEB-INF\classes directory and edit confluence-init.properties by changing the line starting with 'confluence.home='.

3. Modify the location of your war file if need be. If using Tomcat, this is likely in /Conf/Catalina/localhost. You'll want to make sure the docbase attribute is pointing to the right location.
4. This next step is dependent on your database:
   a. For users of the internal database, the content is stored inside the home directory. You should switch to an external database after the transfer is successful.
      i. For databases stored on another server: change the user account or datasource permissions so that the new server has the same network access permissions as the original. Then confirm from the new server that the hostname can be resolved and is listening for database connections on the expected port.
      ii. For external databases stored locally: on the original server, create a manual database backup using a native db dump backup tool. Copy the database backup to the new server.
   b. On the new server, install or upgrade the database version to match the original server.
   c. Import the database backup.
   d. Add a database user account with the same username and password as the original.
   e. Provide the user with the full access to the imported database.
   f. Use a database administration tool to confirm that the user can login from the localhost.
   g. To modify any database connection information, go to the Confluence home directory and edit confluence.cfg.xml. The connection URL is set under hibernate.connection.url. **Ensure it does not point to your production database server.**
   h. If you are using internal user management, skip this step. For users who have JIRA or LDAP integration, provide the new server with network or local access to the same hosts as the original. If this is a true test instance, set up a test of your JIRA instance or LDAP server so as not to disrupt production systems and change the server.xml or atlassian-user.xml files to point to the appropriate test servers. Note that it might be acceptable to use a production connection here, as users won't be logging on to the test system in high volume.
   i. If appropriate, make sure no emails are sent out from the test system.
   j. Start Confluence.
   k. Go to Administration > License Details and add your development license key. You can generate one at http://my.atlassian.com. There are more details in Getting a License for a Staging Environment.
   l. If you configured Confluence as a Windows service, repeat those instructions.
   m. Add your development license key.

Transferring Confluence To Another Server Using a Different Operating System

Using database tools (preferred option)

If you are using the Production backup strategy, follow these steps:

1. Download the proper distribution (**the same one you have from your original instance**) from the Download Archive.
2. Copy your Confluence *home* (not install) directory from your original server (even if it was a different OS).
3. If you are changing the location of the home directory, open the Confluence install\confluence\WEB-INF\classes directory and edit confluence-init.properties by changing the line starting with 'confluence.home='.
4. For external databases stored locally, on the original server, create a manual database backup using a native db dump backup tool.
5. Copy the database backup to the new server.
6. On the new server, install or upgrade the database version to match the original server.
7. Import the database backup.
8. Add a database user account with the same username and password as the original.
9. Provide the user with the full access to the imported database.
10. Use a database administration tool to confirm that the user can login from the localhost.
11. To modify any database connection information, go to the Confluence home directory and edit confluence.cfg.xml. The connection URL is set under hibernate.connection.url. **Ensure it does not point to your production database server.**
12. If you are using internal user management, skip this step. For users who have JIRA or LDAP integration, provide the new server with network or local access to the same hosts as the original.
13. Copy server.xml, atlassian-user.xml, osuser.xml, any patches, and any other customized files velocity or properties files. If you are using internal user management, skip this step. For users who have JIRA or LDAP integration, provide the new server with network or local access to the same hosts as the original. If this is a true test instance, set up a test of your JIRA instance or LDAP server so as not to disrupt production systems and change the server.xml or atlassian-user.xml files to point to the appropriate test servers. Note that it might be acceptable to use a production connection here, as users won't be logging on to the test system in high volume.
14. If appropriate, make sure no emails are sent out from the test system.
15. Start Confluence.
16. Go to Administration > License Details and add your development license key. You can generate one at http://my.atlassian.com. There are more details in Getting a License for a Staging Environment.
17. If you configured Confluence as a Windows service, repeat those instructions.
18. Add your development license key.

For XML backups (only for small to medium sized installations)

If you're not yet using the Production backup strategy, you can do this with your regular XML backup. Create a backup and import into the new server.

1. Create a backup from Confluence:
    a. Go to the Confluence '**Administration Console**'. To do this:
        - Open the '**Browse**' menu and select '**Confluence Admin**'. The 'Administration Console' view will open.
    b. Select '**Backup & Restore**'.
    c. Check the '**Backup Attachments**' option and select '**Backup**'.
2. Identify the current version of Confluence your are using, displayed at the bottom of each Confluence page.
3. Download the same version as you are currently using to the new server, which may be the current Confluence release, or an older version.
4. Go to Administration > License Details and add your development license key. You can generate one at http://my.atlassian.com. There are more details in Getting a License for a Staging Environment.
5. Using the same version, follow the Upgrading Confluence guide.
6. Add your development license key.
7. Restore your XML Backup From <<Administration > Backup and Restore>>.
8. If appropriate, make sure no emails are sent out from the test system.

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:

```
1.SELECT * FROM BANDANA WHERE BANDANAKEY = 'atlassian.confluence.smtp.mail.accounts';
```

2. Disable space-level mail archiving by running the following database query:

```
1.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, JIRA macro, etc.). I They write content, or create unwanted load on external systems, they should be disabled promptly after starting the test instance.

Blog post on Moving Confluence from Windows to Linux
Ricky Sheaves (calebscreek) has written an interesting blog post on Moving Confluence from Windows to (Ubuntu) Linux.

Merging instances

If you wish to merge two instances, you can consider using the remote import plugin. This plugin is currently unsupported. The supported method would be to export a space and then import spaces one by one. The two instances must be the same version.

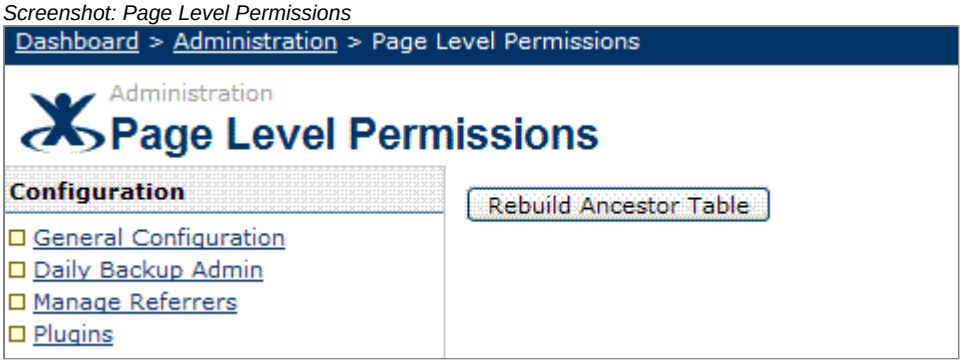
Rebuilding the Ancestor Table

In Confluence, the ancestor table defines what pages are ancestors or descendants of other pages (which can be used by search restrictions with the ancestorids restriction). Occasionally, the ancestor table will become out of sync. When this happens, you can rebuild the table to restore everything to normal.

Simply access this URL:

```
http://yoursite/admin/permissions/pagepermsadmin.action
```

Screenshot: Page Level Permissions



RELATED TOPICS

Administrators Guide Home Confluence Documentation Home

Restoring a Site

CAUTION: Restoring a backup of an entire confluence site (consisting of multiple spaces) will:

- Wipe out all Confluence content in the database. Ensure that your database is backed up.
- Log you out after the restore process. Make sure you know your login details contained in the data being restored.

Atlassian suggests establishing a [backup strategy using a native database tool] for a production instance of Confluence.

Confluence supports **backward compatibility** for site backups. (But **not** for space backups). You can only successfully restore backups of a site from an older version of Confluence to a newer version of Confluence. For example, if you create a site backup in Confluence 2.4.3, it cannot be restored into a Confluence 2.2.2 instance. It can however, be restored into 2.4.5 or 2.5.x, because 2.4.5 and 2.5.x are newer versions of Confluence.

There are two ways to restore a site from a backup file:

1. Restore a site from the Confluence Setup Wizard: This restores the data into a new instance of Confluence.
2. Restore a site from the Administration Console: This restores data into the current instance of Confluence.

If your daily backup zips cannot be restored for whatever reason, but you have backups of both your database and your Confluence home directory, then it is still possible to restore from these backups.

Selective space restore not possible
You cannot select a single space to restore from the entire site backup when the backup contains more than one space.

RELATED TOPICS

Restoring from Backup During Setup
Manually Backing Up The Site
Restoring a Site
Restoring Data from the Administration Console
Restoring a Space
Confluence Docs 3.2

Administrators Guide Home Confluence Documentation Home

Restoring a Space

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

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

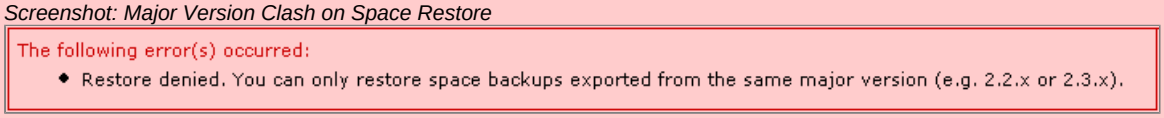
Confluence will only allow you to restore a space if there is not already a space by that name on the site. If you already have a space with the identical name, you will need to delete or rename the existing space before restoring the new one.

Cannot restore to a different major Confluence release
Confluence only supports **forward compatibility** and **backward compatibility** for individual space import and export when executed within the **same major version** of Confluence instances.

Restoration Data Must Share the Same Major Version Number
This means that a space export created in a newer major version of Confluence cannot be imported into an older major version of Confluence. For example, if you create a space export in Confluence 2.4.5, it cannot be imported into a Confluence 2.2.2 instance. It can be however imported into 2.4.6. (because 2.2.2 and 2.4.5 are two different 'major' versions). Similarly, a space export created in 2.2.2 can not be imported into 2.4.5. However, it can be restored in 2.2.10 (since 2.2.2 and 2.2.10 belong to the same major version release).

If such an operation is carried out, an error message similar to the one below will be displayed and the import action will be stopped.

Screenshot: Major Version Clash on Space Restore



Workaround for restoring Spaces between Major Releases

You'll need to set up a test server, download and install the same version of confluence as the version you exported the space from, then import the space into this test server. Next upgrade Confluence on your test installation so it's the right major version so that you can perform the export and import this space into your production confluence successfully. Otherwise, you can try to Change the version of the space export, but please try this on a test instance as well.

You need to have System Administrator permissions in order to perform this function.

To restore a space,

1. Go to the Confluence '**Administration Console**'. To do this:
    - Open the '**Browse**' menu and select '**Confluence Admin**'. The 'Administration Console' view will open.
2. Select '**Backup and Restore**' in the 'Administration' section of the left-hand panel.

You can restore data in one of two ways:

1. **Upload a zipped backup to Confluence:**
    - Browse for the backup file.
    - Uncheck '**Build Index**' if you want to create the index at a later stage.
    - Click '**Upload and Restore**'.
2. **Restore a backup from the file system:**
    - Select the backup file from the form field displayed. If you do not see your backup file, make you sure that it has been copied into the `/opt/java/src/confluence/deployments/conf.atlassian.com/home/restore` directory.
    - Uncheck '**Build Index**' if you want to create the index at a later stage.
    - Click '**Restore**'.

RELATED TOPICS

- Restoring from Backup During Setup
- Manually Backing Up The Site
- Restoring a Site
- Restoring Data from the Administration Console
- Restoring a Space
- Confluence Docs 3.2

Administrators Guide Home Confluence Documentation Home

Changing the version of a space backup

Confluence prevents the import of space backups which aren't from the same major version. The reason for this is that any schema change between the export and imported version of Confluence will cause the import to fail, leaving you with an incomplete import. Even worse, the failure can be database-dependent, so it may work fine on one particular database but your backup will fail to import later.



Do not import a modified space backup on a production server. Import the modified space backup on a test server, then export from the test server to create a pristine space backup for the new version.

To change the version of a space backup, do the following:

- extract the space backup ZIP file
- edit `exportDescriptor.properties` in a text editor
- change the buildNumber to the buildNumber of the Confluence version you wish to import into
- zip up the modified contents of the backup into a ZIP file again.

This will allow you to import a backup into a test instance of Confluence. After checking the imported space for errors, export it cleanly from the test server and import the fresh backup into your production server.

If your import fails on the test server due to Hibernate errors, this indicates a schema incompatibility and cannot be worked around. You will need to restore your entire site on an old version of Confluence, and export the space from there. See the last section of Restoring a Space for details.

Restoring a Test Instance from Production

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.

Upgrading a test Confluence instance with production data

Essentially, we are copying both the production home directory and database to the test instance. We then update the database details on the test instance to point to the test database, leaving all other instance metadata (most importantly the Confluence *build number*) the same as production.

1. Shut down your test instance.
2. Restore the production database to the test database server.
3. Create a backup of the `confluence.cfg.xml` file found in the home directory of the test instance.
4. Copy the production confluence-home directory to the test application server.
5. Open the `confluence.cfg.xml` which has been copied in a text editor. Change the database settings to match the test database server. **Ensure you do not point to your production database.** (You can compare with the backup you made in Step 3 if you need to get the database settings. Don't just copy this file – you need the build number unchanged from production to indicate the database is from an older version of Confluence.)

Before starting your test instance, you need to do the following steps to ensure no contact with production systems.

Ensuring no contact with production systems

To ensure no contact with external systems, you will need to disable both inbound and outbound mail services.

1. Disable global outbound mail by running the following database query:

```
1.SELECT * FROM BANDANA WHERE BANDANAKEY = 'atlassian.confluence.smtp.mail.accounts';
```

2. Disable space-level mail archiving by running the following database query:

```
1.SELECT * FROM BANDANA WHERE BANDANAKEY = 'atlassian.confluence.space.mailaccounts';
```

Change the 'SELECT *' to a 'DELETE' in the above queries once you are sure you want to remove the specified accounts.

Once this is done, you can start your test instance without any mails being sent or retrieved. Think carefully about other plugins which may access production systems (SQL macro, etc.). These should be disabled promptly after starting the test instance.

You can create a developer license for this server and update the License Details after starting up.

See also

Upgrading Confluence
Migrating Confluence Between Servers
Restoring to a Test Instance of Confluence from Production

Restoring Data from other Backups

Typically, Confluence data is restored from the Administration Console or from the Confluence Setup Wizard.

If you are experiencing problems restoring from an zipped XML backup file, it is still possible to restore provided you have:

1. A backup of your home directory.
2. A backup of your database (if you're using an external database).

Instructions for this method of restoring differ depending on whether you are using the embedded database or an external database (like Oracle, MS SQL Server, MySQL or Postgres).

Embedded Database

If you are running against the embedded database, the database is located inside the `database` folder of your Confluence Home Directory. Hence, all you need to do is:

1. Retrieve the most recent backup of your home directory.
2. Unpack the Confluence distribution and point the `confluence-init.properties` file to this directory.

External Database

If you're using an external database, you need to do the following.

1. Prepare backups of your home directory and database (preferably backups that are dated the same). That is, make sure the home directory is accessible on the filesystem and the database available to be connected to.
2. If this database happens to have a different name, or is on a different server, you need to modify the jdbc url in the `confluence.cfg.xml` file inside the Confluence Home Directory. The value of this property is specified as `hibernate.connection.url`.
3. Unpack the Confluence distribution and point the `confluence-init.properties` file to the home directory.

RELATED TOPICS

Important Directories and Files
Migrating to a Different Database

Administrators Guide Home Confluence Documentation Home

Restoring Data from the Administration Console

Use this option if you want to restore data into your current instance of Confluence. If you want to restore data into a new instance, follow the instructions here.

You need to have System Administrator permissions in order to perform this function.

CAUTION: Restoring a backup of an entire Confluence site (consisting of multiple spaces) will do the following:

- Wipe out all Confluence content in the database. Ensure that your database is backed up.
- Log you out after the restore process. Make sure you know your login details contained in the data being restored.

To restore data from backup,

- Go to the Confluence '**Administration Console**'. To do this:
  - Open the '**Browse**' menu and select '**Confluence Admin**'. The 'Administration Console' view will open.
- Select '**Backup and Restore**' in the 'Administration' section of the left-hand panel.

You can restore data in one of two ways:

1. **Upload a zipped backup to Confluence:**
   - Browse for the backup file.
   - Uncheck '**Build Index**' if you want to create the index at a later stage.
   - Click '**Upload and Restore**'.
2. **Restore a backup from the file system:**
   - Select the backup file from the form field displayed. If you do not see your backup file, make sure that it has been copied into the `/opt/java/src/confluence/deployments/conf.atlassian.com/home/restore` directory.
   - Uncheck '**Build Index**' if you want to create the index at a later stage.
   - Click '**Restore**'.

RELATED TOPICS

Restoring from Backup During Setup
Manually Backing Up The Site
Restoring a Site
Restoring Data from the Administration Console
Restoring a Space
Confluence Docs 3.2

Administrators Guide Home Confluence Documentation Home

Retrieve file attachments from a backup

File attachments on pages can be retrieved from a backup without needing to import the the backup into Confluence. This is useful for recovering attachments that have been deleted by users.

Both daily and manual backups allow this, as long as the 'Include attachments' property was set. Users wanting to restore pages, spaces or sites should check out the Confluence Administrator's Guide instead.

Before following the instructions for recovering attachments, please review how backups store file and page information.

How Backups Store File and Page Information

The backup zip file contains entities.xml, an XML file containing the Confluence content, and a directory for storing attachments.

Backup Zip File Structure

Page attachments are stored under the attachments directory by page and attachment id. Here is an example listing:

```
Listing for test-2006033012_00_00.zip
\attachments\98\10001
\attachments\98\10002
\attachments\99\10001
entities.xml
```

Inside the attachment directory, each numbered directory inside is one page, and the numbered file inside is one attachment. The directory number is the page id, and the file number is the attachment id. For example, the file \attachments\98\10001 is an attachment with page id 98 and attachment id 10001. You can read entities.xml to link those numbers to the original filename. Entities.xml also links each page id to the page title.

Entities.xml Attachment Object

Inside the entities.xml is an Attachment object written in XML. In this example, the page id is 98, the attachment id is 10001 and the filename is myimportantfile.doc. The rest of the XML can be ignored:

```
1.<object class="Attachment" package="com.atlassian.confluence.pages">
2.<id name="id">98</id>
3.<property name="fileName"><![CDATA[myimportantfile.doc]]></property>
4....
5.<property name="content" class="Page" package="com.atlassian.confluence.pages"><id name="id"
>10001</id>
6.</property>
7....
8.</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:

```
1.<object class="Page" package="com.atlassian.confluence.pages">
2.<id name="id">98</id>
3.<property name="title"><![CDATA[Editing Your Files]]></property>
4....
5.</object>
```

Instructions for Recovering Attachments

Each file must be individually renamed and re-uploaded back into Confluence by following the instructions below. Choose one of the three methods:

Choice A - Recover Attachments By Filename

Best if you know each filename you need to restore, especially if you want just a few files:

1. Unzip the backup directory and open entities.xml.
2. Search entities.xml for the filename and find the attachment object with that filename. Locate its page and attachment id.
3. Using the page and attachment id from entities.xml, go to the attachments directory and open that directory with that page id. Locate the file with the attachment id.
4. Rename the file to the original filename and test it.
5. Repeat for each file.
6. To import each file back into Confluence, upload to the original page by attaching the file from within Confluence.

Choice B - Restore Files By Page

Best if you only want to restore attachments for certain pages:

1. Unzip the backup directory and open entities.xml.
2. Search entities.xml for the page title and find the page object with that title. Locate its page id.
3. Go to the attachments directory and open that directory with that page id. Each of the files in the directory is an attachment that must be renamed.
4. Search entities.xml for attachment objects with that page id. Every attachment object for the page will have an attachment id and filename.
5. Rename the file with that attachment id to the original filename and test it.
6. Repeat for each page.
7. To import each file back into Confluence, upload to the original page by attaching the file from within Confluence.

Choice C - Restore All Files

Best if you have a small backup but want to restore many or all the attachments inside:



Following process is applicable to **space** export only. Site xml backups do not require page id to be updated manually due to the nature of persistent page_id's.

1. Unzip the backup directory and open entities.xml.
2. Go to the attachments directory and open any directory. The directory name is a page id. Each of the files in the directory is an attachment that must be renamed.
3. Search entities.xml for attachment objects with that page id. When one is found, locate the attachment id and filename.
4. Rename the file with that attachment id to the original filename and test it.
5. Find the next attachment id and rename it. Repeat for each file in the directory.
6. Once all files in the current directory are renamed to their original filenames, search entities.xml for the page id, eg directory name. Find the page object with that page id and locate its page title.
7. Rename the directory to the page title and move on to the next directory. Repeat for each un-renamed directory in the attachments directory.
8. To import each file back into Confluence, upload to the original page by attaching the file from within Confluence.



To obtain detailed information about lost attachments, **location**, **name** and **type** of the attachments, you may use the [findattachments script]

Troubleshooting failed XML site backups



XML site backups are only necessary for migrating to a new database. Setting up a test server or Establishing a reliable backup strategy is better done with an SQL dump.

Seeing an error when creating or importing a backup?

Problem	Solution
Exception while creating backup	Follow instructions below
Exception while importing backup	Follow Troubleshooting XML backups that fail on restore instead

Resolve Errors With Creating An XML Backup

The errors may be caused by a slightly corrupt database. If you're seeing errors such as 'Couldn't backup database data' in your logs, this guide will help you correct the error on your own. We strongly recommend that you backup your database and your Confluence home directory beforehand, so that you can restore your site from those if required. If you are unfamiliar with SQL, we suggest you contact your database administrator for assistance.

Preferable solution

The Production Backup Strategy is a very reliable and more efficient way to do backups. If you are running into problems with XML backups - whether memory related or because of problems like the one described here - use the native backup tool as an alternate solution.

To Identify And Correct The Problem

To work out where the data corruption or problems are, increase the status information reported during backup, then edit the invalid database entry:

1. Stop Confluence.
2. If you have an external database, use a database administration tool to create a manual database backup.
3. Backup your Confluence home directory. You will be able to restore your whole site using this and the database backup.
4. Open the `my_confluence_install/confluence/WEB-INF/classes/log4j.properties` and add this to the bottom and save:

```
1.log4j.logger.com.atlassian.confluence.importexport.impl.XMLDatabinder=DEBUG, confluencelog
2.log4j.additivity.com.atlassian.confluence.importexport.impl.XMLDatabinder=false
```

5. Find your atlassian-confluence.log. Move or delete all existing Confluence logs to make it easier to find the relevant logging output.
6. Restart Confluence and login.
7. Begin a backup so that the error reoccurs.
8. You must now check your log files to find out what object could not be converted into XML format. Open `my_confluence_install/logs/catalina.out`. 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(<generate
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 `catalina.out`, check the first line of the exception. This says there was an error writing the `ContentPermission` object with id 5 into XML. This translates as *the row with primary key 5 in the CONTENTLOCK table* needs fixing. To work out what table an object maps to in the database, here's a rough guide:
    - Pages, blogposts, comments --> CONTENT table
    - attachments --> ATTACHMENTS table
    - More information can be found in the schema documentation
12. Now you must find the primary key of the incorrect row in this table. In this case, you can check the first line and see that the row has a primary key of 5.
13. Each property is written to a column, so the last property that was being written has the incorrect value. The row being written to when the exception was thrown was `CONTENT` (line 5) with a value of `2535` (line 6). Now you know the column and value. This value `2535` is the id of an entry that no longer exists.
14. Using a database administrative tool, login ot the Confluence database. Locate the row in the relevant table and correct the entry. Check other rows in the table for the default column value, which may be null, 0 or blank. Overwrite the invalid row value with the default.
15. Restart Confluence.
16. Attempt the backup again. If the backup fails and you are stuck, please lodge a support request with your latest logs.

Troubleshooting "Duplicate Key" related problems

If you are encountering an error message such as:

```
could not insert:
[bucket.user.propertyset.BucketPropertySetItem#bucket.user.propertyset.BucketPropertySetItem@a70067d3
SQL []; Violation of PRIMARY KEY constraint 'PK_OS_PROPERTYENTRY314D4EA8'. Cannot insert duplicate
key in object 'OS_PROPERTYENTRY'.; nested exception is java.sql.SQLException: Violation of PRIMARY
KEY constraint 'PKOS_PROPERTYENTRY_314D4EA8'. Cannot insert duplicate key in object
'OS_PROPERTYENTRY'.
```

this indicates that the Primary Key constraint 'PK_OS_PROPERTYENTRY_314D4EA8' has duplicate entries in table 'OS_PROPERTYENTRY'.
You can locate the constraint key referring to 'PK_OS_PROPERTYENTRY_314D4EA8' in your table 'OS_PROPERTYENTRY' and locate any duplicate values in it and remove them, to ensure the "PRIMARY KEY" remains unique. An example query to list duplicate entries in the 'OS_PROPERTYENTRY' table is:

```
SELECT ENTITY_NAME,ENTITY_ID,ENTITY_KEY,COUNT(*) FROM OS_PROPERTYENTRY GROUP BY
ENTITY_NAME,ENTITY_ID,ENTITY_KEY HAVING COUNT(*)>1
```

To Help Prevent This Issue From Reoccuring

1. If you are using the embedded database, be aware that it is bundled for evaluation purposes and does not offer full transactional integrity in the event of sudden power loss, which is why an external database is recommended for production use. You should migrate to an external database.
2. If you are using an older version of Confluence than the latest, you should consider upgrading at this point.

RELATED TOPICS

Enabling detailed SQL logging

Administrators Guide Home Confluence Documentation Home

Troubleshooting XML backups that fail on restore



XML site backups are only necessary for migrating to a new database. Upgrading Confluence, Setting up a test server or Production Backup Strategy is better done with an SQL dump.



If migrating from HSQLDB to MySQL, you might have a better experience using the MySQL Migration Toolkit.

Seeing an error when creating or importing a site or space backup?

Problem	Solution
Exception while creating backup	Follow Troubleshooting failed XML site backups instead
Exception while importing backup	Follow instructions below

Resolve Errors When Attempting To Restore An XML Backup

The errors may be caused by a slightly corrupt database. You will need to find the XML backup file entry that is violating the DB rules, modify the entry and recreate the XML backup:

1. On the instance being restored, follow the instructions to disable batched updates (for simpler debugging), log SQL queries and log SQL queries **with parameters** at Enabling detailed SQL logging.
2. Once all three changes have been made, restart Confluence.
3. Attempt another restore.
4. Once the restore fails, check your log files to find out what object could not be converted into XML format. For Confluence Standalone users, check your Confluence install directory under the `/logs/` and check both `atlassian-confluence.log` and `catalina.out` file. The correct file will contain SQL debug output.
5. Scroll to the bottom of the file and identify the last error relating to a violation of the database constraint. For example:

```
2006-07-13 09:32:33,372 ERROR [confluence.importexport.impl.ReverseDatabinder] endElement
net.sf.hibernate.exception.ConstraintViolationException:
  could not insert: [com.atlassian.confluence.pages.Attachment#38]
net.sf.hibernate.exception.ConstraintViolationException: could not insert:
[com.atlassian.confluence.pages.Attachment#38]
...
Caused by: java.sql.SQLException: ORA-01400: cannot insert NULL into
("CONFUSER"."ATTACHMENTS"."TITLE")
at oracle.jdbc.driver.DatabaseError.throwSqlException(DatabaseError.java:112)
at oracle.jdbc.driver.T4CTTIoer.processError(T4CTTIoer.java:331)
at oracle.jdbc.driver.T4CTTIoer.processError(T4CTTIoer.java:288)
```

This example indicates a row in your attachment table with ID = 38 that has a null title.

6. Go to the server that the backup was created on. You must have a copy of the database from which the backup was created. If you do not have this, use a DBA tool to restore a manual backup of the database.
7. Open a DBA tool and connect to the original database instance and scan the table names in the schema. You will have to modify a row in one of these tables.
8. To work out which table, open `catalina.out`, check the first line of the exception. To work out what table an object maps to in the database, here's a rough guide:
    - Pages, blogposts, comments --> CONTENT table.
    - attachments --> ATTACHMENTS table.
9. To correct the example error, go to the attachment table and find that attachment object with id 38. This will have a a null title. Give a title using the other attachments titles as a guide. You may have a different error and should modify the database accordingly.
10. Once the entry has been corrected, create the XML backup again.
11. Import the backup into the new version.
12. If the import succeeds, revert the changes made in your SQL logging to re-enable disable batched updates and turn off log SQL queries and log SQL queries with parameters.
13. Restart Confluence.

Troubleshooting "Duplicate Key" related problems

If you are encountering an error message such as:

```
could not insert:
[bucket.user.propertyset.BucketPropertySetItem#bucket.user.propertyset.BucketPropertySetItem@a70067d3
SQL []; Violation of PRIMARY KEY constraint 'PK_OS_PROPERTYENTRY314D4EA8'. Cannot insert duplicate
key in object 'OS_PROPERTYENTRY'.; nested exception is java.sql.SQLException: Violation of PRIMARY
KEY constraint 'PKOS_PROPERTYENTRY_314D4EA8'. Cannot insert duplicate key in object
'OS_PROPERTYENTRY'.
```

This indicates that the Primary Key constraint 'PK_OS_PROPERTYENTRY_314D4EA8' has duplicate entries in table 'OS_PROPERTYENTRY'.
You can locate the constraint key referring to 'PK_OS_PROPERTYENTRY_314D4EA8' in your table 'OS_PROPERTYENTRY' and locate any duplicate values in it and remove them, to ensure the "PRIMARY KEY" remains unique. An example query to list duplicate entries in the 'OS_PROPERTYENTRY' table is:

```
SELECT ENTITY_NAME,ENTITY_ID,ENTITY_KEY,COUNT(*) FROM OS_PROPERTYENTRY GROUP BY
ENTITY_NAME,ENTITY_ID,ENTITY_KEY HAVING COUNT(*)>1
```

Troubleshooting "net.sf.hibernate.PropertyValueException: not-null" related problems

If you're receiving a message like:

```
ERROR [Importing data task] [confluence.importexport.impl.ReverseDatabinder] endElement
net.sf.hibernate.PropertyValueException: not-null property references a null or transient value:
com.atlassian.user.impl.hibernate.DefaultHibernateUser.name
```

This means there's an unexpected null value in a table. In the above example, the error is in the name column in the USERS table. We've also seen them in the ATTACHMENTS table.

Remove the row with the null value, redo the xml export, and reimport.

To Help Prevent this Issue from Reccuring

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

RELATED TOPICS

Troubleshooting failed XML site backups
Confluence Administrator's Guide

Migrating from HSQLDB to MySQL



If you've gone through Migrate to Another Database and cannot migrate because of a failed xml backup, this page might help.

Disclaimer

MySQL Migration Toolkit is released by the makers of MySQL and as such, problems with the software should be directed to them. Atlassian Support does not offer support for the Migration Toolkit, nor do we provide support for this migration path. These instructions are offered for strictly informational purposes, and your mileage may vary.



Backup Reminder
Please backup your database and your home folder before attempting this.

Resources needed:

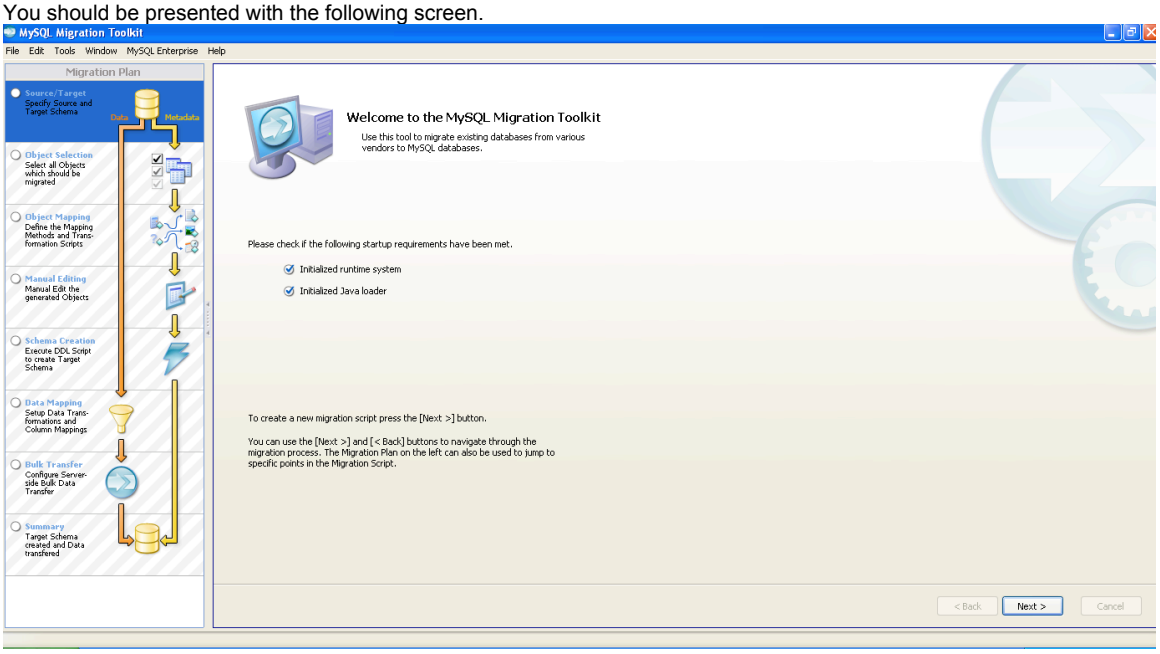
- Empty MySQL DB with appropriate credentials to allow creation, deletion, and insertion of tables and rows.
- A Windows machine that can both communicate to the Confluence server and the destination DB.
- MySQL Migration Toolkit
- HSQL Database Engine

Preparation for migrating to MySQL from HSQLDB

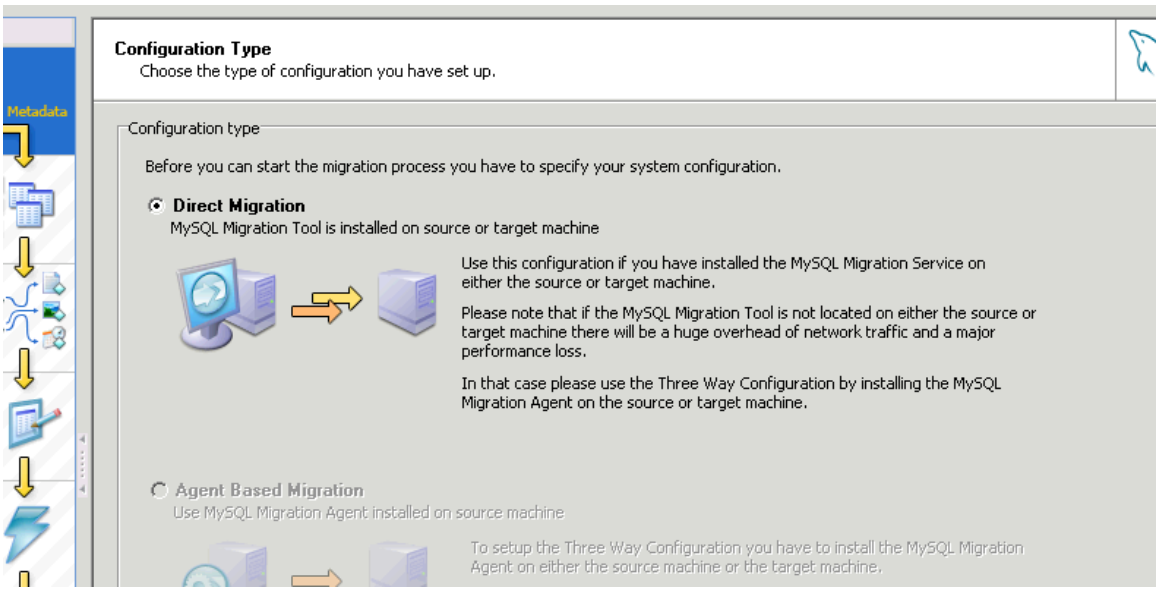
1. Shutdown Confluence
2. Make a copy of the confluence home folder for backup purposes
3. Install the Migration Toolkit
4. Unzip the hsqldb package.
5. Copy the hsqldb.jar from hsqldb/lib into C:\Program Files\MySQL\MySQL Tools for 5.0\java\lib
6. Start the MySQL Migration Toolkit

Running the Migration Toolkit

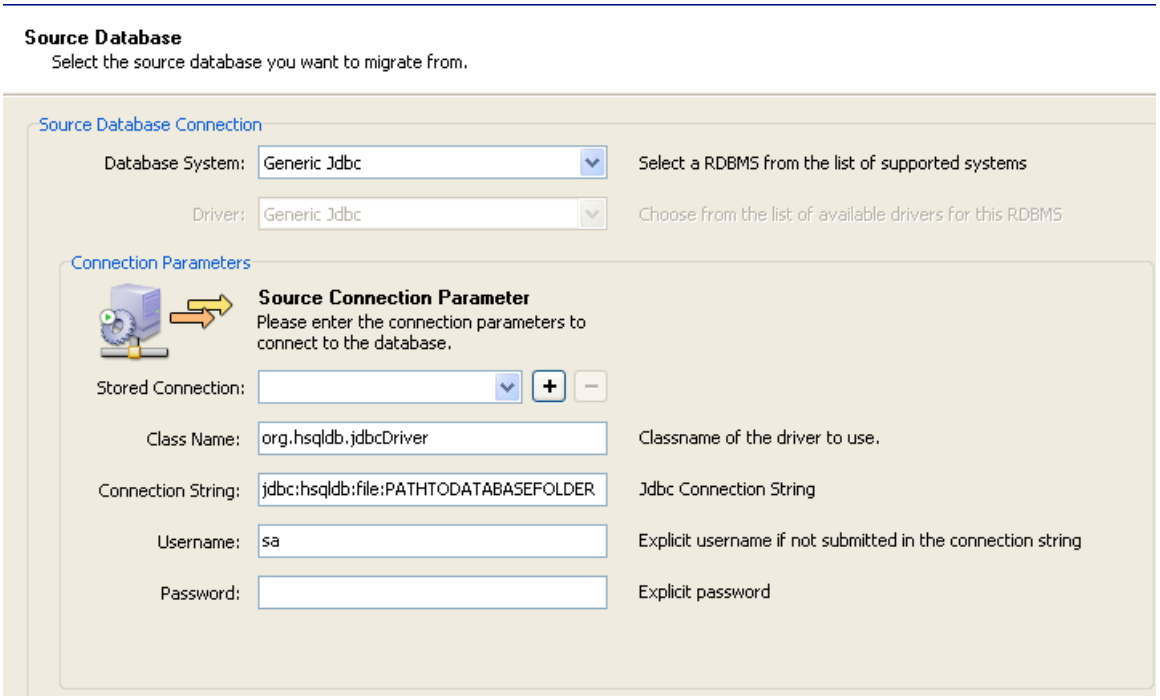
You should be presented with the following screen.



Choose Direct Migration



Source Database

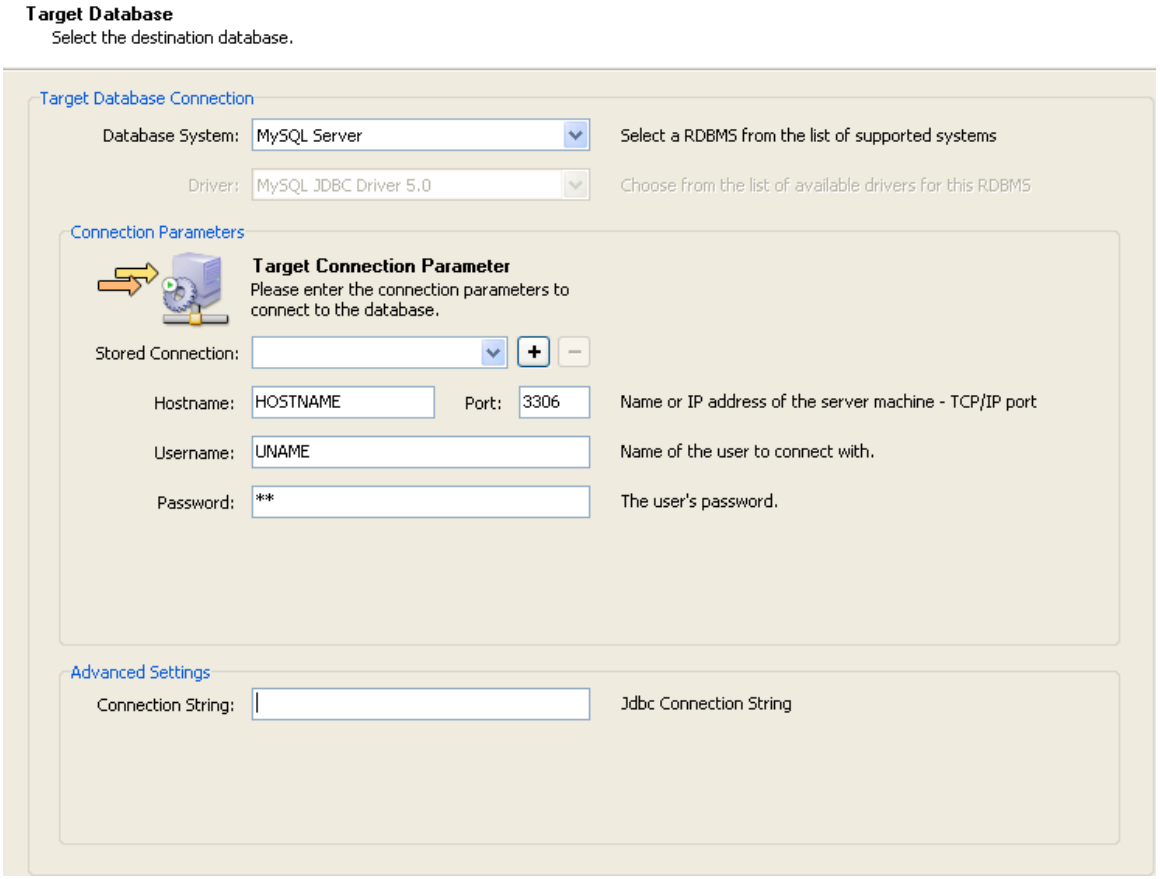


Database System:	Generic JDBC
Connection String:	jdbc:hsqldb:file:PATHTODATABASEFOLDER\confluencedb
Username:	sa
Password:	*No password. Leave this field blank*

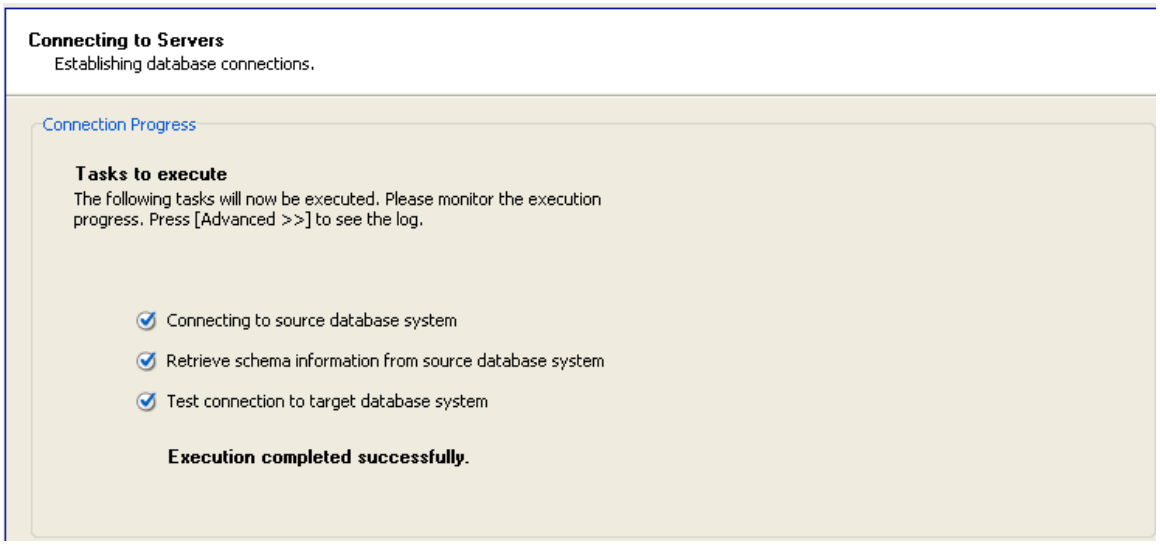
Destination Database

Please make sure that the computer that is running MySQL Toolkit is able to access the MySQL server and that the user listed has the ability to create, drop, insert, and update tables.

If your MySQL user has a $ character in the password (such as 'pa$sword'), please change the password or create a temporary account with full permissions. If you do not, the toolkit will throw an "Illegal group reference" error and you will not be able to proceed with the migration.



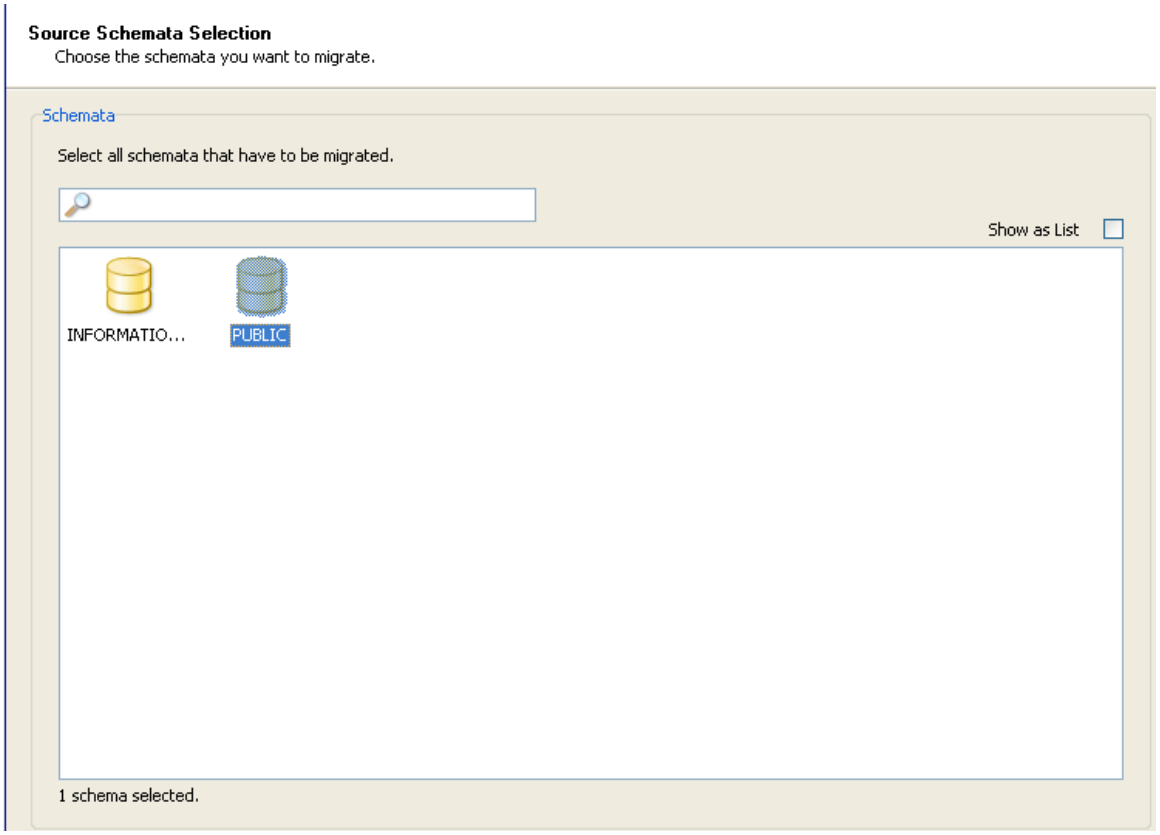
Connecting to Servers



You should see the toolkit trying to connect. If you have problems, please click on the advanced options and sql will show you debugging information. Click Advanced to see the log. If you see "Java Heap Space: Out of Memory", you can start the *MySQL Migration Toolkit* with a -Xmx flag to allocate more memory to the JVM.

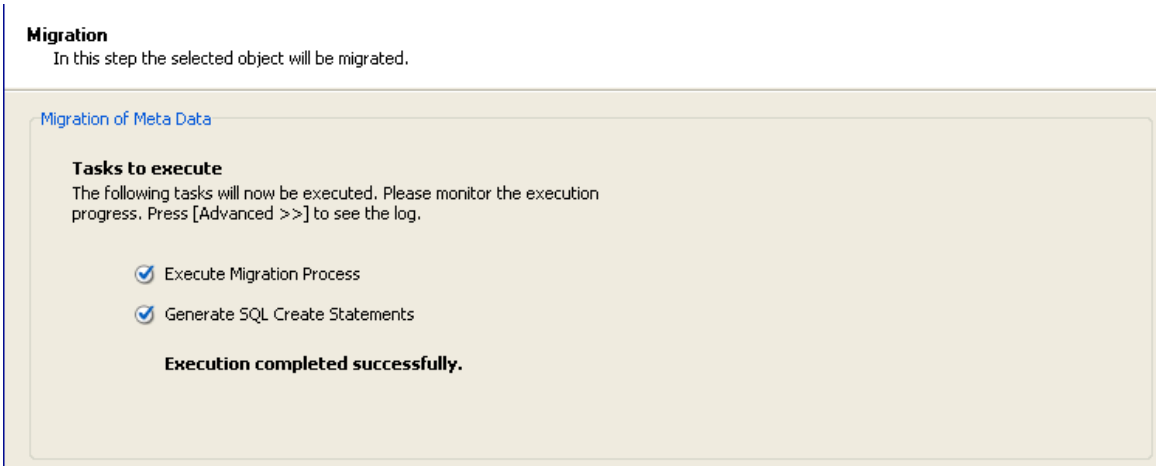
After this screen you should come to reverse engineering. Click next.

Source Schemata Selection



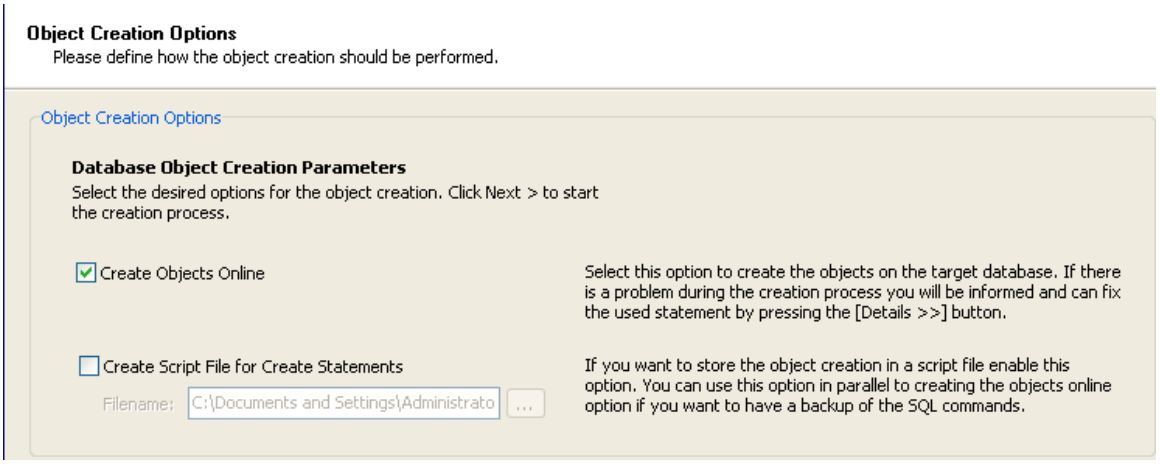
You should see 2 databases, **INFORMATION_SCHEMA** and **PUBLIC**. Choose **PUBLIC**

Object Type Selection



Click Next.

Object Type Mapping



Click **Show Details** on both sections. For **Migration Method for Type Schema**, choose **Multilanguage**. For **Migration Method for Type Table**, choose **Data Consistancy/Multilanguage**

Click **Advanced**. Check **Enabled Detailed Mappings in Next Step**

Detailed Object Mapping

Click to rename the **destination database** to be the one set aside to migrate to.

From this point on, you should be able to click next all the way through to finish the migration.

Viewing and Editing License Details

The '**License Details**' page tells you:

- How many users your Confluence instance is licensed to support, and how many are currently registered.
  Note: The number of registered users only includes users who have '**can use Confluence**' permission. Deactivated users are not included.
  Click the '**Refresh**' button to make sure you see the latest count.
- What type of license you have (e.g. Commercial, Academic, Community).
- How much time remains in your one-year support and upgrades period (for full licenses) or 30-day trial (for trial licenses).
- Your server ID, which:
  - is generated when you install Confluence for the first time
  - exists for the life of the Confluence instance
  - survives an upgrade
  - is held in the database
  - is not bound to a specific license
  - is the same for all servers in a cluster.

To view the details of your Confluence license,

1. Log into Confluence as a user with Confluence Administrator or System Administrator permissions.
2. Go to the Confluence '**Administration Console**'. To do this:
   - Open the '**Browse**' menu and select '**Confluence Admin**'. The 'Administration Console' view will open.
3. Click '**License Details**' under the heading '**Administration** in the left-hand panel.

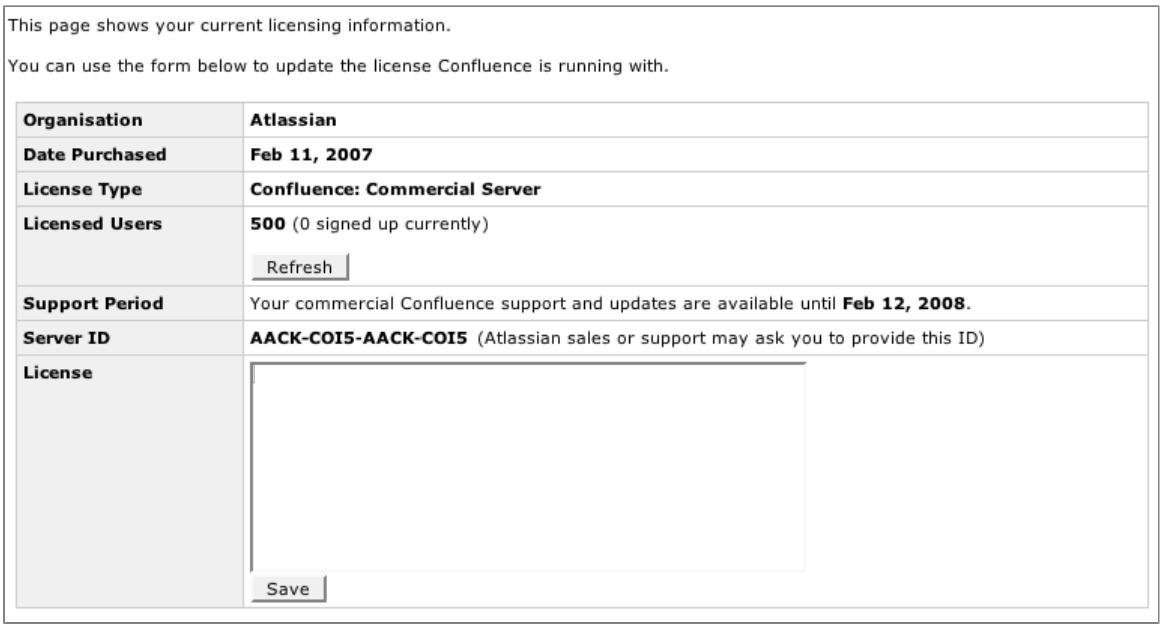
To look up your license details on the Atlassian website, please log in to my.atlassian.com

To update your Confluence license,

1. Log into Confluence as a user with Confluence Administrator or System Administrator permissions.
2. Go to the Confluence '**Administration Console**'. To do this:
    - Open the '**Browse**' menu and select '**Confluence Admin**'. The 'Administration Console' view will open.
3. Click '**License Details**' under the heading '**Administration** in the left-hand panel.
4. Enter your new license details into the '**License**' field and click the '**Save**' button.

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

Screenshot : License Details



Downgrading your Confluence license (reducing the number of allowed users)
If you need to downgrade your Confluence license to one which allows fewer users, please make sure first that your new license covers your current user base.

- View your license details as described above.
- Verify that the number of users '**signed up currently**' is lower than the number allowed by the new license.
- If you currently have more users signed up than the new license allows, please follow these instructions on removing users from your Confluence site.

RELATED TOPICS

- Viewing and Editing License Details
- How Do I Find My License from the File System?
- Getting a License for a Staging Environment
- Cache Performance Tuning
- Cache Performance Tuning for Specific Problems
- Confluence Cache Schemes
- Cache Statistics
- Viewing System Information

Administrators Guide Home Confluence Documentation Home

Viewing System Information

The System Information screen provides information about Confluence's configuration, and the environment in which Confluence has been deployed. Your system configuration information is helpful to us when diagnosing errors you may face using Confluence. If you file a support request or bug report, the more detail you can provide about your installation and environment the faster we will be able to help.

To view your system information,

1. Go to the Confluence '**Administration Console**'. To do this:
    - Open the '**Browse**' menu and select '**Confluence Admin**'. The 'Administration Console' view will open.
2. Click '**System Configuration**' in the 'Administration' section.



The handy **Memory Graph** helps you keep track of Confluence's memory usage.

RELATED TOPICS

- Cache Statistics
- Site Statistics
- Viewing and Editing License Details
- Viewing and Managing Installed Plugins
- Live Monitoring Using the JMX Interface

Administrators Guide Home Confluence Documentation Home

Live Monitoring Using the JMX Interface



Like the Heisenberg uncertainty principle, adding live monitoring to a production instance may have an impact itself on performance!

With the JMX interface (introduced in Confluence 2.8), you can monitor the status of your Confluence instance in real time. This will provide you with useful data such as the resource usage of your instance and its database latency, allowing you to diagnose problems or performance issues. To read the JMX data, you will need to use a JMX client.

Disable JMX



If you experience any problems during Confluence startup that are related to JMX, it is possible to disable the JMX registration process. Please place jmxContext.xml in your `<confluence-install>/confluence/WEB-INF/classes` folder to do so.

What is JMX?

JMX (Java Management eXtensions) is a technology for monitoring and managing Java applications. JMX uses objects called MBeans (Managed Beans) to expose data and resources from your application.

1. Enabling JMX Remote with Tomcat

By default, Confluence uses the Apache Tomcat web server. To use JMX, you must enable it on your Tomcat server, by carrying out the steps under the Apache Tomcat documentation, entitled Enabling JMX Remote. With those steps completed, restart your Tomcat server.

For the stand-alone, add the startup parameter -Dcom.sun.management.jmxremote to setenv.sh or setenv.bat. See instructions for the Windows Service - enter it in the same place as PermGen Memory.

2. Selecting your JMX Client

You need to use a JMX client in order to view the JMX output from Confluence. JConsole is a readily available JMX client that is included with Sun's Java Developer Kit (version 5 onwards). The full name is the '*Java Monitoring and Management Console*', but we will refer to it as JConsole for the purposes of this document.

3. Adding the JMX Client to your Path

You must add the location of the JConsole binary file to your '**path**' environment variable. As JConsole resides in the 'bin' (binaries) folder under your Java directory, the path should resemble something like this:

```
1.JDK_HOME/bin/
```

In this example, replace 'JDK_HOME' with the full system path to your Java directory.

4. Configuring JConsole

To configure JConsole,

1. Run the JConsole application.
2. You will be prompted to create a new connection. Choose '**remote process**', enter the hostname of your Confluence instance and a port of your choosing.



To connect easily, add the startup parameters to setenv.bat or setenv.sh:
-Dcom.sun.management.jmxremote -Dcom.sun.management.jmxremote.port=8086
-Dcom.sun.management.jmxremote.authenticate=false
Port 8086 is unlikely to be used. Then, connect remotely using port 8086.



JConsole, or any JMX client, will not see applications which are not owned by the same user. For example under Windows, if an application is started as a service, it is the System User which owns the process, and not the Current User.

3. Click '**Connect**'.

Note: Other JMX clients besides JConsole can read JMX information from Confluence.

What can I monitor with JMX?

The JMX interface allows you to see live internal information from your Confluence instance, via the following MBeans:

IndexingStatistics

This MBean shows information related to search indexing.

Property name	Function	Values
Flushing	Shows state of cache (i.e. flushing, or not).	True/False
LastElapsedMilliseconds	Time taken during last indexing.	Milliseconds
LastElapsedReindexing	Time taken during last re-indexing.	Milliseconds
TaskQueueLength	Shows number of tasks in the queue.	Integer

SystemInformation

This MBean shows information related to database latency. It also contains most of the information presented on the System Information page.

Property name	Function	Values
DatabaseExampleLatency	Shows the latency of an example query performed against the database.	Milliseconds

RequestMetrics

This MBean shows information related to system load and error pages served.

Property name	Function	Values
AverageExecutionTimeForLastTenRequests	Average execution time for the last ten requests.	Milliseconds
CurrentNumberOfRequestsBeingServed	Number of requests being served at this instant.	Integer
ErrorCount	Number of times the Confluence error page was served.	Integer

MailServer-SMTPServer

This MBean shows information related to email dispatch attempts and failures. There will be an MBean for every SMTP Mailserver that has been configured in the Confluence instance.

Property name	Function	Values
EmailsAttempted	The number of email messages Confluence has tried to send.	Integer
EmailsSent	The number of email messages sent successfully.	Integer

MailTaskQueue

This MBean shows information related to the email workload.

Property name	Function	Values
ErrorQueueSize	Number of errors in the queue.	Integer
Flushing	Shows state (i.e. flushing, or not)	True/False
FlushStarted	Time that operation began.	Time
RetryCount	The number of retries that were performed.	Integer
TaskSize	Number of email messages queued for dispatch.	Integer

SchedulingStatistics

This MBean shows information related to current jobs, scheduled tasks and the time that they were last run.

High CPU consuming threads

For Java 1.6, add the Top Threads Plugin to monitor whether CPU is spiking. Download it to a directory and run JConsole like this:
JConsole -pluginpath /pathto/topthreads.jar

This works only with jdk 1.6, but that can be on the remote machine if the server is running a lower version.

RELATED TOPICS

- Viewing System Information
- Cache Statistics
- Viewing and Editing License Details
- Viewing and Managing Installed Plugins

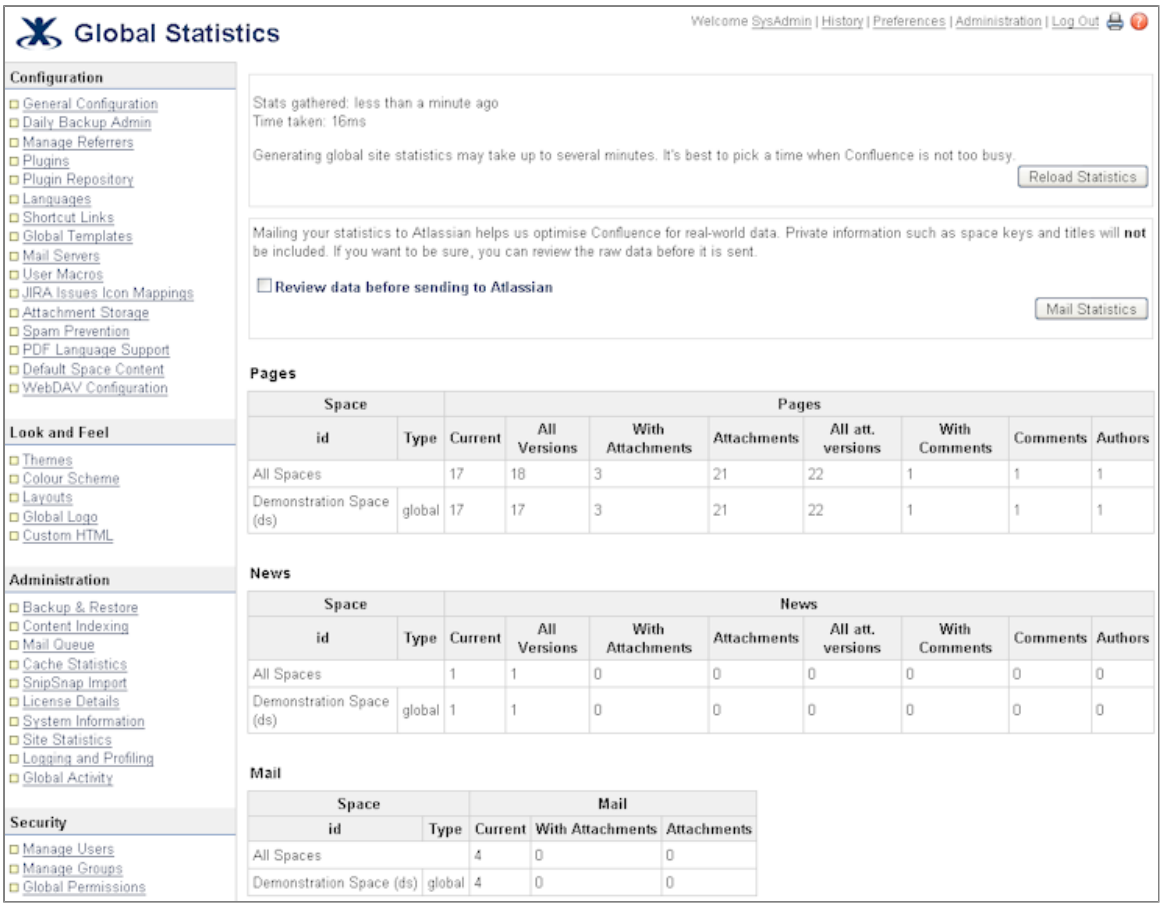
Site Statistics

Understanding Site Statistics

Site Statistics allows you to view a breakdown of pages and editing activity on your Confluence instance.

You can also choose to email your site statistics information to Atlassian, with the '**Mail Statistics**' button. This may be useful for troubleshooting purposes, or when requesting Atlassian support. You can also review the data before it is sent (if you are concerned that it may contain sensitive information).

Screenshot: Confluence Site Statistics



Viewing Confluence Site Statistics

To view Confluence Site Statistics,

1. Log into Confluence as a user with Confluence Administrator or System Administrator permissions.
2. Go to the Confluence '**Administration Console**'. To do this:
    - Open the '**Browse**' menu and select '**Confluence Admin**'. The 'Administration Console' view will open.
3. Select '**Site Statistics**' under the heading 'Administration' in the left-hand panel.
4. Click '**Load Statistics**'. For a large Confluence installation, this may affect the server's performance. For this reason, we suggest you do this at a 'quiet time' for the system.
5. Site Statistics will be displayed.

The Global Statistics Plugin is unsupported.

If the '**Site Statistics**' option is not shown, you may have to install the **Global Statistics Plugin**. In this case, follow the instructions for installing plugins and look for the '**Global Statistics Plugin**'.

If you want to view page-view statistics, this is possible with the SQL Plugin. Some examples are available on the Confluence Community Pages. Please note that some plugins are not officially supported by Atlassian.

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 Confluence Standalone

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.

RELATED TOPICS

Editing Files within JAR Archives
Where are the files that used to be in my Confluence installation directory?

Administrators Guide Home Confluence Documentation Home

Configuring Confluence

- Site Configuration
    - Editing the Site Title
    - Editing the Global Logo
    - Configuring the Site Home Page
    - Editing the Site Welcome Message
    - Configuring the Server Base URL
    - Configuring the Site Support Address
    - Configuring the Destination of View Space Links
    - Customising Default Space Content
    - Showing Link Icons
- Configuring Encoding
    - Character encodings in Confluence
    - Troubleshooting Character Encodings
        - "€" Euro character
        - MySQL 3.x Character Encoding Problems
- Configuring Mail
    - Configuring a Server for Outgoing Mail
    - Enabling the 'Mail Page' plugin
    - The Mail Queue
- Optional Settings
    - Attachment Storage Configuration
        - Hierarchical File System Attachment Storage
    - Configuring Quick Navigation
    - Enabling CamelCase Linking
    - Enabling OpenSearch
    - Enabling Remote APIs
    - Enabling Rich Text Editing Option
    - Enabling the Did You Mean Feature
    - Enabling Threaded Comments
    - Enabling Trackback
    - Making Rich Text Editing default
    - WebDAV Configuration
- Other Settings
    - Configuring Attachment Size
    - Configuring Character Encoding
    - Configuring HTTP Timeout Settings
    - Configuring Indexing Language
    - Configuring JIRA Issues Icon mappings
    - Configuring Number Formats
    - Configuring Shortcut Links
    - Configuring Time and Date Formats
    - Number of Ancestors to Show in Breadcrumbs
    - Thumbnail Settings
- Configuring System Properties
    - Recognised System Properties
- Configuring Logging
- External Gadgets

Administrators Guide Home Confluence Documentation Home

Site Configuration

- Editing the Site Title
- Editing the Global Logo
- Configuring the Site Home Page
- Editing the Site Welcome Message
- Configuring the Server Base URL
- Configuring the Site Support Address
- Configuring the Destination of View Space Links
- Customising Default Space Content
- Showing Link Icons

Editing the Site Title

The site title appears in your browser's title bar. By default, it is set to 'Confluence'.

To change the title of your Confluence site,

1. Go to the '**Administration Console**' and click '**General Configuration**' in the left-hand panel.
2. Click '**Edit**' at the top of the '**Site Configuration**' screen.
3. Enter a new title for your site in the input field next to '**Site Title**'.
4. Click '**Save**'.

RELATED TOPICS

- Configuring the Site Home Page
- Configuring the Destination of View Space Links
- Configuring the Server Base URL
- Configuring the Site Support Address
- Customising Default Space Content
- Editing the Global Logo
- Editing the Site Title
- Editing the Site Welcome Message
- Showing Link Icons

Administrators Guide Home Confluence Documentation Home

Editing the Global Logo

By default, the global logo appears beside the page title on all pages in the site. You can disable the logo or replace it with one of your own.

To edit the global logo,

1. From the '**Administration Console**' click on '**Global logo**' under the heading '**Look and Feel**' in the left panel.
2. In the screen displayed, select '**Off**' to disable logo.
3. To upload a new logo, click '**Browse**' to select a new image and click '**Upload Logo**'.

RELATED TOPICS

- Configuring the Site Home Page
- Configuring the Destination of View Space Links
- Configuring the Server Base URL

Configuring the Site Support Address
Customising Default Space Content
Editing the Global Logo
Editing the Site Title
Editing the Site Welcome Message
Showing Link Icons

Administrators Guide Home Confluence Documentation Home

Configuring the Site Home Page

You can configure Confluence to direct users to any of the space home pages on the site when they log in, rather than to the dashboard.

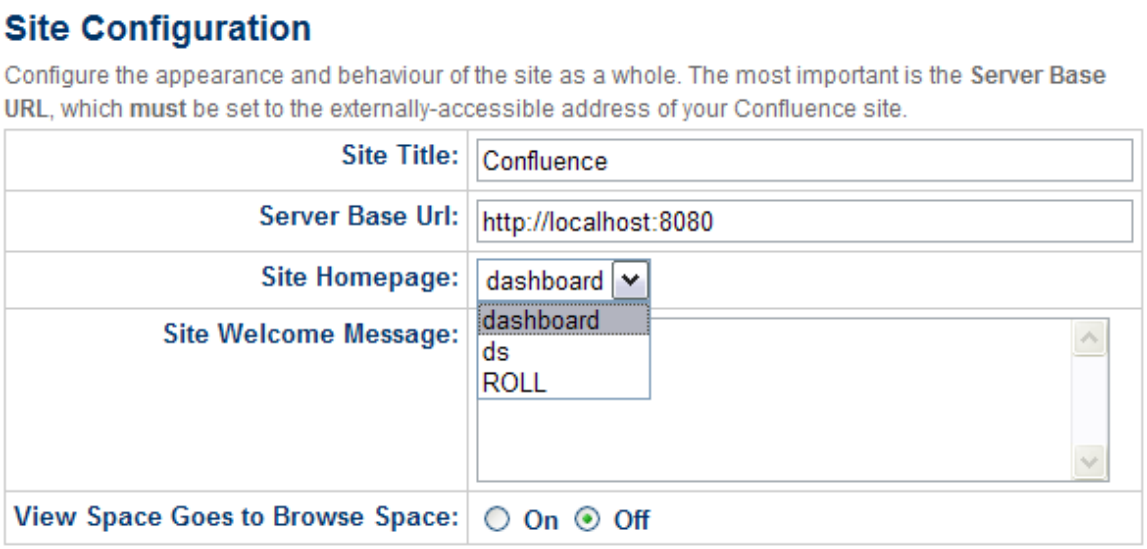
To configure the site-wide home page,

1. Go to the '**Administration Console**' and click '**General Configuration**' in the left-hand panel.
2. Click '**Edit**' next to the '**Site Configuration**' panel.
3. Select a space from the '**Site Homepage**' dropdown menu. When users log in, Confluence will open the home page of the space you choose here.
4. Ensure that the '**View Space Goes to Browse Space**' option is set to '**Off**' if you want users to be sent to the space **home page** and not the **space summary page**.
5. Click the '**Save**' button at the bottom of the screen.

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

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

Screenshot : Configuring the site homepage



RELATED TOPICS

Configuring the Site Home Page
Configuring the Destination of View Space Links
Configuring the Server Base URL
Configuring the Site Support Address

Customising Default Space Content
Editing the Global Logo
Editing the Site Title
Editing the Site Welcome Message
Showing Link Icons

Administrators Guide Home Confluence Documentation Home

Editing the Site Welcome Message

The site welcome message appears at the top left of the Confluence dashboard, between the site logo and the list of spaces. You can use it to display an introduction to the site or a message of the day.

To edit the site welcome message,

1. Go to the '**Administration Console**' and click '**General Configuration**' in the left-hand panel.
2. Click '**Edit**' at the top of the '**Site Configuration**' section.
3. Enter your text in the text box next to '**Site Welcome Message**'. You can enter any text or Confluence wiki markup.
4. Click '**Save**'.

On this page:

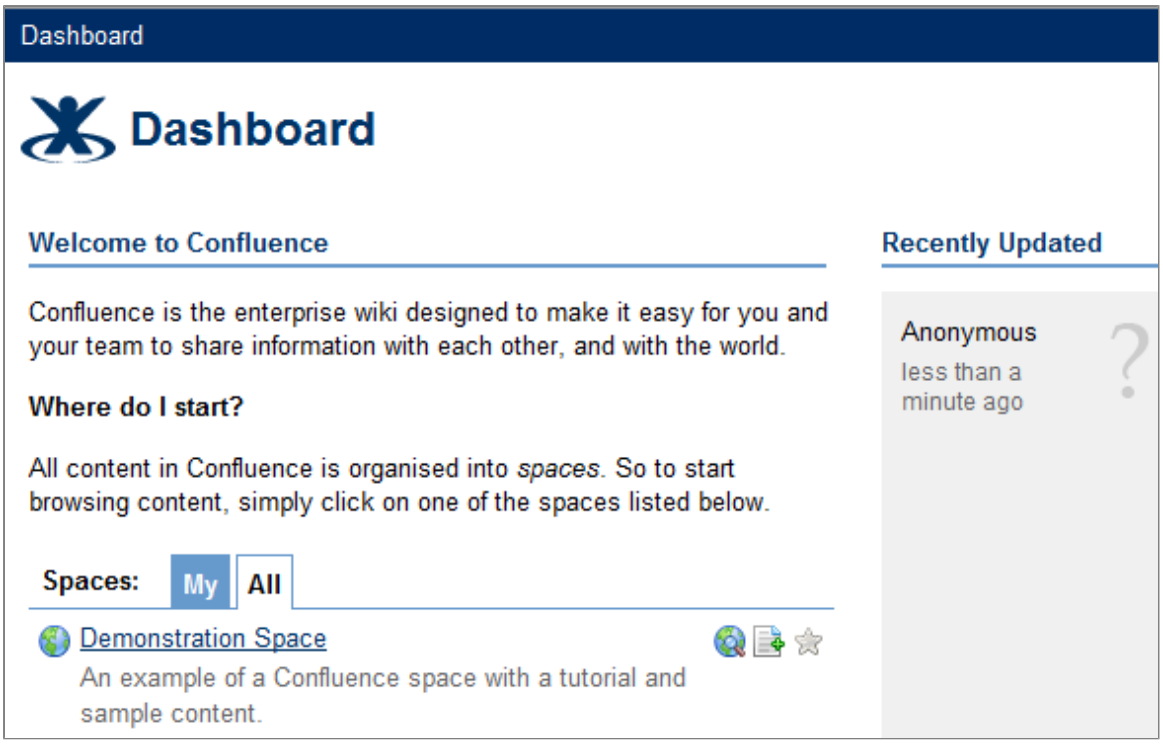
- The Default Site Welcome Message
- Example 1. Adding a Simple Welcome Message
- Example 2. Formatting your Welcome Message
- Example 3: Including Content from Another Page
- Example 4. Adding Blog Posts Filtered by Labels to your Welcome Message
- How we Use the Site Welcome Message at Atlassian

The Default Site Welcome Message

By default, the site welcome message looks more or less like the screenshot below, starting with the words '**Welcome to Confluence**' and ending above the list of spaces.

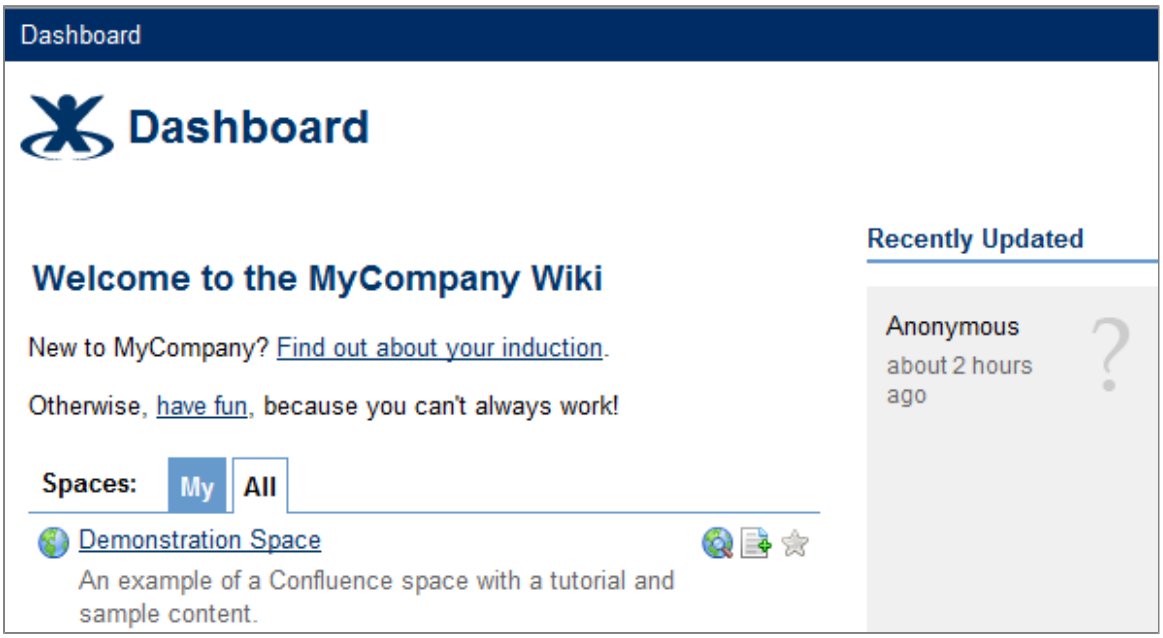
To restore the default site welcome message and remove your customised message, just delete the text in the '**Site Welcome Message**' text box. Provided that you have not customised Confluence, your Confluence users will see the default message if there is no text in the '**Site Welcome Message**' text box in your Administration Console.

Screenshot: Site welcome message at top left of the dashboard



Example 1. Adding a Simple Welcome Message

Let's say you want to display a simple message like this at the top of your dashboard:



To produce the above welcome message, follow the step-by-step instructions above and add the following wiki markup into the '**Site Welcome Message**' text box:

```
h2. Welcome to the MyCompany Wiki

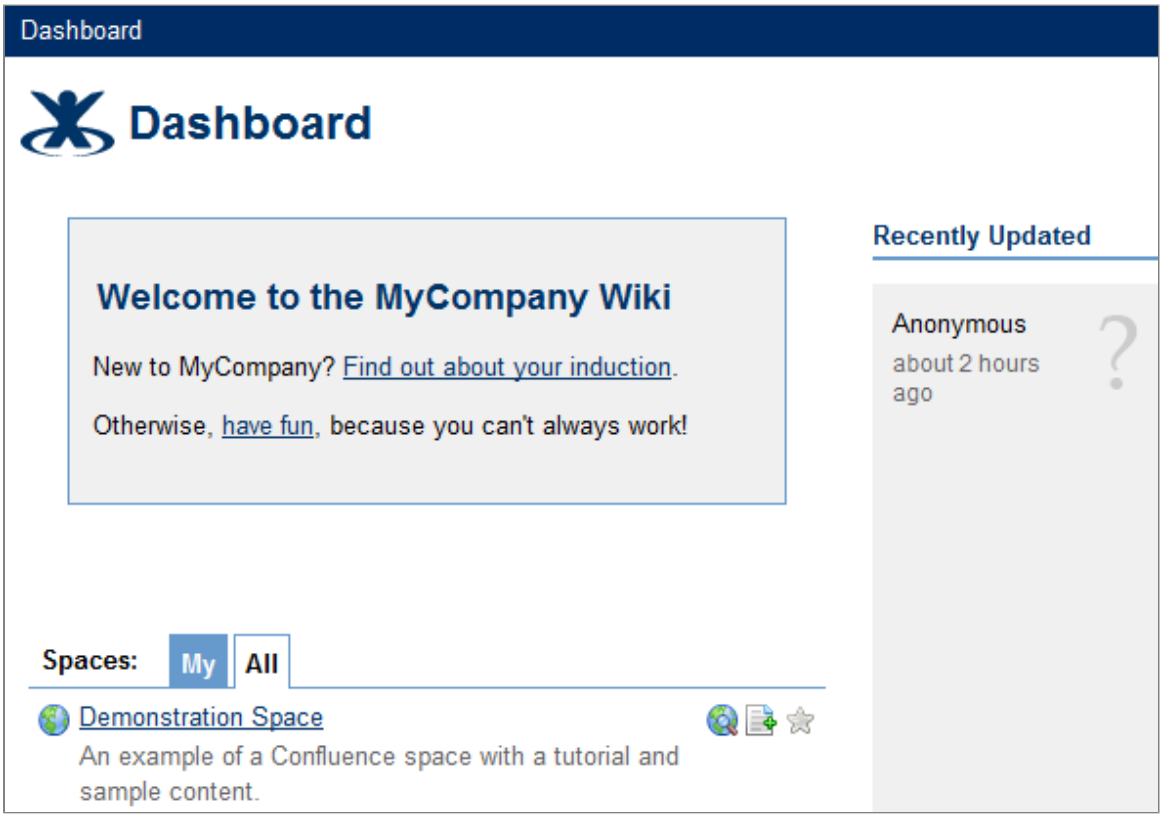
New to MyCompany? [Find out about your induction|DS:Company Induction].

Otherwise, [have fun|DS:Have Fun], because you can't always work!
```

In our example, the links point to two pages in the Confluence Demonstration Space, 'DS'. If your Confluence site does not have a 'DS' space, the links will be broken. That's OK, because you will want to replace them with links to your own pages anyway. This is just an example.

Example 2. Formatting your Welcome Message

Now let's say you want to put the words into a panel and add some spacing, so that your dashboard looks like this:



To produce the above welcome message, follow the step-by-step instructions above and add the following wiki markup into the '**Site Welcome Message**' text box:

```
{panel}
h2. Welcome to the MyCompany Wiki

New to MyCompany? [Find out about your induction|DS:Company Induction].

Otherwise, [have fun|DS:Have Fun], because you can't always work!
\\
\\
{panel}
\\
\\
```

Example 3: Including Content from Another Page

It may be easier to write your welcome message on a normal Confluence page and include the page into the '**Site Welcome Message**' text box. Using a normal page means that you can:

- Write the message using the Rich Text Editor rather than wiki markup.
- Preview the content of the welcome message before saving it, using the page editor's preview feature.
- Allow other people, who are not Confluence administrators, to edit the welcome message.

To include content from another page:

1. Create a Confluence page as usual and add your welcome message as the page content. Remember to limit the size of the content, because it must fit nicely onto the dashboard. For this example, let's assume you put your page in the '**DS**' space and the title of your page is '**Dashboard Welcome Message**'.
2. Add page permissions or space permissions to suit your requirements. You may want to restrict the editing of the page to a group of people, or you may want to allow any employee to edit the page. This will determine who can update the welcome message on the dashboard.
3. Follow the step-by-step instructions above and add the following wiki markup into the '**Site Welcome Message**' text box:

   ```
   {include:DS:Dashboard Welcome Message}
   ```

4. Save the site welcome message. The dashboard will display the content of the page immediately. Similarly, if you or anyone else edits the page, the welcome message on the dashboard will change as soon as you save the page.

In the above example we use the {include} macro to display the content from the given page. See the guide to the include macro. In our

example, the space key 'DS' and the page name 'Dashboard Welcome Message' are variable. You can use any space and page you like.

Example 4. Adding Blog Posts Filtered by Labels to your Welcome Message

Looking for more advanced ideas? This video shows you how to display a list of blog posts on your dashboard and how to choose the blog posts by labelling them.

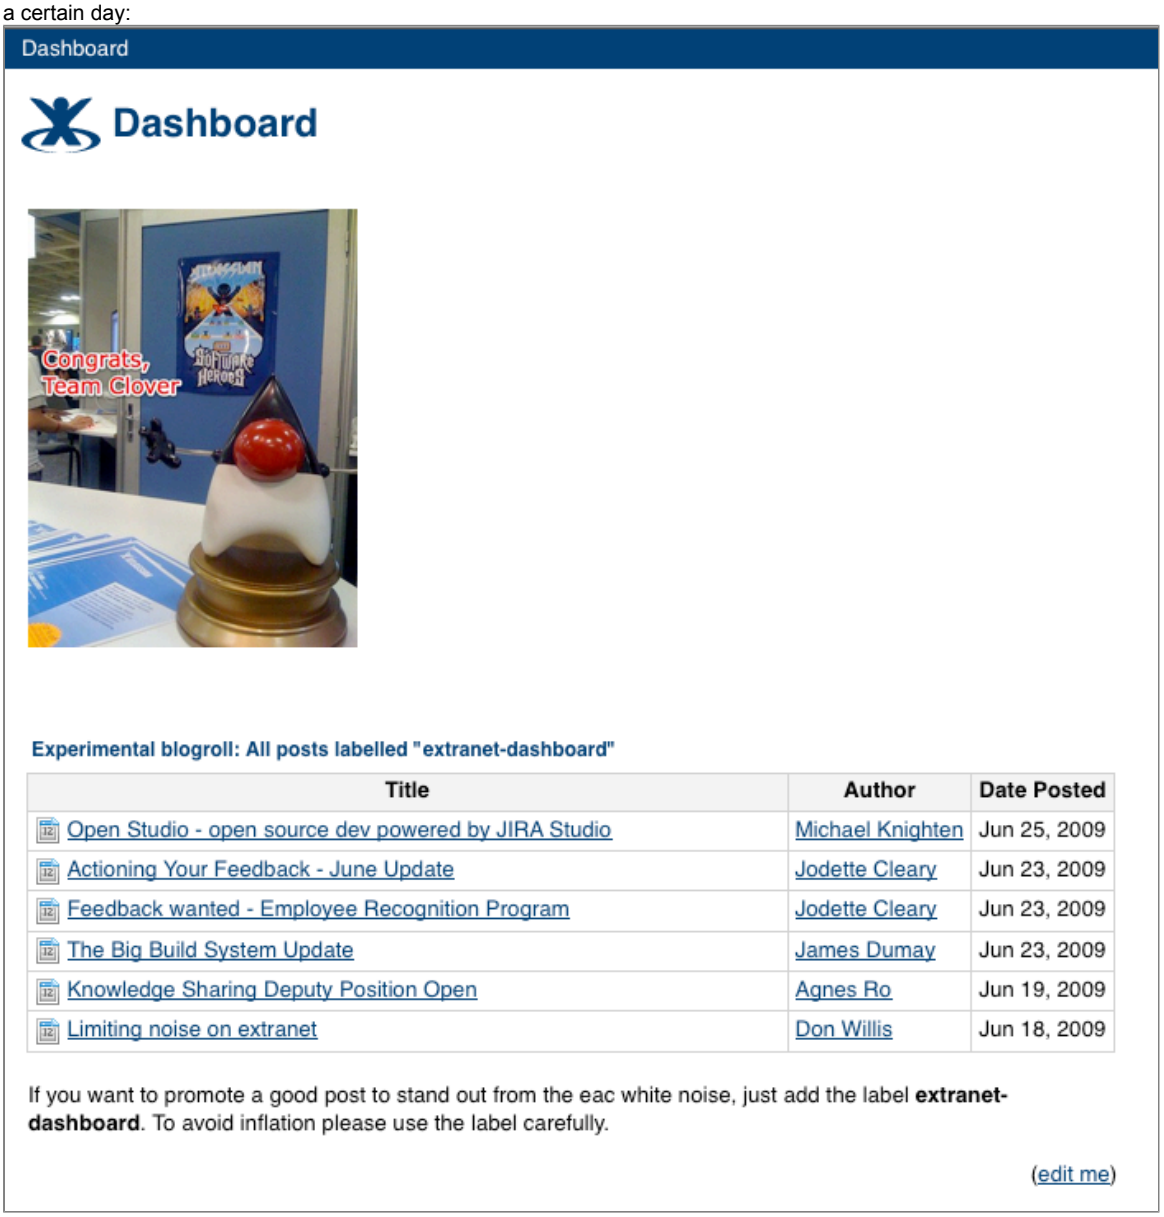
Video title: '*Bring "Must Read" Content to the Dashboard*'

Summary of the procedure shown in the video:

1. Create a page containing the {`blog-posts`} macro. Choose to display only the blog posts that are labelled with '**dashboard-blog**'. (This is just an example of a label. You can choose any label text you like.) See the guide to the Blog Posts macro.
2. Add the label to a blog post. (In the video, we just add the label to one blog post. You will probably want to add it to a number of posts.)
3. Edit your site welcome message to include the above page, using the include macro.

How we Use the Site Welcome Message at Atlassian

Atlassian makes great use of the welcome message on our internal Confluence wiki. Here is an example of the dashboard as it appeared on a certain day:



The welcome message itself contains just an {`include`} macro:

```
{include:STAFF:Extranet Homepage}
```

The include macro allows you to include the content an entire page onto another page. This particular page lives in the `STAFF` space, where anyone can edit it. It usually shows some amusing picture or company-wide notice. The featured photo generally changes once a week or so

– whenever someone feels like changing it. The page itself has over 600 edits by many different people.

The page also includes an edit link, for quick access to change the welcome message. We have the Composition plugin installed which allows you to use the {float} macro.

Our wiki markup in the 'Extranet Homepage' page looks something like this:

```
!Clover Dukey.jpg|width=200!

{nodisplay}
This is the content that goes on the Extranet homepage, above the spaces list.

NOTE: KEEP YOUR PICTURES SMALL (<80KB) -- USE JPG FOR PICTURES, WIDTH 400
{nodisplay}
h4. Experimental blogroll: All posts labelled "extranet-dashboard"

{blog-posts:content=titles|labels=extranet-dashboard|spaces=@all|max=10}
If you want to promote a good post to stand out from the eac white noise,
just add the label *extranet-dashboard*. To avoid inflation please use the
label carefully.

{float-right}
([edit me|http://extranet.atlassian.com/pages/editpage.action?pageId=603422736])
{float-right}
```

RELATED TOPICS

- Configuring the Site Home Page
- Configuring the Destination of View Space Links
- Configuring the Server Base URL
- Configuring the Site Support Address
- Customising Default Space Content
- Editing the Global Logo
- Editing the Site Title
- Editing the Site Welcome Message
- Showing Link Icons

Administrators Guide Home Confluence Documentation Home

Configuring the Server Base URL

The **Server Base URL** is the URL via which users access Confluence. The base URL **must** be set to the same URL by which browsers will be viewing your Confluence site.

Confluence will automatically detect the base URL during setup, but you may need to set it manually if your site's URL changes or if you set up Confluence from a different URL to the one that will be used to access it publicly.

You need to have System Administrator permissions in order to perform this function.

To configure the Server Base URL,

1. In Confluence, open the '**Browse**' menu and select '**Confluence Admin**'. The 'Administration Console' will open.
2. Click '**General Configuration**' in the left-hand panel.
3. Click the '**Edit**' button next to '**Site Configuration**'.
4. Enter the new URL in the '**Server Base URL**' text box.
5. '**Save**' your changes.

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.

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`, the server base URL should be `http://www.foobar.com/confluence`.

RELATED TOPICS

Configuring the Site Home Page

Configuring the Destination of View Space Links

Configuring the Server Base URL

Configuring the Site Support Address

Customising Default Space Content

Editing the Global Logo

Editing the Site Title

Editing the Site Welcome Message

Showing Link Icons

Configuring the Site Support Address

The **Site Support Address** is an email address which points to a JIRA instance configured to receive and handle support requests by email.

By default, the site support address is set to the Atlassian Support System. In most cases, there is no need to change it.

In order to use the site support address, ensure that SMTP email is set up on your Confluence instance.

To configure the site support address,

1. Go to the **Administration Console** and click '**General Configuration**' in the left-hand panel.
2. Click the '**Edit**' button next to 'Site Configuration'.
3. Enter the new '**Site Support Address**'.
4. Click the '**Save**' button at the bottom of the screen.

RELATED TOPICS

Troubleshooting Problems & Requesting Technical Support
Site Configuration
Administrators Guide Home Confluence Documentation Home

Configuring the Destination of View Space Links

By default, when you click a space link in order to view the space, you are taken to the space's home page. If you wish, you can configure Confluence to redirect all space links on the site to the '**Browse Space**' view of the space instead.

To direct the space link to the 'browse space' view,

1. Go to the '**Administration Console**' and click '**General Configuration**' in the left-hand panel.
2. Click '**Edit**' at the top of the '**Site Configuration**' screen.
3. Select '**On**' next to '**View Space goes to Browse Space**'.
4. Click '**Save**'.

RELATED TOPICS

Configuring the Site Home Page

Configuring the Destination of View Space Links

Configuring the Server Base URL

Configuring the Site Support Address

Customising Default Space Content
Editing the Global Logo
Editing the Site Title
Editing the Site Welcome Message
Showing Link Icons

Administrators Guide Home Confluence Documentation Home

Customising Default Space Content

Confluence Administrators can define **default content for a space home page**. This content will appear on the home page whenever someone adds a new space. You can define different content for global spaces and for personal spaces.

The default content will appear only for new spaces created after you have defined the content. Content in existing home pages will not be changed.

To define default content for home pages in global spaces,

1. Go to the '**Administration Console**' and click '**Default Space Content**' under 'Configuration' in the left-hand panel.
2. The '**Space Home Pages**' tab will open on the 'Default Space Content' page. Enter the content which you want to appear on the home page for new global spaces. You can use special characters within the content as variables (place holders). Confluence will replace the curly brackets and digits with the corresponding information as shown below:
    - **{0}** — The space name.
3. Click the '**Save**' button.

To define default content for home pages in personal spaces,

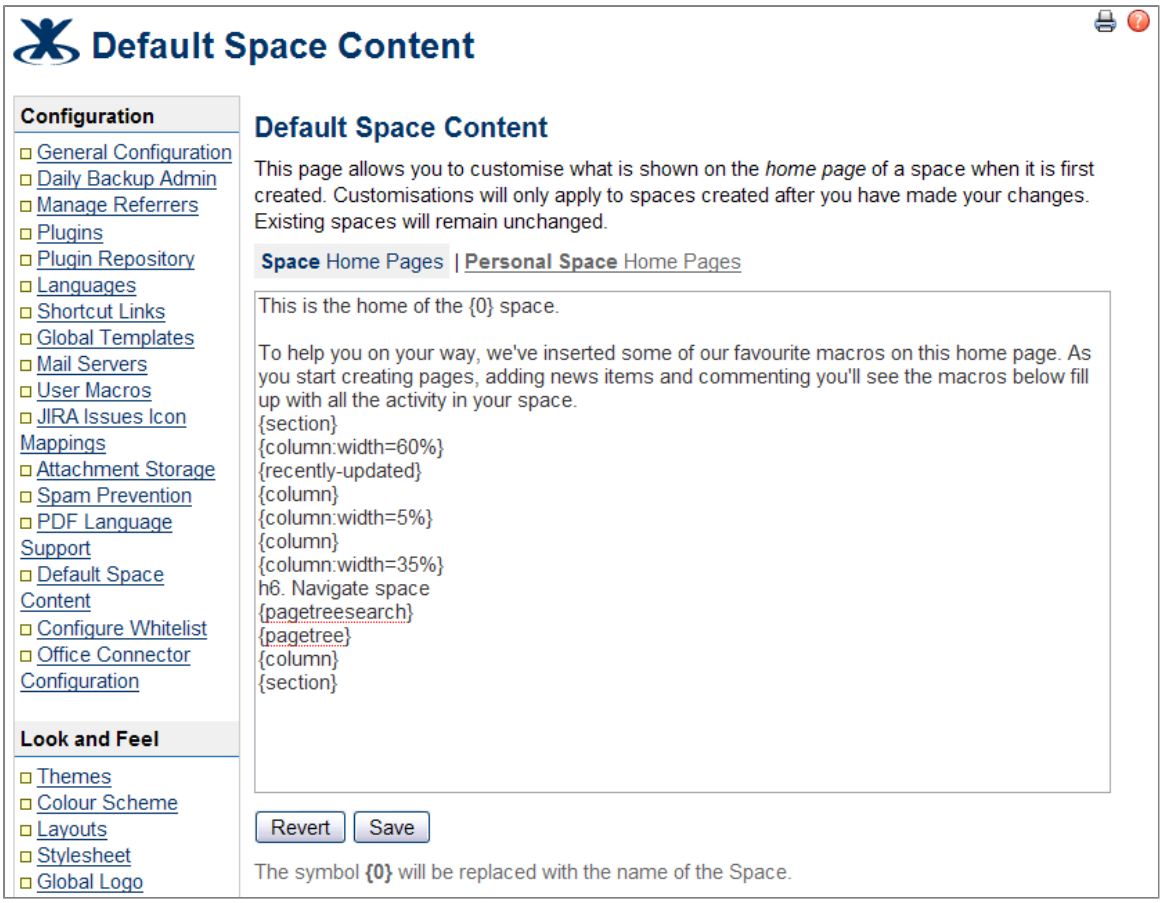
1. Go to the '**Administration Console**' and click '**Default Space Content**' under 'Configuration' in the left panel.
2. The 'Space Home Pages' tab will open on the 'Default Space Content' page. Click the '**Personal Space Home Pages**' tab.
3. Enter the content which you want to appear on the home page for new personal spaces. You can use special characters within the content as variables (place holders). Confluence will replace the curly brackets and digits with the corresponding information as shown below:
    - **{0}** — The space owner's full name.
    - **{1}** — The space owner's e-mail address.
    - **{2}** — Any personal information the space owner has entered on their user profile in the 'Information about me' section.
4. Click the '**Save**' button.

You can also undo all customisations of the default home page content, and go back to the default content as originally supplied with Confluence.

To restore the original default content,

1. Go to the '**Administration Console**' and click '**Default Space Content**' under 'Configuration' in the left panel.
2. Select either the 'Space Home Pages' tab or the '**Personal Space Home Pages**' tab, as required.
3. Click the '**Revert**' button.

Screenshot : Defining default space content



RELATED TOPICS

- Configuring the Site Home Page
- Configuring the Destination of View Space Links
- Configuring the Server Base URL
- Configuring the Site Support Address
- Customising Default Space Content
- Editing the Global Logo
- Editing the Site Title
- Editing the Site Welcome Message
- Showing Link Icons

Administrators Guide Home Confluence Documentation Home

Showing Link Icons

In order to distinguish external links, user links and email links in wiki content, the Confluence Administrator can configure Confluence to show a small icon in the top right-hand corner of each link.

http://www.google.com

Company Documentation

sales@atlassian.com

To show link icons,

1. From the '**Administration Console**', click '**General Configuration**' under the heading '**Configuration**' in the left-hand panel.
2. Click the '**Edit**' button next to '**Formatting and International Settings**'.
3. Beside '**Show Link Icons**', select '**On**' to enable the feature. Select '**Off**' to disable it.
4. Click '**Save**'.

RELATED TOPICS

- Configuring the Site Home Page
- Configuring the Destination of View Space Links
- Configuring the Server Base URL
- Configuring the Site Support Address
- Customising Default Space Content
- Editing the Global Logo
- Editing the Site Title
- Editing the Site Welcome Message
- Showing Link Icons

Administrators Guide Home Confluence Documentation Home

Configuring Encoding

Confluence allows the configuration of which character encoding is used to deliver pages.

While different character encodings are supported, we strongly recommend that **UTF-8** is used. Confluence is heavily tested on UTF-8, and users are likely to have less problems with this encoding than others.

Mac Users
Mac Users please note that **MacRoman** encoding is compatible with UTF-8. You do not need to change your encoding settings if you are already using MacRoman.

To avoid problems with character encoding, make sure the encoding used across the different components of your system are the same:

- Configuring Database Character Encoding
- Application Server URL encoding
- Confluence Character Encoding

If you are having problems with the character encoding in Confluence, please see the Troubleshooting Character Encodings page.

Character encodings in Confluence

Where character encoding is used

There are three places that character encoding matters to Confluence:

1. **Database encoding** - usually the most important; it is where almost all user data is stored.
2. **Filesystem encoding** - important for attachment storage (pre-2.2), reading Velocity templates and writing exported files.
3. **HTTP request and response encoding** - important for form parsing, correct rendering by the browser and browser interpretation of encoded URLs.

Problems generally arise when Confluence thinks one of the above encoding is different to what it actually is. For example, Confluence might believe the database is using ISO-8859-1 encoding, when in fact it is UTF-8 encoded.

Java character encoding

Java *always* uses the double-byte UCS-2 character encoding for all `char` and `String` data. This means that each of the encodings above defines how, at that particular point, characters are converted to and from Java's native UCS-2 format into some other format that the browser, filesystem or database might understand.

So when a request comes in to Confluence, we convert it from the request encoding to UCS-2. Then we store that data into the database, converting from UCS-2 to the database's encoding. Retrieving information from the database and sending it back to the browser is the same process in the opposite direction.

Problems with character encodings

If Confluence has the wrong idea about encoding for one of the above, it manifests itself in different ways:

1. Incorrect database encoding - user data is corrupted between saving and restoring from the database. This often happens after a delay, as we cache data as it is written to the database and only later retrieve the corrupted copy from the database.
2. Incorrect/non-Unicode filesystem encoding - international filenames break attachment download/upload/removal (pre-2.2); exports break with international content or attachments.
3. Incorrect HTTP encoding - incorrect encoding selected by browser, resulting in incorrect rendering of characters. Changing browser encoding causes page to render properly. Broken URLs when linking to pages or attachments with non-ASCII characters.

Configuration of character encodings

The **Confluence character encoding** is a configuration setting found in `Administration > General Configuration`, and at runtime available in Settings.defaultEncoding. It is subsequently used in the following parts of the system:

- ConfluenceWebWorkConfiguration sets `webwork.i18n.encoding` to the this encoding, which WebWork uses in the response Content-Type header.
- AbstractEncodingFilter sets the HTTP request encoding to this encoding. This seems unnecessary, since the Content-Type header from the client should include the encoding used. This affects form submissions and file uploads.
- VelocityUtils reads in Velocity templates using this encoding when reading templates from disk.
- AbstractXmlExporter creates its output using this encoding.
- GeneralUtil uses this encoding when doing URLEncode and URLDecode. Different browsers have different support for character sets in URLs, so it's uncertain how much benefit this provides.

In summary, changing the Confluence character encoding will change your **HTTP request and response encoding** and your **Filesystem encoding** as used by exports and velocity templates.

The **database encoding** is the responsibility of your JDBC drivers. The drivers are responsible for reading and writing from the database in its native encoding and translating this data to and from Java Strings (which are UCS-2). For some drivers, such as MySQL, you must set Unicode encoding explicitly in the JDBC URL. For others, the driver is smart enough to determine the database encoding automatically.

Ideally, your database itself should be in a Unicode encoding (and we recommend doing this for the simplest configuration), but that is not necessary as long as:

- the database encoding supports all the characters you want to store in Confluence
- your JDBC drivers can properly convert from the database encoding to UCS-2 and vice-versa.

The **filesystem encoding** is mostly ignored by Confluence, except for the cases where the above configuration setting above plays a part (exports, velocity). When attachments are uploaded, they are written as a stream of bytes directly to the filesystem. It is the same when they are downloaded: the bytes from the file InputStream are written directly to the HTTP response.

In some places in Confluence, we use the *default filesystem encoding* as determined by the JVM and stored in the `file.encoding` system property (it can be overridden by setting this property at startup). This encoding is used by the Java InputStreamReader and InputStreamWriter classes by default. This encoding should probably never be used; for consistent results across all filesystem access we should be using the encoding set in the General Configuration.

In certain cases we explicitly hard-code the encoding used to read or write data to the filesystem. Two important examples are:

- importing Mbox mailboxes which are known to be ISO-8859-1
- Confluence Bandana config files are always stored as UTF-8.

Some application servers, Tomcat for example, have an encoding setting that modifies Confluence URLs before they reach the application. This can prevent access to international pages and attachments (really anything with international characters in the URL). See configuring your Application Server URL encoding.

Advice

In general, **always set all character encodings to UTF-8**. That includes database, JDBC drivers, application server, filesystem and Confluence.

In certain isolated cases (e.g. Microsoft Windows), it might not be possible to use a fully Unicode filesystem (that is, a default Windows install doesn't support Unicode filenames properly). If so, stick with UTF-8 for the other two and be aware that your operating system might have limitations around international attachments (pre-2.2), backup and restore of international data, etc.

RELATED TOPICS:

- Configuring Database Character Encoding
- Troubleshooting Character Encodings

Troubleshooting Character Encodings

Often users may have problems with certain characters in a Confluence instance. Symptoms may include:

- Non-ASCII characters appearing as question marks (?)
- Page links with non-ASCII characters not working
- Single characters being displayed as two characters
- Garbled text appearing

In most cases, it is due to a mis-configuration in one of the components that Confluence uses.

Follow these steps to diagnose the problem:

1. Run the encoding test

Confluence includes an encoding test that can reveal problems with your configuration.

To perform the test, access the Encoding Test page via the `<confluence base-url>/admin/encodingtest.action` page on your Confluence instance. You will be required to copy and paste a line of text and submit a form. The test will take the text and pass it through Confluence, the application server and the database, and return the results.

You should also test pasting some sample text (Japanese for example) if you are experiencing problems with a specific language.

Example:

```
http://confluence.atlassian.com/admin/encodingtest.action
```

or

```
http://<host address>:<port>/admin/encodingtest.action
```



If the text displayed in the encoding test is different to what was entered, then there are problems with your character encoding settings.

A successful test looks like the following:


The encoding test has now been run. Below, you can compare the raw text delivered from Confl
a round-trip through the database. All the test results should appear identical.

| Iñtërnâtiônàlizætiøn | This image is how all of the test results below *should* appear.
this page, and all of your System Information. |
|---|---|

Test 1: Raw text

This is the test string generated in Confluence

Iñtërnâtiônàlizætiøn

Test 2: Form submission

This is the test string pasted by you into the web form and submitted back to Confluence

Iñtërnâtiônàlizætiøn

Test 3: Database round-trip (select as LOWER)

This is the string from Test 2 after being stored in the database and then retrieved

iñtërnâtiônàlizætiøn

Expected result (converting Java string to lowercase)

iñtërnâtiônàlizætiøn

Test 4: Database round-trip (select as UPPER)

This is the string from Test 2 after being stored in the database and then retrieved

IÑTËRNÂTIÔNÀLIZÆTIØN

Expected result (converting Java string to uppercase)

IÑTËRNÂTIÔNÀLIZÆTIØN

MySQL 3.x
MySQL 3.x is known to have some problems with the upper- and lower-casing of some characters, and may fail the last two tests. For more information, see MySQL 3.x Character Encoding Problems.

2. Ensure the same encoding is used across all components

As mentioned in the Configuring Encoding document, the same character encoding should be used across the database, application server and web application (Confluence).

- To change the character encoding used in **Confluence**, see Configuring Character Encoding.
- To change the character encoding used in the **application server**, please ensure you set the Application Server URL encoding and view your application server's documentation on any other settings required to enable your encoding.
- To change the character encoding used in the **database**, see Configuring Database Character Encoding.

3. Requesting support

If there are still problems with character encoding after following the above steps, create a support request, and our support staff will aid in solving your problem.

Entering in the following details will help us to identify your problem:

- Attach screenshots of the problem
- Attach the results of the encoding test (above)
- Select which application server (and version) you are using
- Select which database (and version) you are using
- Copy the contents of the System Information page into the 'Description' field

"€" Euro character

The € (euro) symbol is a three byte character, with byte values in file (UTF-8) of 0xE2, 0x82, 0xAC.

Sometimes, if the character encoding is not set consistently among all participating entities of the system, Confluence, server and the database, one may experience strange behaviour.

> ...
> I write a page with a Euro sign in it (€). All is well, the Euro sign shows up in the wiki markup text-box, and the preview, and the display of the saved page.
> One day later, the Euro sign has changed into a question mark upside down!
> ...
> What is going on? Why does the Euro sign mysteriously change? How do I prevent it?

Interestingly enough the character encoding test passes with no problems, demonstrating that Confluence and the connected Database both recognise the € symbol.

There are two potential reasons for this behaviour:

Database and Confluence is using utf-8 encoding. The connection is not.

When data transferred to it via the connection which does not use utf-8 encoding gets encoded incorrectly. Hence, updating the connection encoding may resolve this problem from now on, yet it probably would not affect already existing data.

Database is not using utf-8. Confluence and your connection are.

If your Database encoding is not set to UTF-8, yet is using some other encoding such as *latin1*, it could be one of the potential reasons why you lose the "€" characters at some stage. It could be occurring due to **caching**. When Confluence saves data to the database, it may also keep a local cached copy. If the database encoding is set incorrectly, the Euro character may not be correctly recorded in the database, but Confluence will continue to use its cached copy of that data (which is encoded correctly). The encoding error will only be noticed when the cache expires, and the incorrectly encoded data is fetched from the database.

> For instance the *latin1* encoding would store and display all 2-byte UTF8 characters correctly except for the euro character which is replaced by '?' before being stored. As Confluence's encoding was set to UTF-8, the 2-byte UTF-8 characters were stored in *latin1* database assuming that they were two *latin1* different characters, instead of one utf8 character. Nevertheless, this is not the case for 3-byte utf8 characters, such as the Euro symbol.

Please ensure that you set the character encoding to UTF-8 for all the entities of your system as advised in this guide.

MySQL 3.x Character Encoding Problems

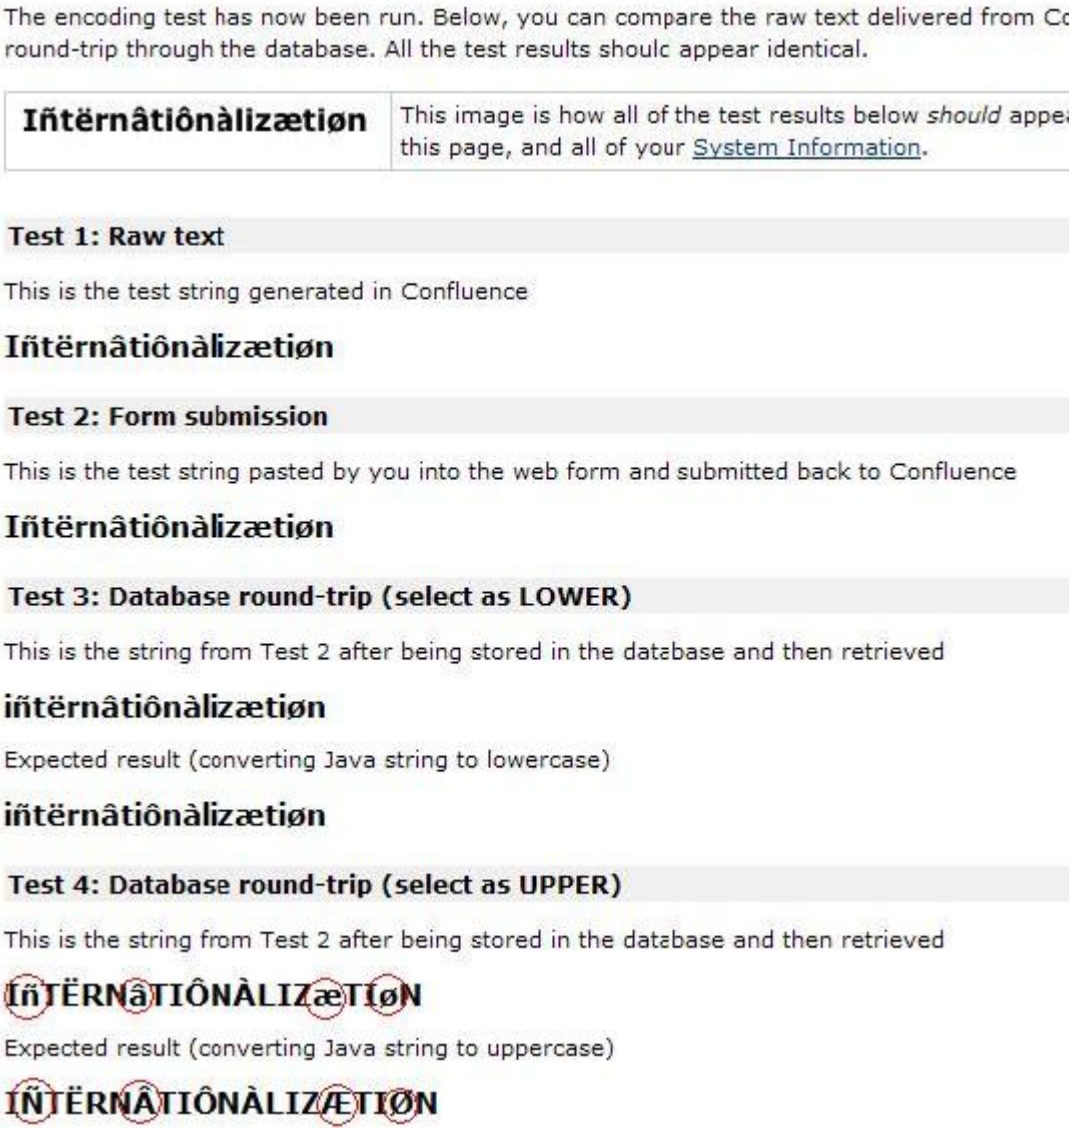
MySQL 3.x is known to have some problems upper- and lower-casing certain (non-ASCII) characters.

Diagnosing the problem

1. Follow the instructions for Troubleshooting Character Encodings.
2. If the upper- and lower-cased strings displayed on the Encoding Test are different, then your database is probably affected.

An example (faulty) output of the Encoding Test is shown below:

Screenshot: Encoding Test Output



Solution

Upgrade to a newer version of MySQL. (4.1 is confirmed to work.)

Configuring Mail

- Configuring a Server for Outgoing Mail
- Enabling the 'Mail Page' plugin
- The Mail Queue

- Customising the eMail Templates

Administrators Guide Home Confluence Documentation Home

Configuring a Server for Outgoing Mail

Configuring your Confluence server to send outgoing mail allows your Confluence users to:

- Receive Daily Reports.
- Send a page via email.

You need to have System Administrator permissions in order to perform this function.

To configure Confluence Standalone to send outgoing mail,

1. Go to the Confluence '**Administration Console**'. To do this:
    - Open the '**Browse**' menu and select '**Confluence Admin**'. The 'Administration Console' view will open.
2. Select '**Mail Servers**' under '**Configuration**' in the left panel. This will list all currently configured SMTP servers.
3. Click '**Add New SMTP Server**' (or **edit** an existing server).
    - **Name**: By default, this is set to 'SMTP Server'.
    - **From**: Enter the email address that will be displayed in the 'from' field for email messages originating from this server.
    - **Subject Prefix**: Enter a subject prefix, if required.
4. Configuring the Host Address, Username and Password:
    - Manually enter your '**Host Address**', '**Username**' and '**Password**' details in the form fields displayed (recommended).
      OR
    - Specify the '**JNDI**' location of a mail session configured in your application server in the form field displayed.

Troubleshooting

If you experience problems with these configurations, please check that your `<Confluence-Install>/confluence/WEB-INF/lib` contains only one copy of the following JAR files:

1. activation-x.x.x.jar
2. mail-x.x.x.jar

Ideally, these should be:

- activation-1.0.2.jar
- mail-1.3.2.jar (or later)

You will then need to move these into the proper directory:

Standalone distribution: Please move the above three jar files from the `<Confluence-Install>/confluence/WEB-INF/lib` directory to `<confluence-install>/lib` (for Confluence version 2.10 onwards) or `<Confluence-Install>/common/lib` (for earlier product versions) and restart Confluence.

RELATED TOPICS

Enabling the 'Mail Page' plugin

Configuring a Server for Outgoing Mail

The Mail Queue

Enabling the 'Mail Page' plugin

The 'Mail Page' plugin allows anyone with the 'View' space permission to email a Confluence page.

The 'Mail Page' plugin is disabled by default. This is because, when someone emails a Confluence page, they can select from a list of all Confluence users and groups (note, however, that email addresses are not visible), or even mail the page to arbitrary addresses. If you have enabled anonymous access or self-signup, visitors could potentially use this feature to send spam or nuisance email through your Confluence server.


This plugin only works when the **mail server** is configured.


You need to have System Administrator permissions in order to perform this function.


Confluence versions 2.4 and later come with the 'Mail Page' plugin preinstalled.

To enable the 'Mail Page' plugin,

1. Go to the Confluence '**Administration Console**'. To do this:
    - Open the '**Browse**' menu and select '**Confluence Admin**'. The 'Administration Console' view will open.
2. Select '**Plugins**' under '**Configuration**' in the left-hand panel.
3. This will list all plugins that are currently installed in your Confluence system. Click '**Mail Page Plugin**'.
4. This will display the '**Mail Page Plugin**' details. To enable the 'Mail Page' plugin, click '**Enable plugin**'.
5. Ensure that both of the following are enabled:
    - '**Mail Page Link**' — displays the 'E-mail' link next to the 'Copy' link on the 'Page Info' screen (see 'E-mailing a page')
    - '**mailpageactions**' — enables the e-mail operation.

RELATED TOPICS

Enabling the 'Mail Page' plugin

Configuring a Server for Outgoing Mail

The Mail Queue

The Mail Queue

Email messages waiting to be sent out are queued in a mail queue and periodically flushed from Confluence once a minute. A Confluence administrator can also manually flush emails from the mail queue.

If there is an error sending messages, the failed emails are sent to an error queue from which you can either try to resend them or delete them.

To view the mail queue,

1. Go to the Confluence '**Administration Console**'. To do this:
    - Open the '**Browse**' menu and select '**Confluence Admin**'. The 'Administration Console' view will open.
2. Select '**Mail Queue**' in the left-hand panel. This will display the emails currently in the queue.
3. Click '**Flush Mail Queue**' to send all emails immediately.
4. Click '**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 '**Delete**' them from here.

RELATED TOPICS

Enabling the 'Mail Page' plugin

Configuring a Server for Outgoing Mail

The Mail Queue

RELATED TOPICS

Administrators Guide Home Confluence Documentation Home

Optional Settings

- Attachment Storage Configuration
- Configuring Quick Navigation
- Enabling CamelCase Linking
- Enabling OpenSearch
- Enabling Remote APIs
- Enabling Rich Text Editing Option
- Enabling the Did You Mean Feature
- Enabling Threaded Comments
- Enabling Trackback
- Making Rich Text Editing default
- WebDAV Configuration

Attachment Storage Configuration

Confluence allows you to store attachments in one of three places:

- Filesystem - locally in the Confluence home directory
- Database - in Confluence's configured database
- WebDAV - remotely on a WebDAV server (***deprecated***)

A System Administrator can configure Confluence's attachment storage via the '**Attachment Storage**' option on the '**Administration Console**'.

You need to have System Administrator permissions in order to perform this function.

Attachment Storage Options

Local File System

By default, Confluence stores attachments in the `attachments` directory within the configured Confluence home folder. If you are looking to run Confluence Clustered, attachments must be stored in the database.

Database

Confluence gives administrators the option to store attachments in the database that Confluence is configured to use.

Here are some reasons why, as an administrator, you may want to choose this storage system:

- Ease of backup.
- Avoiding issues with certain characters in attachment file names.

> While storing attachments in the database can offer some advantages, please be aware that the amount of space used by the database will increase because of the greater storage requirements.

WebDAV

Confluence also allows administrators to set an external WebDAV repository as the location for attachment storage.

> **WebDAV attachment manager deprecated**
> The option to store Confluence attachments on a WebDAV server has never worked in a useful fashion, and has not been maintained for many versions.
>
> - The WebDAV attachment manager will be **deprecated** from Confluence 2.7, and will be removed from a later version of Confluence.
> - If you store attachments on external WebDAV servers, we recommend that you migrate to file-system or database-backed attachment storage as soon as possible. Refer to CONF-9313 and CONF-2887.
> - This DOES NOT affect the operation of the WebDAV plugin.

Migration between Attachment Storage Systems

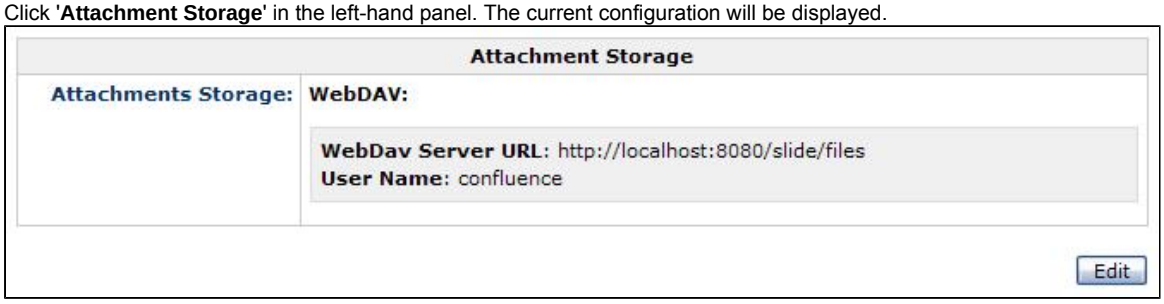
You can 'migrate' your attachments from one storage system to another. All existing attachments will be moved over to the new attachment storage system.

> When the migration occurs, all other users will be locked out of the Confluence instance. This is to prevent modification of attachments while the migration occurs. Access will be restored as soon as the migration is complete.

> When migrating attachments from your database to a filesystem, the attachments are removed from the database after migration. However, when migrating attachments from a filesystem to your database, the attachments remain on the filesystem after migration. If you wish to change this function's behaviour from 'copy' to 'move', please see CONF-14802 and cast your vote.

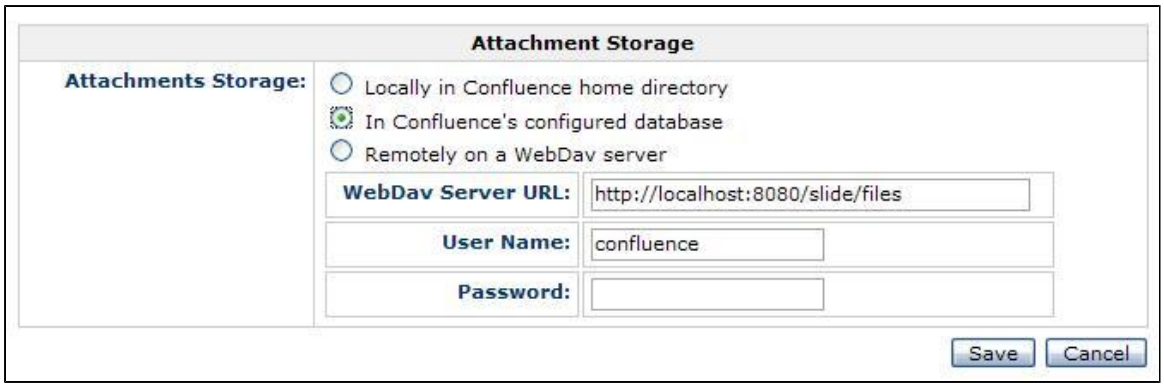
To perform a migration, follow the steps below:

1. Go to the Confluence '**Administration Console**'. To do this:
    - Open the '**Browse**' menu and select '**Confluence Admin**'. The 'Administration Console' view will open.
2. Click '**Attachment Storage**' in the left-hand panel. The current configuration will be displayed.



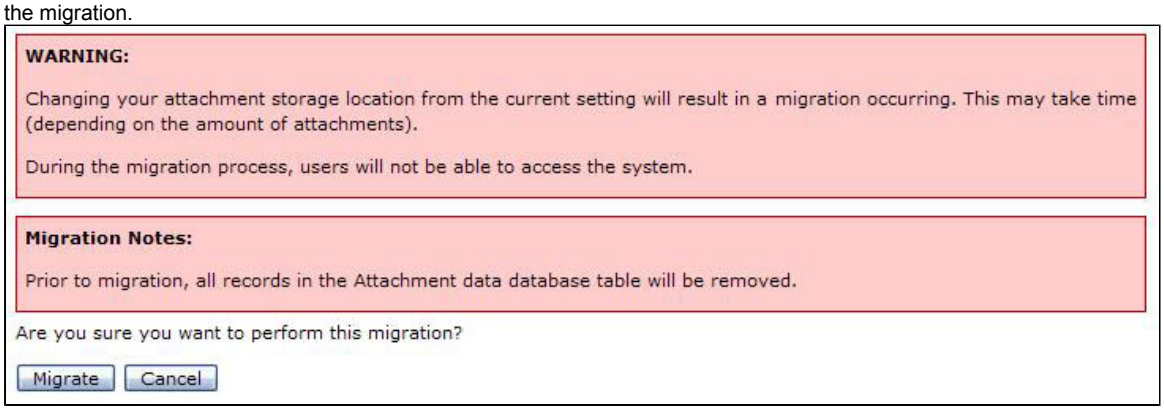
Attachment storage configuration

3. Click the '**Edit**' button to modify the configuration.
4. Select the storage system you desire.



Edit attachment storage

5. Click the '**Save**' button to save the changes.
6. A screen will appear, asking you to confirm your changes. Clicking 'Migrate' will take you to a screen that displays the progress of the migration.



Migration warning

Troubleshooting

To enable debug logging for WebDAV attachment storage, add the following to the bottom of `WEB-INF/classes/log4j.properties` and restart Confluence:

```
log4j.logger.com.atlassian.confluence.pages.persistence.dao=DEBUG,confluencelog
log4j.additivity.com.atlassian.confluence.pages.persistence.dao=false

log4j.logger.org.apache.webdav=DEBUG,confluencelog
log4j.additivity.org.apache.webdav=false
```

RELATED TOPICS

Important Directories and Files

Administrators Guide Home Confluence Documentation Home

Hierarchical File System Attachment Storage

Introduction

For Confluence version 3.0, the structure of attachments stored on the filesystem was changed. In versions of Confluence prior to 3.0, attachments were stored in directories corresponding to the id of the content to which they belong. The more content in Confluence with attachments, the more directories you would have immediately beneath your configured attachments directory. This directory structure has been changed in Confluence 3.0 and since the default configuration of Confluence is to store attachments in the filesystem, this change is likely to have relevance to administrators of most existing Confluence installations.

If you are installing Confluence for the first time, there will be no consequences as a result of this change. If you are upgrading from a previous version of Confluence, the migration to this new filesystem structure should happen automatically during the upgrade.

The reason for introducing this change was to address the issue CONF-13004. Certain file systems have a limit on the number of files that can be stored in a directory and large Confluence installations were reaching this limit. In addition, storing too many files at a single directory level can cause performance degradation in some circumstances. This new attachment storage strategy ensures this will no longer be the case.

Backup Confluence Home
Before upgrading to Confluence 3.0, as with any upgrade you must ensure you have a backup of your Confluence home directory before you proceed.

The New Directory Layout

The attachment storage layout was chosen to fulfil the following main requirements:

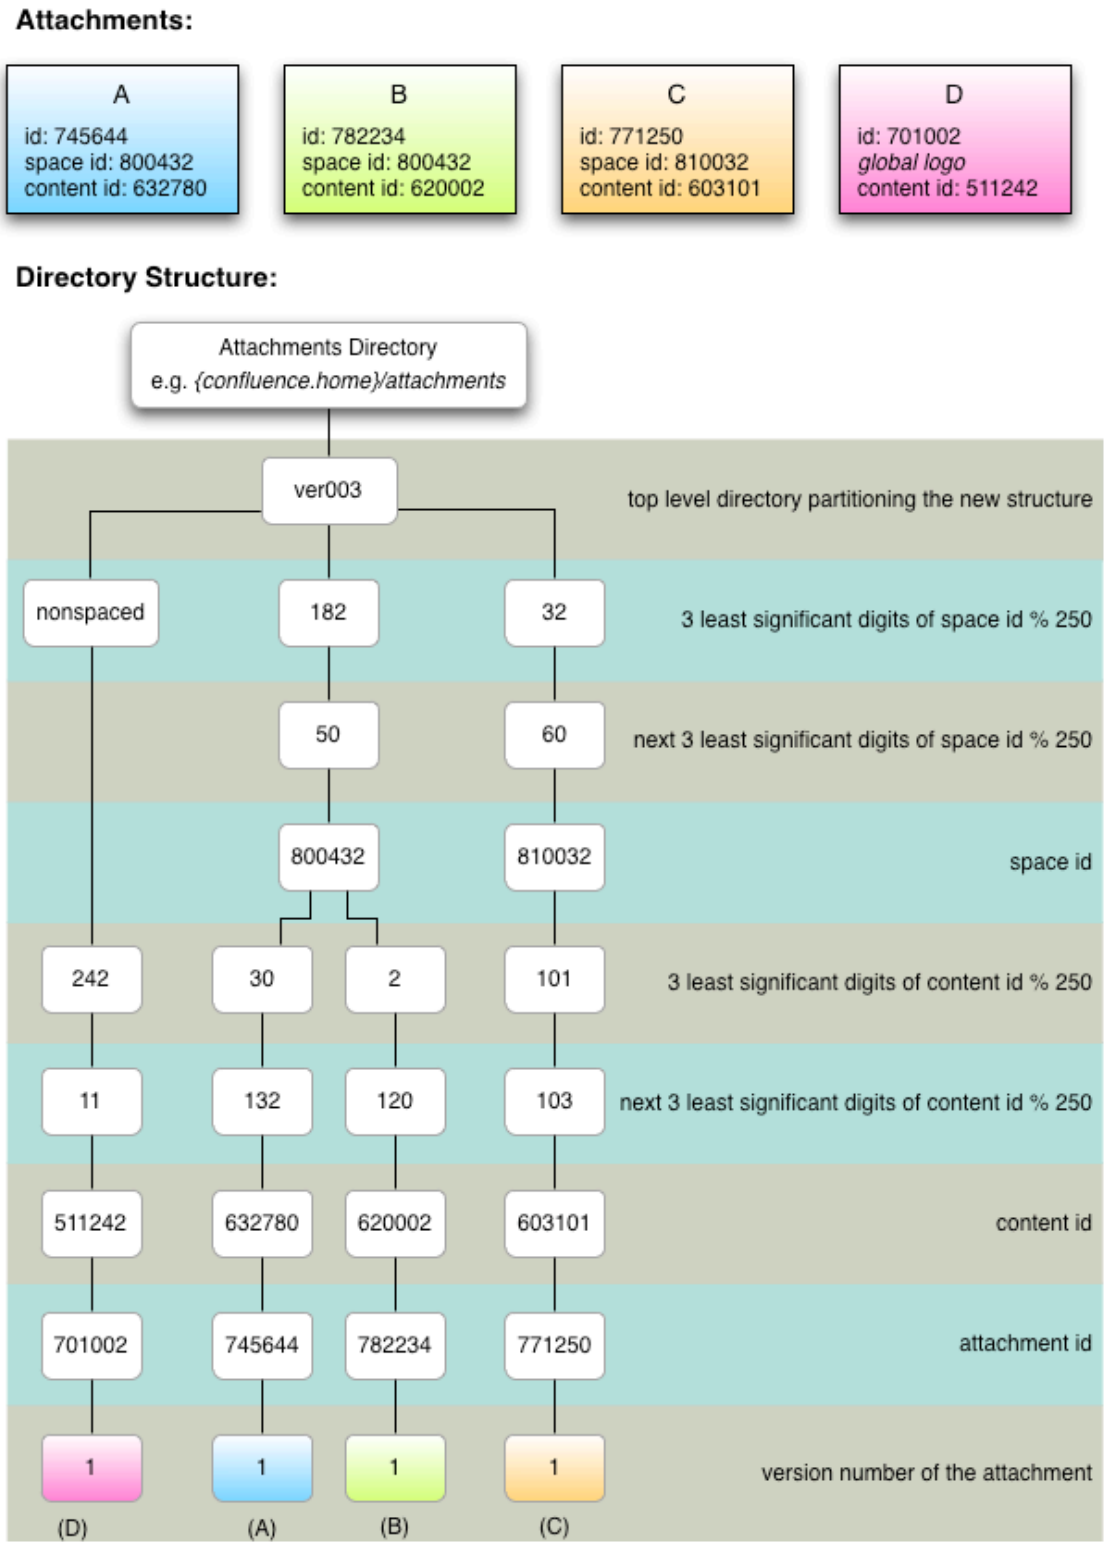
1. Limit the number of entries at any single level in a directory structure.
2. Partition attachments per space making it possible for a system admin to selectively back up attachments from particular spaces (see the JIRA issue for more details).

An attachment in Confluence can be thought of as having a number of identifying attributes: *id*, *space id* and *content id*. That is to say, the attachment logically belongs to a piece of content which logically belongs in a space (not all content belongs to a space). For attachments within a space in Confluence, the directory structure is typically 8 levels, with the name of each directory level based on the following algorithm:

level	Derived From
1 (top)	Always 'ver003' indicating the Confluence version 3 storage format
2	The least significant 3 digits of the *space id*, modulo 250
3	The next 3 least significant digits of the *space id*, modulo 250
4	The full *space id*
5	The least significant 3 digits of the *content id*, modulo 250
6	The next 3 least significant digits of the *content id*, modulo 250
7	The full *content id*
8	The full *attachment id*

Within the 8th level will be a file for each version of that attachment, named to match the version number e.g. 1

An example:



To find the directory where attachments for a particular space are stored, you can use the JSP findspaceattachments.jsp at the location `<confluence url>/admin/findspaceattachments.jsp`. This JSP requires a space key and returns the directory on the file system where attachments for that space are stored.

Attachment D in the above diagram is stored in a slightly different structure. Attachments that are not conceptually within a space replace the level 2 - 4 directories with a single directory called 'nonspaced'. Examples of such attachments are the global site logo and also attachments on draft content.

Upgrading to the new attachment storage structure

As mentioned previously, this upgrade is only necessary if you have Confluence configured to store attachments on the file system.

If migration is not necessary due to a different storage configuration (for example, because attachments are stored in the database), then no migration will occur during upgrade and the Confluence log will simply show the following messages -

```
INFO [main] [AbstractUpgradeManager] upgradeStarted Starting automatic upgrade of Confluence
INFO [main] [UpgradeTask] isUpgradeNeeded The configured attachmentDataDao does not store
    attachment data on the file system so the HierarchicalFileSystemAttachmentUpgradeTask is not
necessary.
INFO [main] [AbstractUpgradeManager] upgradeFinished Upgrade completed successfully
```

Should migration be required, it will occur automatically during upgrade and the log will show output similar to this -

```
INFO [main] [UpgradeTask] doUpgrade Beginning HierarchicalFileSystemAttachmentUpgradeTask.
Depending on the size of the
    attachment data this may take some time.
INFO [main] [UpgradeTask] run 4023 pages may have attachments to be moved to a new hierarchical
structure.
INFO [main] [UpgradeTask] run 0 of 4023 pages have had their attachments moved to the new
structure
INFO [main] [UpgradeTask] run 500 of 4023 pages have had their attachments moved to the new
structure
INFO [main] [UpgradeTask] run 1000 of 4023 pages have had their attachments moved to the new
structure
INFO [main] [UpgradeTask] run 1500 of 4023 pages have had their attachments moved to the new
structure
INFO [main] [UpgradeTask] run 2000 of 4023 pages have had their attachments moved to the new
structure
INFO [main] [UpgradeTask] run 2500 of 4023 pages have had their attachments moved to the new
structure
INFO [main] [UpgradeTask] run 3000 of 4023 pages have had their attachments moved to the new
structure
INFO [main] [UpgradeTask] run 3500 of 4023 pages have had their attachments moved to the new
structure
INFO [main] [UpgradeTask] run 4000 of 4023 pages have had their attachments moved to the new
structure
INFO [main] [UpgradeTask] run Successfully moved the attachments for all 4023 pages to the new
hierarchical structure.
INFO [main] [UpgradeTask] doUpgrade Completed HierarchicalFileSystemAttachmentUpgradeTask.
INFO [main] [AbstractUpgradeManager] upgradeFinished Upgrade completed successfully
```



It should be noted that for most implementations of Java, the migration to the new data structure involves moving the files (not copying them). Hence, there should not be a need to have additional disk space available. It also means that the migration should be relatively fast.

Have you previously applied the CONF-8298 patch?

The patch or workaround on the CONF-8298 issue changed the structure of attachment storage but not to the most efficient possible structure. So during the Confluence 3.0 upgrade process this intermediate (CONF-8298) structure will be detected and automatically upgraded.

Troubleshooting the upgrade



It should be noted that in the event of a failure, your attachment directory may be in an inconsistent state and your first step in troubleshooting should be to restore the backup of your home directory.

There are a number of reasons the migration could fail. This will be shown in the log with a message similar to "`Failed to move the attachments for all pages to the new hierarchical structure.`".

Immediately preceding this message in the log will be entries for each page whose attachments could not be moved. The following table shows examples of these messages and offers some possible explanations.

Example Message	Description
The configured attachment directory `<directory name>` could not be found or was not a directory.	The configured Confluence attachment directory is not accessible. Check confluence home for the attachment directory and ensure the permissions are correct to allow reading and writing for this directory.
It is not possible to migrate the attachments to the new structure since files already exist which the attachment process may need to create.	Your attachments directory contains files or directories which the upgrade task wants to create. That is, a top level directory called ver003 containing directories or files with names containing up to 3 digits (e.g. 1, 213). This could be due to a previous failed attempt to migrate the attachments. You should restore a previous good copy of your attachments directory and remove any files or directories with this naming pattern before retrying.

Couldn't find current Confluence content for the id `<content Id>`. The attachment is a non-spaced attachment (e.g. global logo, draft attachment, etc) and will be migrated to the nonspaced directory.	This is a normal message indicating that the attachment being migrated does not belong to a space e.g. global logo.
Problem while accessing the database for content id `content Id` so its attachments will not be migrated.	It was not possible to access the database at this point during the migration. You will need restore your Confluence attachment directory from the backup and attempt the upgrade again, once the database is accessible again.
Could not create the new attachment directory `directory`.	The upgrade task could not create the new directory to contain the attachment being moved. Does the server user have sufficient permission to perform this operation in the indicated directory? Is there sufficient disk space?
Failed to move the current attachment directory `<some path>` to the new location of `<some other path>`.	The upgrade task could not move the directory. Does the server user have sufficient permission to perform this operation in the indicated directory?

Configuring Quick Navigation

When a user is searching Confluence (see Using the Quick Navigation Aid) the quick navigation aid automatically offers a dropdown list of pages and other items, matched by title to the search query. By default, this feature is enabled, with the maximum number of simultaneous quick navigation requests set to 40. However, these options can be modified as described below.



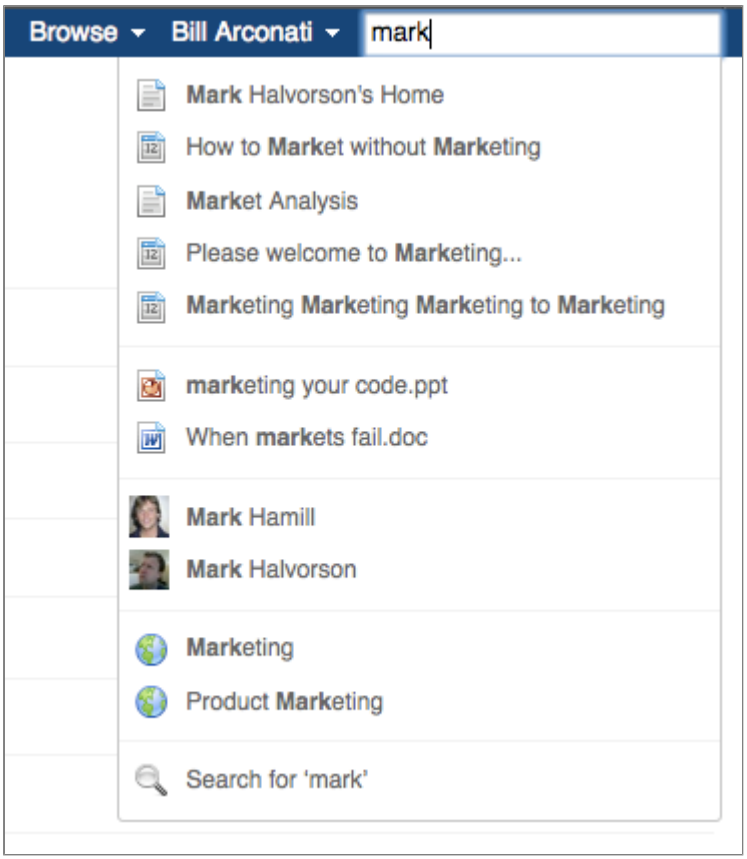
The maximum number of simultaneous quick navigation requests defines the maximum number of individuals who can use this feature simultaneously on the same Confluence server. If your Confluence server serves a large number of individuals who use this feature regularly, some of whom are being denied access to it, you may wish to increase this value.

To modify the quick navigation feature's options,

1. Go to the Confluence '**Administration Console**'. To do this:
    - Open the '**Browse**' menu and select '**Confluence Admin**'. The 'Administration Console' view will open.
2. Select '**General Configuration**' in the left-hand panel.
3. In the '**General Configuration**' screen, click '**Edit**'.
4. To disable this feature, select '**Off**' beside '**Quick Navigation**'.
5. To modify the maximum number of simultaneous quick navigation requests, enter the appropriate number in the field beside '**Max Simultaneous Requests**'.
6. Click '**Save**'.

The following screenshot demonstrates the user interface of the quick navigation aid.

Screenshot: The quick navigation aid showing titles matching the query 'mark'



RELATED TOPICS

Searching Confluence

Administrators Guide Home Confluence Documentation Home

Enabling CamelCase Linking

CamelCase linking is a form of markup used in many wikis where words are capitalised and compounded together without spaces, 'LikeThis', in order to create links automatically.

By default, CamelCasing is not enabled in Confluence. To use CamelCasing, a Confluence administrator will need to enable this option from the '**Administration Console**'.

To enable CamelCasing,

1. Go to the Confluence '**Administration Console**'. To do this:
    - Open the '**Browse**' menu and select '**Confluence Admin**'. The 'Administration Console' view will open.
2. Select '**General Configuration**' in the left-hand panel.
3. Click '**Edit**' on the '**General Configuration**' screen.
4. Select '**On**' beside '**CamelCase Links**'.
5. Click '**Save**'.

RELATED TOPICS

- Enabling CamelCase Linking
- Enabling Remote APIs
- Enabling Rich Text Editing Option
- Attachment Storage Configuration
- Enabling Trackback
- WebDAV Configuration
- Enabling Threaded Comments

[!Administration Guide Attachments directory^adminhome.gif!] Confluence Documentation Home

Enabling OpenSearch

With OpenSearch autodiscovery, you can add Confluence search to your Firefox or IE7 search box (see Searching Confluence from your Browser's Search Box). By default, OpenSearch autodiscovery is enabled. This feature can be enabled or disabled as described below.

To enable or disable OpenSearch autodiscovery,

1. Go to the Confluence '**Administration Console**'. To do this:
    - Open the '**Browse**' menu and select '**Confluence Admin**'. The 'Administration Console' view will open.
2. Select '**General Configuration**' in the left-hand panel.
3. In the '**General Configuration**' screen, click '**Edit**'.
4. Select '**On**' beside '**Open Search**' to enable this feature, or '**Off**' to disable it.
5. Click '**Save**'.

RELATED TOPICS

Searching Confluence

Administrators Guide Home Confluence Documentation Home

Enabling Remote APIs

Confluence provides XML-RPC and SOAP remote APIs. You need to enable the APIs from the **Administration Console** before you can access Confluence remotely.

You need to have System Administrator permissions in order to perform this function.

To enable the remote API,

1. Go to the Confluence '**Administration Console**'. To do this:
    - Open the '**Browse**' menu and select '**Confluence Admin**'. The 'Administration Console' view will open.
2. Select '**General Configuration**' in the left-hand panel.
3. Click '**Edit**' next to '**Site Configuration**'.
4. Select '**On**' next to '**Remote API (XML-RPC & SOAP)**'.
5. Click '**Save**' to retain your changes.

RELATED TOPICS

No content found for label(s) remoteapi.

Confluencer.NET

Administrators Guide Home Confluence Documentation Home

Enabling Rich Text Editing Option

With Confluence 2.0 and later versions, users have the option of using the Rich Text editor to create pages. By default, this is set to 'On'. If desired, a Confluence administrator can disable '**Rich Text Editing**' from the **Administration Console**.

To disable Rich Text editing,

1. Go to the Confluence '**Administration Console**'. To do this:
    - Open the '**Browse**' menu and select '**Confluence Admin**'. The 'Administration Console' view will open.
2. Select **General Configuration**' in the left-hand panel.
3. In the '**General Configuration**' screen, click '**Edit**'.
4. Select '**Off**' beside '**Rich Text Editing**'.
5. Click '**Save**'.

RELATED TOPICS

Enabling Rich Text Editing Option

Making Rich Text Editing default

Rich Text Editor Overview

Administrators Guide Home Confluence Documentation Home

Enabling the Did You Mean Feature

When you perform a full Confluence search, Confluence may offer you an alternative spelling of your search query. The alternative spelling will appear next to the words 'Did you mean'. By default, this feature is disabled and currently only supports the English language. However, it can be enabled as described below.

To enable the 'Did You Mean' feature,

1. Go to the Confluence '**Administration Console**'. To do this:
    - Open the '**Browse**' menu and select '**Confluence Admin**'. The 'Administration Console' view will open.
2. Select '**General Configuration**' in the left-hand panel.
3. In the '**General Configuration**' screen, click '**Edit**'.
4. Select '**On**' beside '**Did You Mean**'.
   If you have no 'Did you mean' feature index or you have not yet created it, this option will not be available. To create this index, click '**build the did-you-mean index**' and on the subsequent page, click '**Build**' in the '**Did You Mean Index**' section. Then return to the '**General Configuration**' screen in Edit mode.
5. Click '**Save**'.

RELATED TOPICS

Searching Confluence

Administrators Guide Home Confluence Documentation Home

Enabling Threaded Comments

Comments on pages or news items 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. The Confluence administrator can enable or disable the threaded view for the entire Confluence site.

To enable or disable the threaded view,

1. Go to the Confluence '**Administration Console**'. To do this:
    - Open the '**Browse**' menu and select '**Confluence Admin**'. The 'Administration Console' view will open.
2. Select '**General Configuration**' in the left-hand panel.
3. In the '**Feature Settings**' section, click '**Edit**'.
4. Select '**On**' beside '**Threaded Comments**' to enable threaded mode.
   Select '**Off**' to disable threaded mode and display all comments in flat mode.
5. Click '**Save**'.

RELATED TOPICS

Viewing Comments

Commenting on a Page

Administrators Guide Home Confluence Documentation Home

Enabling Trackback

When Trackback is enabled, any time you link to an external webpage that supports Trackback Autodiscovery, Confluence will send a trackback ping to that page to inform it that it has been linked to.

Confluence pages also support Trackback Autodiscovery and when Trackback is enabled, can receive trackback pings sent by other sites.

To enable trackback,

1. Go to the Confluence '**Administration Console**'. To do this:
    - Open the '**Browse**' menu and select '**Confluence Admin**'. The 'Administration Console' view will open.
2. Select '**General Configuration**' in the left panel.
3. In the '**Feature Settings**' screen, click '**Edit**'.
4. Select "**On**' beside '**Trackback**' and click '**Save**'.

RELATED TOPICS

- Hiding the People Directory (Confluence Docs 3.2)
- Anonymous Access to Remote API (Confluence Docs 3.2)
- Spam Prevention via Captcha (Confluence Docs 3.2)
- Hiding External Links From Search Engines (Confluence Docs 3.2)
- User Email Visibility (Confluence Docs 3.2)
- Ignoring External Referrers (Confluence Docs 3.2)
- Managing External Referrers (Confluence Docs 3.2)
- Enabling or Disabling Public Signup (Confluence Docs 3.2)
- Hiding external referrers (Confluence Docs 3.2)
- Excluding external referrers (Confluence Docs 3.2)
- Adding SSL for Secure Logins and Page Security (Confluence Docs 3.2)

Administrators Guide Home Confluence Documentation Home

Making Rich Text Editing default

A Confluence administrator can configure whether the default mode of editing on the site is '**Rich Text**' or '**Wiki Markup**'.

Users will still be able to configure their individual preferences from the '**Edit**' tab of a page.

To make Rich Text Editing the default,

1. Go to the Confluence '**Administration Console**'. To do this:
    - Open the '**Browse**' menu and select '**Confluence Admin**'. The 'Administration Console' view will open.
2. Select '**General Configuration**' in the left-hand panel.
3. Click '**Edit**' on the '**General Configuration**' screen.
4. Select '**On**' beside '**Users see Rich Text Editor by default**'.
   (Select '**Off**' to set '**Wiki Markup**' editing as the default.)
5. Click '**Save**'.

RELATED TOPICS

- Enabling Rich Text Editing Option
- Making Rich Text Editing default
- Rich Text Editor Overview

[!Administration Guide Attachments directory^adminhome.gif!] Confluence Documentation Home

WebDAV Configuration



If you are looking for more information on configuring WebDAV clients, please refer to Using a WebDAV Client to Work with Pages.

On this page:

- Introduction to Confluence's WebDAV Client Integration
- Restricting WebDAV Client Write Access to Confluence
- Disabling Strict Path Checking
- Virtual Files and Folders

Introduction to Confluence's WebDAV Client Integration

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.

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**' page, you can:

- Deny a WebDAV client write permissions to a Confluence installation using a regular expression (regex).
- Disable or enable strict path checking.
- Enable or disable access to specific virtual files/folders.



- The 'WebDav Configuration' page is only be available if the WebDAV plugin has been enabled. (Refer to Installing and Configuring Plugins using the Plugin Repository Client for more information on enabling Confluence plugins). Note that this plugin is bundled with Confluence, and can be enabled or disabled by the System Administrator.
- The settings on the 'WebDav Configuration' page do **not** apply to external attachment storage configuration.

Restricting WebDAV Client Write Access to Confluence

In earlier versions of the WebDAV plugin, separate options for restricting a WebDAV client's write permissions (that is, create/move, edit and delete actions), were available. However, in the current version of this plugin, they have been simplified and combined into a general write permission restriction that covers all of these actions.

WebDAV clients are now denied write permission to your Confluence installation by setting a regex that matches specific content within the WebDAV client's user agent header. Upon setting a regex, it will be added to a list of restricted WebDAV clients. Any WebDAV clients whose user agent header matches a regex in this list will be denied write permission to your Confluence installation.

Example: A PROPFIND method header generated by a Microsoft Web Folder WebDAV client, showing the user agent header field:

```
01.PROPFIND /plugins/servlet/confluence/default HTTP/1.1
02.Content-Language: en-us
03.Accept-Language: en-us
04.Content-Type: text/xml
05.Translate: f
06.Depth: 1
07.Content-Length: 489
08.User-Agent: Microsoft Data Access Internet Publishing Provider DAV
09.Host: 127.0.0.1:8082
10.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 selectively.

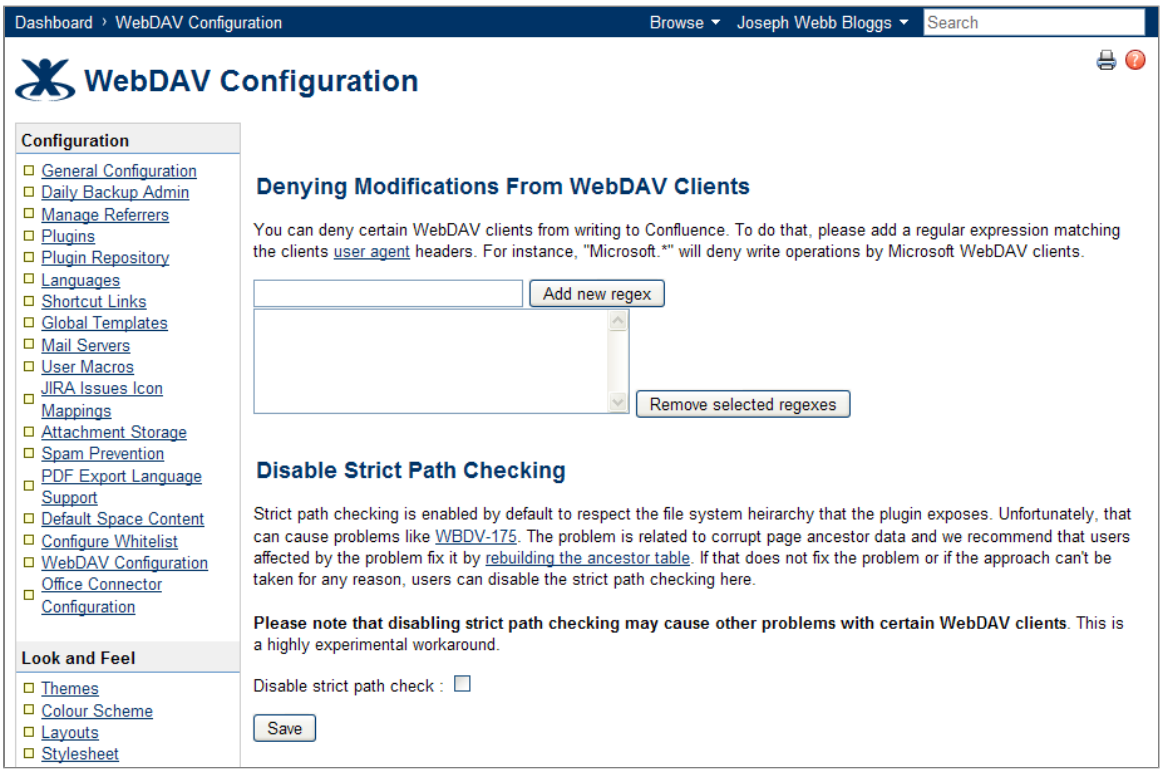
To restrict a WebDAV client's write access permissions to your Confluence installation,

1. Go to the Confluence '**Administration Console**'. To do this:
    - Open the '**Browse**' menu and select '**Confluence Admin**'. The 'Administration Console' view will open.
2. Click '**WebDav Configuration**' under 'Configuration' in the left panel. The 'WebDAV Configuration' page is displayed.
3. Enter a regex that matches a specific component of the user agent header sent by the WebDAV client you want to restrict.
4. Click the '**Add new regex**' button. The regex is added to the list of restricted WebDAV clients.
   You can repeat steps 3 and 4 to add a regex for each additional WebDAV client you want to restrict.
5. Click the '**Save**' button to save the configuration changes.

To restore one or more restricted WebDAV client's write access permissions to your Confluence installation,

1. Go to the Confluence '**Administration Console**'. To do this:
    - Open the '**Browse**' menu and select '**Confluence Admin**'. The 'Administration Console' view will open.
2. Click '**WebDav Configuration**' under 'Configuration' in the left panel. The 'WebDAV Configuration' page is displayed.
3. Select the regex(es) from the list that match(es) the user agent header sent by the restricted WebDAV client(s) you want to restore.
4. Click the '**Remove selected regexes**' button. The regexes you had selected are removed from the list of restricted WebDAV clients.
5. Click the '**Save**' button to save the configuration changes.

Screenshot: WebDAV configuration



Disabling Strict Path Checking

If you observe any idiosyncrasies with your WebDAV client, such as a folder that does exist on your Confluence site but is missing from the client, you can disable the WebDAV plugin's strict path checking option, which may minimise these problems.

To disable the WebDAV plugin's strict path checking option,

1. Go to the Confluence '**Administration Console**'. To do this:
    - Open the '**Browse**' menu and select '**Confluence Admin**'. The 'Administration Console' view will open.
2. Click '**WebDav Configuration**' under 'Configuration' in the left panel. The 'WebDAV Configuration' page is displayed.
3. Clear the 'Disable strict path check' check box.
   You can re-enable this option at a later point in time by simply selecting this check box.
4. Click the '**Save**' button to save this configuration change.

Virtual Files and Folders

In the unlikely event that you observe any problems with the WebDAV client's performance or stability, you can enable access to automatically generated (that is, virtual) files and folders.

> ⚠ By default, these options are hidden on the 'WebDAV Configuration' page. To make them visible, you must append the parameter `?hiddenOptionsEnabled=true` to the end of your URL and reload the page. For example:
>
> ```
> 1.<Confluence base URL>/admin/plugins/webdav/config.action?hiddenOptionsEnabled=true
> ```

Screenshot: The Hidden Virtual Files and Folders Option



To enable or disable access to virtual files and folders,

1. Go to the Confluence '**Administration Console**'. To do this:
    - Open the '**Browse**' menu and select '**Confluence Admin**'. The 'Administration Console' view will open.
2. Click '**WebDav Configuration**' under 'Configuration' in the left panel. The 'WebDAV Configuration' page is displayed.
3. Amend your URL as described in the note above and reload the 'WebDav Configuration' page.
4. Select or clear the check box options in the 'Virtual Files and Folders' section as required.
5. Click the '**Save**' button to save the configuration changes.

RELATED TOPICS

Attachment Storage Configuration (Confluence Docs 3.2)

WebDAV Configuration (Confluence Docs 3.2)

Important Directories and Files (Confluence Docs 3.2)

No content found for label(s) data-storage,webdav.

Administrators Guide Home Confluence Documentation Home

Other Settings

- Configuring Attachment Size
- Configuring Character Encoding
- Configuring HTTP Timeout Settings
- Configuring Indexing Language
- Configuring JIRA Issues Icon mappings
- Configuring Number Formats
- Configuring Shortcut Links
- Configuring Time and Date Formats
- Number of Ancestors to Show in Breadcrumbs
- Thumbnail Settings

Configuring Attachment Size

Confluence gives you the option of limiting the maximum size of a single file attachment. Confluence administrators should keep in mind that

the amount of disk space used by Confluence is directly proportional to the number and size of attachments put into the system.

To configure the maximum size allowed for an attachment

1. Go to the '**Administration Console**' and click '**General Configuration**' in the left-hand panel.
2. Click '**Edit**' on the '**General Configuration**' screen.
3. Enter the maximum size next to '**Attachment Maximum Size**'. The default is 10 MB.
4. '**Save**' your changes.

To configure the maximum index-able size of attachments

By default, large attachment is defined as greater than 1 MB.
The threshold for attachments that won't get excerpts can be modified using the system property `atlassian.indexing.contentbody.maxsize`, which takes a size in bytes.

Example

To specify 250 kb you would use the following JVM parameter:
`-Datlassian.indexing.contentbody.maxsize=256000`

Outcomes of Limiting Attachment Indexing Size

Limiting the size of attachment indexing has the following effects:

- Decreases the size of the index when large attachments are present.
- Decreases the memory used in indexing large attachments.
- Prevent excerpts of large attachments being displayed in search results.

For more details, please refer to the following JIRA issue.

RELATED TOPICS

Configuring Character Encoding
Configuring Time and Date Formats
Configuring Attachment Size
Number of Ancestors to Show in Breadcrumbs
Thumbnail Settings
Configuring Number Formats
Configuring HTTP Timeout Settings
Configuring Indexing Language
Recognised System Properties

Administrators Guide Home Confluence Documentation Home

Configuring Character Encoding

Confluence uses UTF-8 character encoding to deliver its pages.

While it is possible to change the character encoding, unless you are certain of what you are doing, we recommend that you leave this as it is.

To change the character encoding,

1. Go to the '**Administration Console**' and click on '**General Configuration**' in the left panel.
2. Click '**Edit**' at the bottom of the '**Formatting and International Settings**' screen. For Confluence version earlier than 2.6.2, look for the '**Options and Settings**' screen.
3. Beside '**Encoding**', enter the new character encoding of your choice.
4. '**Save**' your changes.

RELATED LINKS

Joel Spolsky: The Absolute Minimum Every Software Developer Absolutely, Positively Must Know About Unicode and Character Sets (No Excuses!)

RELATED TOPICS

Configuring Character Encoding

Configuring Time and Date Formats

Configuring Attachment Size

Number of Ancestors to Show in Breadcrumbs

Thumbnail Settings

Configuring Number Formats

Configuring HTTP Timeout Settings

Configuring Indexing Language

Recognised System Properties

Administrators Guide Home Confluence Documentation Home

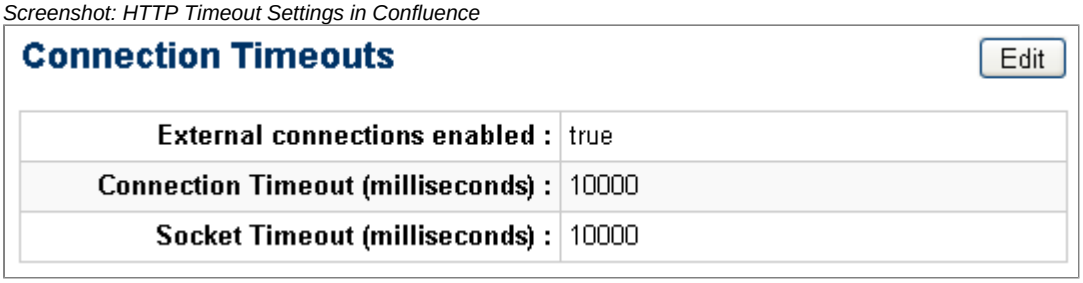
Configuring HTTP Timeout Settings

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. Go to the Confluence '**Administration Console**'. To do this:
    - Open the '**Browse**' menu and select '**Confluence Admin**'. The 'Administration Console' view will open.
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.

Screenshot: HTTP Timeout Settings in Confluence



HTTP Timeout Settings

The settings for HTTP Timeout in Confluence are as follows:

- **Adjust External connections enabled**: This setting allows system administrators to disable external connections so macros like the RSS Macro wont be allowed to make connections to an external server. It's 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).

This feature is available in 2.2.8 and later versions of Confluence. Versions prior to 2.7 have a different method for adjusting these settings (link leads to legacy documentation).

Configuring Indexing Language

Changing the **Indexing Language** setting may improve the accuracy of Confluence search results if the majority of the content of your site is in some language other than English. Since Confluence v. 2.2.5, the content indexing support is provided in ***German, Russian, Chinese,***

CJK, French, Brazilian, Czech and Greek besides English (default).

To configure a different indexing language,

1. Go to the '**Administration Console**' and click on '**General Configuration**' in the left panel.
2. Click '**Edit**' on the right hands side of the '**Formatting and International Settings**' section.
3. There is a drop-down list of '**Indexing Language**' currently supported in Confluence.
4. '**Save**' your changes.

RELATED TOPICS

- Working with Macros
- Creating a Lowercase Page Title Index
- Configuring Indexing Language
- Rebuild the Content Indices from scratch
- Content Index Administration

Configuring JIRA Issues Icon mappings

If you are using the {jiraissues} macro to retrieve information from a JIRA server, you will have to tell Confluence where to find the icons for any custom statuses or issue types you have configured in JIRA.

Confluence is configured by default with all JIRA's standard issue type and status icons. You will only need to change these settings if you have customised additional statuses or issue types for JIRA or have changed JIRA's default icons.

To configure custom icons,

1. Go to the '**Administration Console**' and click on '**Jira Issue Icon Mappings**' in the left panel.
2. For each icon you wish to configure, enter the *name* of the issue type or status into the **Jira entity** field, and the filename of its icon into the **filename** field.

   Ensure that the icon with that filename is located in the `/images/icons` directory of the JIRA server.
3. You may edit existing icon mappings by clicking on the `remove` link by an existing mapping, then re-adding it with a new icon filename.

RELATED TOPICS

- Configuring Character Encoding
- Configuring Time and Date Formats
- Configuring Attachment Size
- Number of Ancestors to Show in Breadcrumbs
- Thumbnail Settings
- Configuring Number Formats
- Configuring HTTP Timeout Settings
- Configuring Indexing Language
- Recognised System Properties

Administrators Guide Home Confluence Documentation Home

Configuring Number Formats

To change the number formats,

1. Go to the '**Administration Console**' and click on '**General Configuration**' in the left panel.
2. Click '**Edit**' at the bottom of the '**Options and Settings**' screen.
    - There are two number format settings:
        - Long Number Format
        - Decimal Number Format
3. Change the formats using the guidelines in this document.
4. '**Save**' your changes.

RELATED TOPICS

- Configuring Character Encoding
- Configuring Time and Date Formats
- Configuring Attachment Size
- Number of Ancestors to Show in Breadcrumbs
- Thumbnail Settings
- Configuring Number Formats
- Configuring HTTP Timeout Settings
- Configuring Indexing Language
- Recognised System Properties

Administrators Guide Home Confluence Documentation Home

Configuring Shortcut Links

Shortcut links provide a quick way of linking to resources frequently referenced from Confluence. When you create a shortcut link, you are assigning a key to a URL so that when a user edits Confluence documents they can type the key instead of the complete URL.

Here is an example:

Most Google searches look like this: `http://www.google.com/search?q=`. If you create a shortcut for this search with the key 'google', every time a user needs to use `http://www.google.com/search?q=`**`searchterms`**, they can just type `[`**`searchterms`**`@google]` instead.

Here is a screenshot showing the shortcuts currently defined on http://confluence.atlassian.com:

Key	Expanded Value	Default Alias	Operations
cache	http://www.google.com/search?q=cache:		Remove
imdb	http://us.imdb.com/Title?		Remove
jira	http://jira.atlassian.com/secure/QuickSearch.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**.

To create a shortcut link,

1. Go to the '**Administration Console**' and click '**Shortcut Links**' in the left panel.
2. Enter a '**Key**' for your shortcut. This is the shortcut name a user will use to reference the URL.
3. 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.
4. *(Optional. Available in Confluence version 2.3 and later.)* Enter a '**Default Alias**'. This is the text of the link which will be displayed on the page where the shortcut is used, with the user's text being substituted for '%s'.
5. Click '**Save**'.

Using Shortcut Links

Specify in the link what should go on the end of the shortcut URL, followed by an at-sign (@) and the key of the shortcut. Shortcut names are case-insensitive. So, for example, using the keys shown in the above screenshot:

To link to...	Type this	Resulting URL	Demonstration
a JIRA issue	[CONF-1000@JIRA]	http://jira.atlassian.com/secure/QuickSearch.jspa?searchString=CONF-1000	CONF-1000
a Google search	[Atlassian Confluence@Google]	http://www.google.com/search?q=Atlassian+Confluence	Atlassian Confluence@Google

Shortcut links can have titles just like any other link:

To link to...	Type this	Resulting URL	Demonstration
Internet Movie Database	[Fight Club|tt0137523@IMDB]	http://us.imdb.com/Title?tt0137523	Fight Club
Google Cache	[Cached Homepage|www.atlassian.com@cache]	http://www.google.com/search?q=cache:www.atlassian.com	Cached Homepage

Deleting Shortcut Links

Once you have created a shortcut link, it is listed under '**Shortcut Links**' in the '**Administration Console**'. Click '**Remove**' to delete the shortcut.

RELATED TOPICS

Administrators Guide Home Confluence Documentation Home

Configuring Time and Date Formats

Confluence allows you to localise the formats used to display dates and times within the web interface. The settings use the syntax of Java's SimpleDateFormat class (described below).

To change the time and date formats,

1. Go to the '**Administration Console**' and click on '**General Configuration**' in the left panel.
2. Click '**Edit**' at the bottom of the '**Options and Settings**' screen.
    - There are three time and date format settings:
        - Time Format : displaying only the time of day (for example, when each news item is posted)
        - Date Time Format : displaying both the date and the time of day (for example, in historical versions of pages)
        - Date Format : displaying only the date (for example, the creation and most recent modification dates of pages)
3. Change the formats using the guidelines in this document.
4. '**Save**' your changes.

RELATED LINKS

- Java 1.4.2 SimpleDateFormat API

RELATED TOPICS

Configuring Character Encoding
Configuring Time and Date Formats
Configuring Attachment Size
Number of Ancestors to Show in Breadcrumbs
Thumbnail Settings
Configuring Number Formats
Configuring HTTP Timeout Settings
Configuring Indexing Language
Recognised System Properties

Administrators Guide Home Confluence Documentation Home

Number of Ancestors to Show in Breadcrumbs

Whenever there are three or more page links to be displayed in the breadcrumbs, Confluence will use an ellipsis like this '...' and display only the topmost and lowermost page links. Clicking on the ellipsis will display the page links in between.

Note that the Dashboard and space homepage links are always displayed at the start of the breadcrumbs, and are not counted as ancestors for the purpose of this setting.

Screenshot: Breadcrumbs



You can configure how many immediate ancestors you want displayed in the breadcrumbs when you go to the page.

To configure the number of ancestors to show in the breadcrumbs,

1. Go to the '**Administration Console**' and click on '**General Configuration**' in the left panel.
2. Click '**Edit**' at the bottom of the '**Options and Settings**' screen.
3. Beside '**Number of Ancestors to show in Breadcrumbs**', enter a number. For example, if you enter 2, two immediate ancestors for the page will be displayed following the dots.
4. '**Save**' your changes.

Configuring Character Encoding
Configuring Time and Date Formats
Configuring Attachment Size
Number of Ancestors to Show in Breadcrumbs
Thumbnail Settings
Configuring Number Formats
Configuring HTTP Timeout Settings
Configuring Indexing Language
Recognised System Properties

Administrators Guide Home Confluence Documentation Home

Thumbnail Settings

The thumbnail settings allow you to define the height and width of images when they are displayed as thumbnails. This affects the images displayed by the Gallery macro and the Thumbnail macro.

To configure thumbnail settings,

1. Go to the '**Administration Console**' and click '**General Configuration**' in the left-hand panel.
2. Click '**Edit**' on the '**General Configuration**' screen.
3. Under the heading '**Attachment Settings**', enter a value in pixels for:
   - **Thumbnail maximum height** — The default setting is 200 pixels.
   - **Thumbnail maximum width** — The default setting is 200 pixels.
4. '**Save**' your changes.

RELATED TOPICS

Displaying a Thumbnail Image

Uploading a Profile Picture

Gallery Macro

Administrators Guide Home Confluence Documentation Home

Configuring System Properties

In general, you can configure system properties by providing an argument of `-Dprop=value` to a Java program when it starts up.

Because Confluence is a Java web application, the Java program is typically your application server. Therefore, you need to configure system properties in your application server's start-up script.

Below is a general example of how system properties are configured for any Java application. The system property arguments can appear anywhere in the argument list.

```
java ... -Dhttp.proxyHost=proxy.example.org -Dhttp.proxyPort=8080
```

The above example configures two system properties: `http.proxyHost` and `http.proxyPort` with values `proxy.example.org` and `8080` respectively.

Sometimes instructions will say simply 'set system property X', without saying what it should be set *to*. In this case, it is usually sufficient to simply use `-Dprop`, without a value. For example:

```
java ... -Datlassian.mail.disable
```

Application Server Examples

Please consult your application server documentation for how to provide system properties to the Java runtime or to a particular application server.

Below are some examples of how to start up different application servers with these system properties set:

AppServer	Startup Script	Variable to Edit	Example
Confluence Stand-alone .zip or .tar.gz Distribution	bin\setenv.bat (Windows) bin/setenv.sh (Unix, Linux OS X)	JAVA_OPTS	`set JAVA_OPTS= -Dhttp.proxyHost=proxy.example.com -Dhttp.proxyPort=3128` -Xms256m -Xmx512m -XX:MaxPermSize=256m (Windows) `export JAVA_OPTS="-Dhttp.proxyHost=proxy.exampl -Dhttp.proxyPort=3128` -Xms256m -Xmx512m -XX:MaxPermSize=256m" (Unix)
Tomcat Windows Service	n/a	n/a	See below

Tomcat .war installation	bin\setenv.bat (Windows) bin/setenv.sh (Unix) (files must be created)	JAVA_OPTS	set JAVA_OPTS=%JAVA_OPTS% `-Dhttp.proxyHost=proxy.example.com -Dhttp.proxyPort=3128` -Xms256m -Xmx512m -XX:MaxPermSize=256m (Windows) `export JAVA_OPTS="-Dhttp.proxyHost=proxy.example.com -Dhttp.proxyPort=3128` -Xms256m -Xmx512m -XX:MaxPermSize=256m" (Unix)
Weblogic	$BEA_HOME/domain/<domain>/startWebLogic.sh	JAVA_OPTIONS	`export JAVA_OPTIONS="$JAVA_OPTIONS -Dhttp.proxyHost=proxy.example.com -Dhttp.proxyPort=3128` -Xms256m -Xmx512m -XX:MaxPermSize=256m"
WebSphere	n/a	Application Server > Process Definition > JVM Arguments	`-Dhttp.proxyHost=proxy.example.com -Dhttp.proxyPort=3128` -Xms256m -Xmx512m -XX:MaxPermSize=256m

Configuring System Properties in Tomcat as a Windows Service

Defining JAVA_OPTS in your `setenv.bat` file will not be sufficient to configure system properties if you are running Confluence with Tomcat as a Windows Service. Rather, there are two ways to configure system properties:

- through the Tomcat configuration application, normally found in the Windows system tray
- updating the service configuration directly in your Windows Registry.

As an example of the latter, for Tomcat 5 you will need to modify `HKEY_LOCAL_MACHINE -> SOFTWARE -> Apache Software Foundation -> Procrun 2.0 -> Tomcat5 -> Parameters -> Java -> Options` and add the JAVA_OPTS parameter there with the necessary system properties. JvmMx and JvmMs are listed separately from the additional options. See Editing the Windows Registry for details.

Displaying the System Properties

To see what Confluence is using, check Displaying System Properties.

RELATED TOPICS

Recognised System Properties

Recognised System Properties

Confluence has a small number of obscure 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. (Refer to the full instructions.)

Property	Since	Default Value	Module...	Effect
`atlassian.forceSchemaUpdate`	1.0	`true`	atlassian-config	By default, Confluence will only ru database schema update when it detects that it has been upgraded flag will force Confluence to perfo schema update on system startup
`confluence.home`	1.0	Any filesystem path	Confluence and atlassian-config	If this system property is set, Confluence will ignore the conten the `confluence-init.proper` file, and use this property as the s for the Confluence Home director
`confluence.devmode`	1.0	`true`	Confluence	Enables additional debugging opt that may be of use to Confluence developers. Do not enable this fla production system.
`confluence.disable.mailpolling`	2.4	`false`	Confluence	If set to "true", will prevent Conflu from retrieving mail for archiving v spaces. Manually triggering "chec new mail" via the web UI will still v This property has no effect on out mail

confluence.i18n.reloadbundles	1.0	true	Confluence	Setting this property will cause Confluence to reload its i18n resc bundles every time an internation string is looked up. This can be us when testing translations, but will Confluence run *insanely slowly*.
confluence.ignore.debug.logging	1.0	true	Confluence	Confluence will normally log a sev error message if it detects that DE level logging is enabled (as DEBU logging generally causes a signifi degradation in system performan Setting this property will suppress error message.
confluence.jmx.disabled	3.0	false	Confluence	If set to "true", will disable Conflue JMX monitoring. This has the san effect as setting the "enabled" pro to false in WEB-INF/classes/jmxContex
confluence.optimize.index.modulo	2.2	20	Confluence	Number of index queue flushes b the index is optimised.
confluence.plugins.bundled.disable	2.9	false	Confluence	Starts confluence without bundlec plugins. May be useful in a develc environment to make Confluence quicker, but since bundled plugins necessary for some of Confluenc core functionality, this property sh not be set on a production system
atlassian.mail.fetchdisabled	1.0	false	atlassian-mail	Disables mail fetching services fo and POP
atlassian.mail.senddisabled	1.0	false	atlassian-mail	Disables sending of mail
atlassian.disable.caches	2.4	true	atlassian-plugins, atlassian-cache-servlet	Setting this property will disable conditional get and expires: head some web resources. This will significantly slow down the user experience, but is useful in devlop if you are frequently changing sta resources and don't want to conti flush your browser cache.
org.osgi.framework.bootdelegation	2.10	empty	atlassian-plugins	Comma-separated list of package names to provide from applicatior OSGi plugins. Typically required v profiling Confluence. For example "com.jprofiler.**,com.yourkit.**".
confluence.diff.timeout	3.1	1000	Confluence	Number of milliseconds to wait fo operation (comparing two page versions) to complete before abor with an error message.
atlassian.user.experimentalMapping	2.10	false	Confluence	Setting this property changes the relationship between local users a local groups to reduce performan degradation when adding a local a local group with a large number users. Please note, setting this pr can slow down other user manag functions. We recommend that yc only if you are experiencing performance problems when addi local users to large local groups. refer to CONF-12319, fixed in Confluence 3.1.1.
confluence.import.use-experimental-importer	3.2	false	Confluence	Setting this property changes Confluence to use the Experimen XML Importer. It is designed to be more stable implementation but, a time of the release of 3.2, the imp largely untested and thus not sup

RELATED TOPICS

Configuring System Properties

Configuring Logging

We recommend that you configure Confluence's logging to your own requirements. You can change the log settings in two ways:

- Configure logging in Confluence Administration – Your changes will be in effect only until you next restart Confluence.
- Edit the properties file – Your changes will take effect next time you start Confluence, and for all subsequent sessions.

Both methods are described below.

Terminology: In log4j, a 'logger' is a named entity. Logger names are case-sensitive and they follow a hierarchical naming standard. For example, the logger named `com.foo` is a parent of the logger named `com.foo.Bar`.

Configure logging in Confluence Administration

You can change some of Confluence's logging behaviour via the **Administration Console** while Confluence is running. Any changes made in this way will apply only to the currently-running Confluence lifetime. The changes are not written to the `log4j.properties` file and are therefore discarded when you next stop Confluence.

Not all logging behaviour can be changed via the Administration Console. For logging configuration not mentioned below, you will need to stop Confluence and then edit the logging properties file instead.

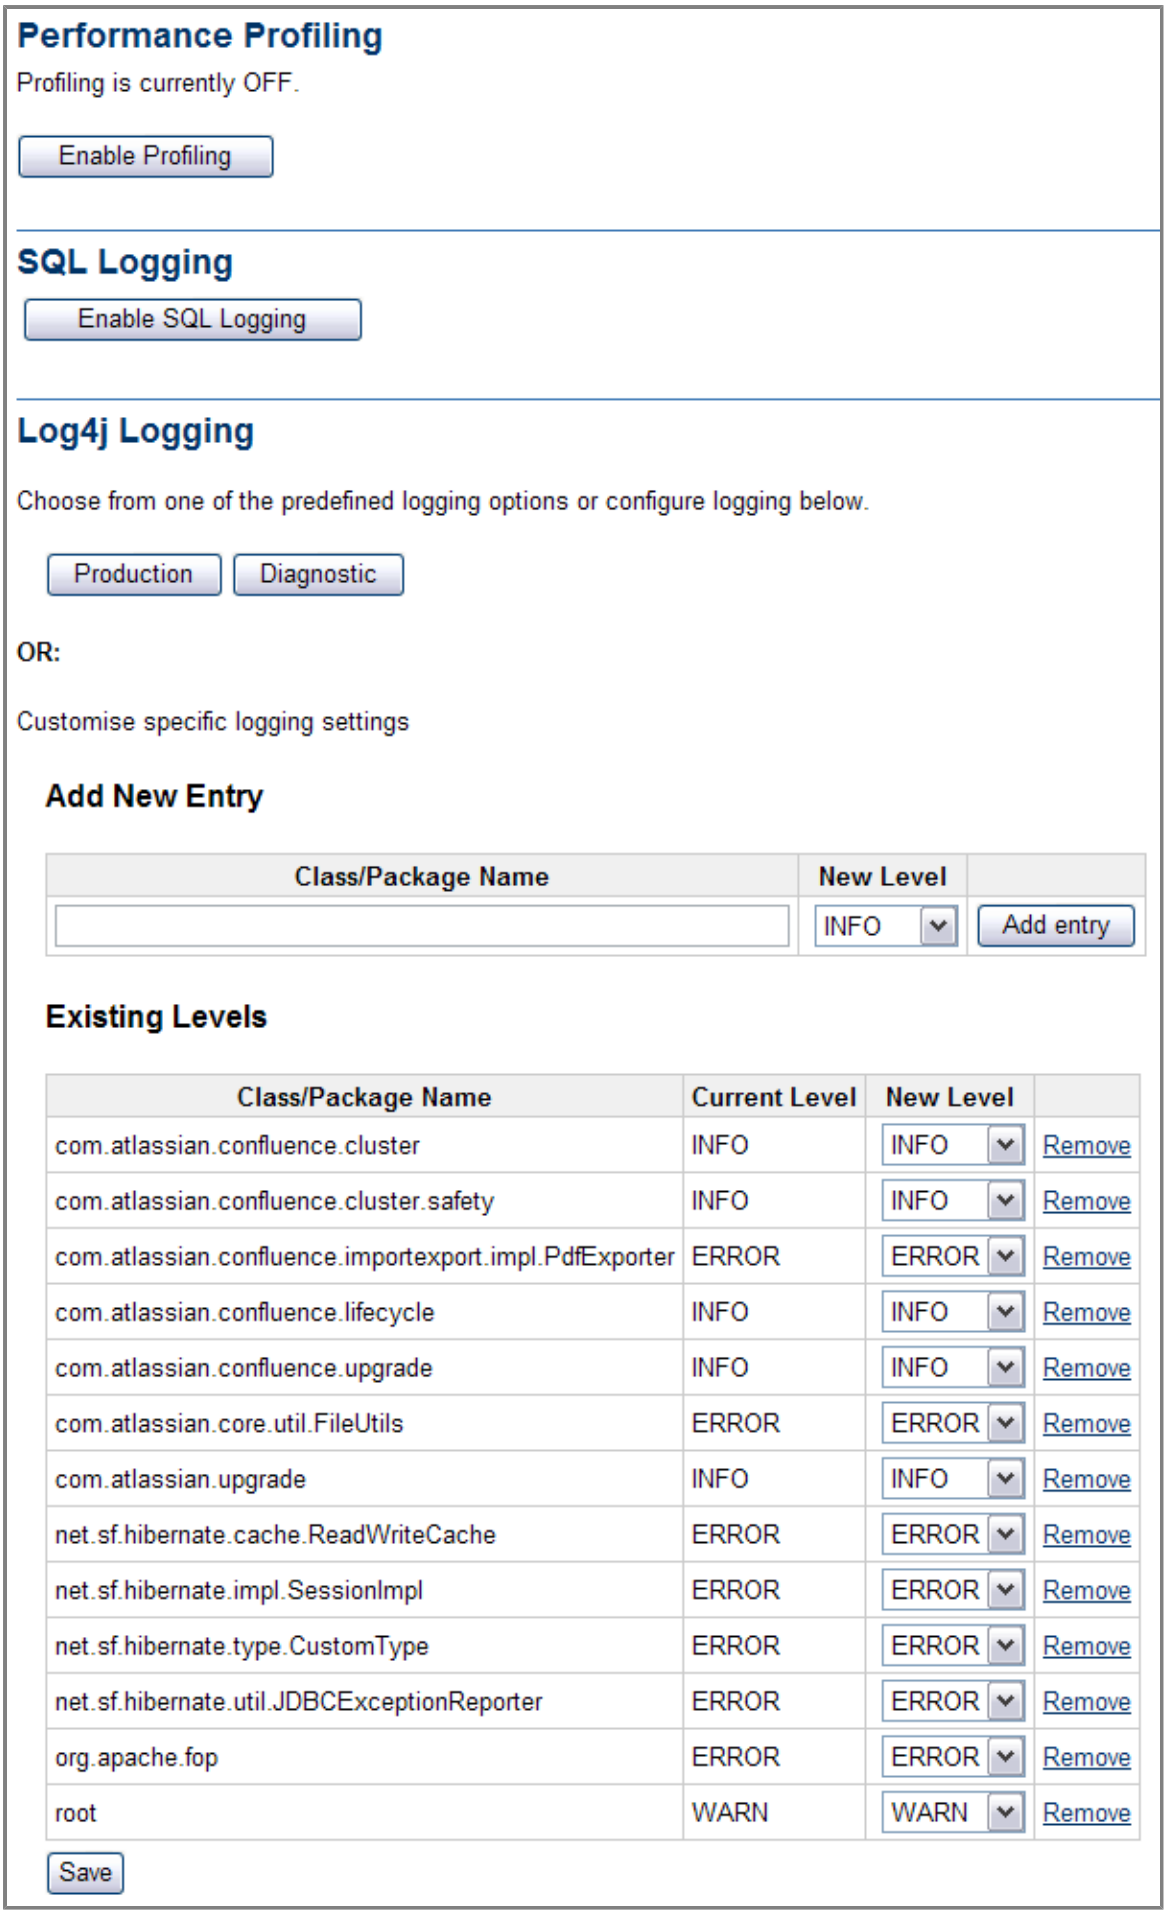
The '**Logging and Profiling**' screen shows a list of all currently defined loggers. On this screen you can:

- Turn page profiling on or off.
- Turn detailed SQL logging on or off.
- Add a new logger for a class/package name.
- Remove a logger for a class/package name.
- Set the logging level (INFO, WARN, FATAL, ERROR or DEBUG) for each class or package name.
- Reset all logging levels to a predefined profile.

Changing the logging configuration

1. Go to the Confluence '**Administration Console**'. To do this:
    - Open the '**Browse**' menu and select '**Confluence Admin**'. The 'Administration Console' view will open.
2. Select '**Logging and Profiling**' in the '**Administration**' section of the left-hand panel.
   You need to have System Administrator permissions in order to perform this function.
3. The '**Logging and Profiling**' screen appears, as shown below. Use the following guidelines to change the logging behaviour while Confluence is running:
    - '**Performance Profiling**' — See Page Request Profiling.
    - '**SQL Logging**' — Click the '**Enable SQL Logging**' button to log the details of SQL requests made to the database.
      If you need to enable logging of SQL parameter values, you will need to change the setting in the properties file. This option is not available via the Administration Console.
    - '**Log4j Logging**' — Click one of the profile buttons to reset all your loggers to the predefined profiles:
        - The '**Production**' profile is a fairly standard profile, recommended for normal production conditions.
        - The '**Diagnostic**' profile gives more information, useful for troubleshooting and debugging. It results in slower performance and fills the log files more quickly.
    - '**Add New Entry**' — Type a class or package name into the text box and click the '**Add Entry**' button. The new logger will appear in the list of '**Existing Levels**' in the lower part of the screen.
    - '**Existing Levels**' - These are the loggers currently in action for your Confluence instance.
        - You can change the logging level by selecting a value from the '**New Level**' dropdown list. Read the Apache documentation for a definition of each level.
        - Click the '**Remove**' link to stop logging for the selected class/package name.
4. Click the '**Save**' button to save any changes you have made in the '**Existing Levels**' section.

Screenshot: Changing Log Levels and Profiling



Performance Profiling

Profiling is currently OFF.

Enable Profiling

SQL Logging

Enable SQL Logging

Log4j Logging

Choose from one of the predefined logging options or configure logging below.

Production | Diagnostic

OR:

Customise specific logging settings

Add New Entry

Class/Package Name	New Level	
	INFO	Add entry

Existing Levels

Class/Package Name	Current Level	New Level	
com.atlassian.confluence.cluster	INFO	INFO	Remove
com.atlassian.confluence.cluster.safety	INFO	INFO	Remove
com.atlassian.confluence.importexport.impl.PdfExporter	ERROR	ERROR	Remove
com.atlassian.confluence.lifecycle	INFO	INFO	Remove
com.atlassian.confluence.upgrade	INFO	INFO	Remove
com.atlassian.core.util.FileUtils	ERROR	ERROR	Remove
com.atlassian.upgrade	INFO	INFO	Remove
net.sf.hibernate.cache.ReadWriteCache	ERROR	ERROR	Remove
net.sf.hibernate.impl.SessionImpl	ERROR	ERROR	Remove
net.sf.hibernate.type.CustomType	ERROR	ERROR	Remove
net.sf.hibernate.util.JDBCExceptionReporter	ERROR	ERROR	Remove
org.apache.fop	ERROR	ERROR	Remove
root	WARN	WARN	Remove

Save

Editing the Properties File

To configure the logging levels and other settings on a permanent basis, you need to stop Confluence and then change the settings in the `log4j.properties` file, described above.

The properties file contains a number of entries for different loggers that can be uncommented if you are interested in logging from particular components. Read more in the Apache log4j documentation.

See Working with Confluence Logs for some guidelines on specific configuration options you may find useful.

External Gadgets

The External Gadgets section allows you to register gadgets served from external web applications (such as JIRA 4.0+) or websites (such as iGoogle or Gmail) with your Confluence installation, so that they:

- Appear in the macro browser
- Can be added and used in Confluence pages or blog posts via a gadget macro

On this page:

- Obtaining the External Gadget's URL
- Registering an External Gadget for Use in Confluence
- Removing Access to an External Gadget in Confluence

Obtaining the External Gadget's URL

Before registering an external web application's gadget with Confluence, you will need to obtain that gadget's URL and copy it to your computer's clipboard.

If your web application is another Atlassian application such as Confluence 3.1+ or JIRA 4.0+, please refer to the appropriate documentation to obtain the gadget URL from your other Atlassian application:

- Obtaining a gadget URL from JIRA 4.0+
- Obtaining a gadget URL from another Confluence 3.1+ server

If your external gadget comes from a non-Atlassian web application or web site, please consult the relevant documentation for that application to obtain the gadget's URL.

Registering an External Gadget for Use in Confluence

To register an external web application's gadget for use in Confluence,

1. Go to the Confluence '**Administration Console**'. To do this:
    - Open the '**Browse**' menu and select '**Confluence Admin**'. The 'Administration Console' view will open.
2. Click 'External Gadgets' under 'Configuration' in the left panel. The 'External Gadgets' page is displayed.
3. In the 'Add a new Gadget' section, paste your gadget's URL into the '**Gadget Specification URL**' field.
4. Click the '**Add**' button. Your gadget will be shown in the list of registered gadgets below and it will also become available in the macro browser.

Do I need to establish an OAuth or Trusted Application relationship too?

In addition to registering an external gadget for use in Confluence, you may also need to establish an OAuth or Trusted Application relationship between the application that serves the gadget (the service provider) and Confluence (the consumer). OAuth and Trusted Application relationships are usually only required for gadgets that access user-restricted data from the external web application. Refer to Configuring OAuth for further information.

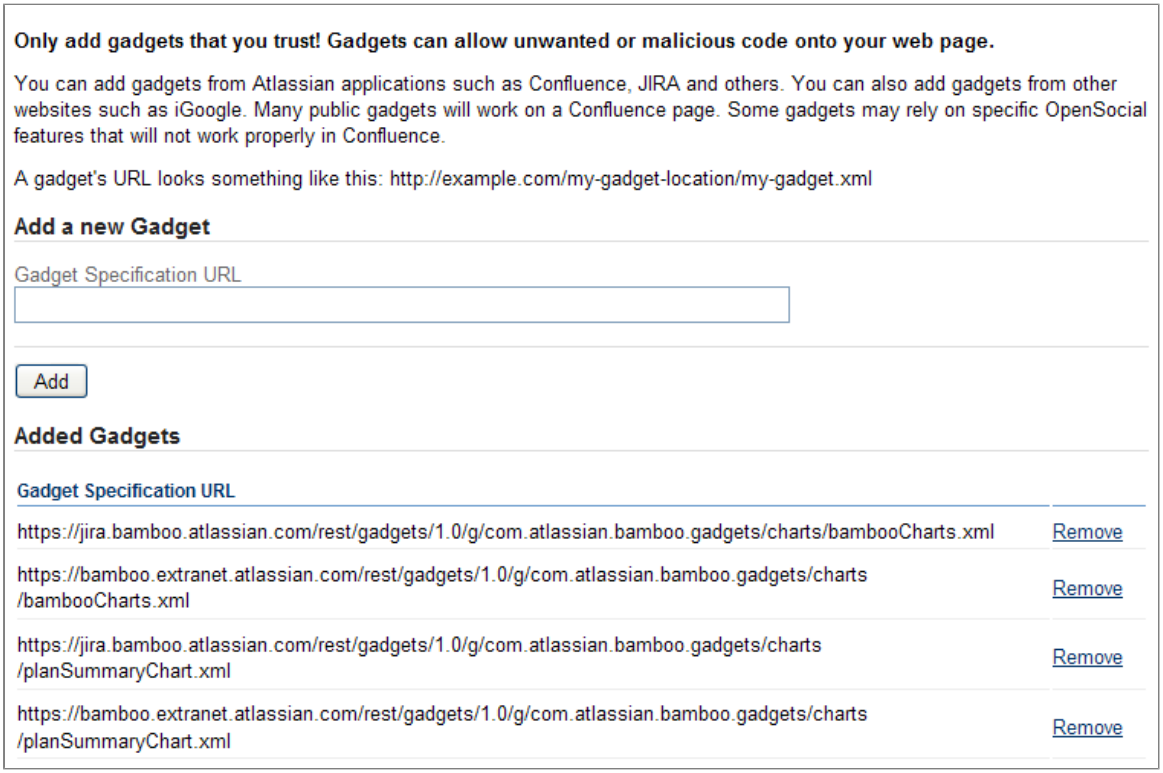
If an external web application provides anonymous access to all or some of its data and that is the only data you need to access, then establishing an OAuth or Trusted Applications relationship may be unnecessary.

Removing Access to an External Gadget in Confluence

To remove Confluence's access to an external web application's gadget,

1. Go to the Confluence '**Administration Console**'. To do this:
    - Open the '**Browse**' menu and select '**Confluence Admin**'. The 'Administration Console' view will open.
2. Click 'External Gadgets' under 'Configuration' in the left panel. The 'External Gadgets' page is displayed.
3. In the 'Added Gadgets' section, click '**Remove**' next to the URL of the external gadget whose access in Confluence is to be removed. The gadget will be removed from the 'Added Gadgets' list and will also be removed from the macro browser.

Screenshot: External Gadgets page



RELATED TOPICS

The big list of Atlassian gadgets

Confluence and JIRA

- Add Confluence EAR-WAR to JIRA Standalone
- Integrating JIRA and Confluence
- Override properties in JIRA to Confluence Bridge
- Setting Up Trusted Communication between JIRA and Confluence

RELATED TOPICS

- Configuring JIRA Issues Icon mappings
- JIRA Issues Macro
- JIRA Portlet Macro

Add Confluence EAR-WAR to JIRA Standalone



Please note that Atlassian Support no longer recommends or supports this configuration. Choose either the EAR/WAR installation or two separate stand-alone installations.

Alternatives

Separate Stand-Alone Installations

A preferable configuration is to run JIRA and Confluence in separate Tomcat instances running behind an Apache frontend server (see guides for Confluence and JIRA).

Advantages

- Each app can be restarted without affecting the other.
- If one webapp hangs for any reason (eg. running out of memory), it doesn't affect the other.
- Any problems can be debugged more easily. Logs are separate and product-specific, rather than everything going to catalina.out. Thread and heap dumps are smaller and more relevant.
- It reduces the likelihood of jar conflicts (eg. jars that must be installed in `common/lib` or `lib` for Confluence running off Apache Tomcat version 6 or above), particularly if you later want to install a third webapp not from Atlassian.

Disadvantages

Offsetting this is the extra complexity of having to run Apache.

Two WAR files in One Application Server

Alternatively, you can use the WAR deployment for both applications in one application server.

Advantages

- The additional memory required by the application server would be shared.

Disadvantages

- The disadvantage is that poor performance in one application would affect the other.

If you are deploying JIRA and Confluence on same web server then please apply the fix described in this KB article to fix the bug with Confluence Browse Menu

See Installing the Confluence EAR-WAR Edition and Installing JIRA WAR-EAR for details.

Integrating JIRA and Confluence



Take a look at the technical guide to the process of adding your Confluence EAR-WAR to JIRA Standalone

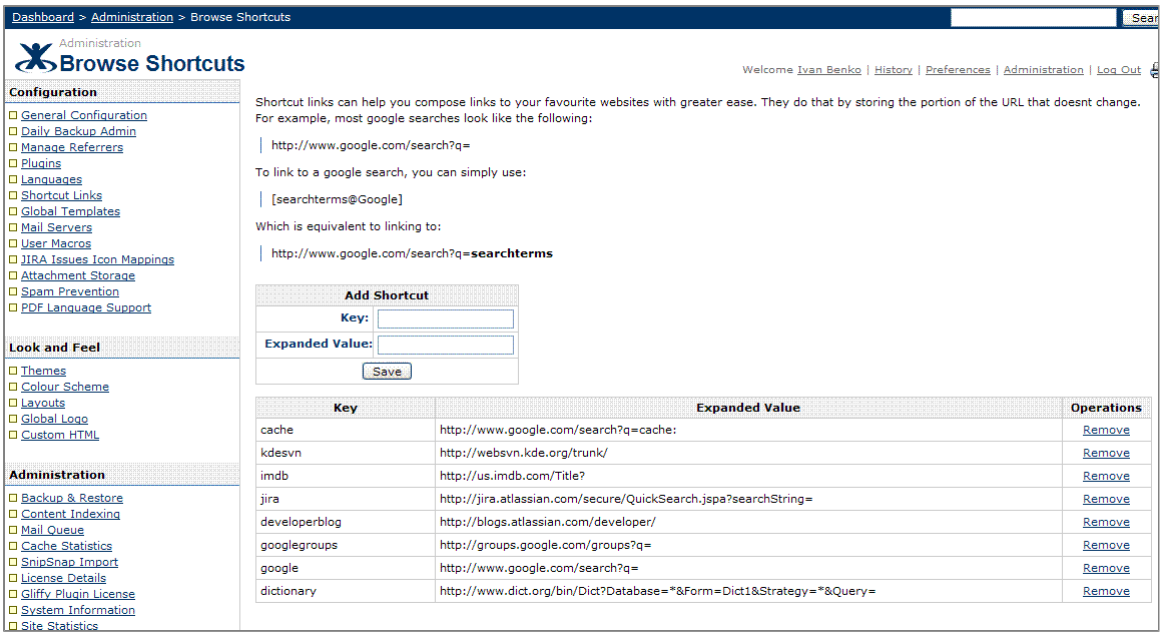
JIRA and Confluence were designed to complement each other. We've all seen projects where people try to store all their knowledge in the issue tracker, and we've seen projects where people have suffered trying to track issues in a knowledge management tool. We say: collect your team's thoughts, plans and knowledge in Confluence, track your issues in JIRA, and let the two applications work together to help you get your job done.

Here are four ways you can get JIRA and Confluence working together: use Confluence shortcuts to make easy links to JIRA issues, use trackback for two-way linking between Confluence and JIRA, use macros to include JIRA reports in Confluence pages, and integrate your JIRA and Confluence user management.

Combine Confluence Shortcuts and JIRA Quick Search

The simplest ideas can often be the most useful. In our Confluence site's global configuration - `Administration > Shorcut Links`, we have the following shortcut defined:

```
JIRA: http://jira.atlassian.com/secure/QuickSearch.jspa?searchString=
```



This way, it's simple to create links using Confluence's shortcut notation. Link directly to JIRA issues: CONF-1000, or use JIRA's intuitive quick-search functionality to create links to particular groups of issue: CONF open improvements will link to a list of all open issues in the Confluence project of type "Improvement" (try it and see!)

Use Trackback for easy two-way linking

Activate Trackback in JIRA and Confluence, and if someone makes a link from one application to the other, the link will automatically lead both ways: create a link from a JIRA issue to an example in a Confluence page, and the Confluence page will automatically know to link back to the JIRA issue, and vice versa. This is the perfect way to keep discussion connected to an issue.

- Document your user stories or use-cases in Confluence, and see at a glance which issues affect each use-case.
- If a JIRA issue requires more discussion or thought than can be conveniently held in comments, link them to a Confluence page.

(Note: as of Confluence 1.0 and JIRA 2.6, there is no mechanism for trackback to log in to JIRA or Confluence, so the use of trackback is limited to pages that are visible to anonymous visitors. In a protected Intranet environment, you may wish to open up Anonymous access to JIRA and Confluence to allow trackback to take place. Future revisions of the applications will give you the opportunity to allow Confluence to "log in" to JIRA and vice versa, avoiding this limitation)

View Confluence content in JIRA or JIRA content in Confluence

Using Gadgets: Confluence 3.1 and JIRA 4.0

Several Confluence macros can be embedded in JIRA's dashboard. Likewise, JIRA gadgets can be rendered on a Confluence page. See Adding a Confluence Gadget to a JIRA Dashboard or Gadget Macro for information on how to set up gadgets for viewing content.

Prior to Confluence 3.1 and JIRA 4.0: use the {jiraissues} and {jiraportlet} macros to embed JIRA reports and portlets into your Confluence site

Any JIRA search-result can be embedded in a Confluence page using the {jiraissues} macro with your choice of included fields and field ordering, and any JIRA dashboard portlet can be embedded in a Confluence page using the {jiraportlet} macro.

This way you can incorporate information from JIRA into the normal flow of your knowledge management. Combined with other macros like {junitreport}, {rss} and {html-include} and the FatCow suite, you can create dashboards in Confluence consolidating information from across your project, with Confluence and JIRA at the centre.

For **Confluence 2.7.0 and later**, an administrator can configure JIRA (3.12.0 or later) and Confluence to communicate in a trusted way, so that Confluence can request information from JIRA on behalf of the currently logged-in user. JIRA will not ask the user to log in again or to supply a password.

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

Read more about trusted communication.

Link to Confluence pages from JIRA issues

While it is possible to simply paste links to Confluence pages into text fields of an issue (e.g. descriptions), the JIRA Linker Plugin provides a custom field that helps you find the correct page.

Integrate JIRA and Confluence user-management

To save you having to enter users into both JIRA and Confluence, you may benefit from using Atlassian Crowd as the user-repository for both applications.

Alternatively you can configure Confluence to use JIRA's user database (this requires that you are using JIRA with an external database; it will not work if you are using JIRA with an embedded HSQL database).

Some useful extensions

1. JIRA Confluence portlet - Display a Confluence page on the JIRA dashboard.
2. Atlassian Activity Stream Plugin - Activity Stream collects information from JIRA, Confluence, FishEye and Crucible.
3. AppLinks Plugin - Allows you to link projects, spaces and repositories between JIRA, Confluence, FishEye, Crucible and SVN applications without the need for long URLs.

And much more coming...

When you buy a license for JIRA or Confluence, you are automatically entitled to a year of updates. We listen to our customers needs, and having our products complement and work well with each other is very important to us. So if there is any way you think Confluence and JIRA could be made to work better, suggest it in our discussion space, and it may very well end up in a future version.

You might also like to take a look at our beyond JIRA page or watch the short video overview on some of these points in .mov format.

Override properties in JIRA to Confluence Bridge

Overriding properties used in the JIRA and Confluence Bridge

If, for some reason, you need to override the name of a column or a table used in Confluence's bridge to JIRA, you may do so in osuser.xml (see below).

This is most likely something you would consider doing if columns names were failing because your database is case sensitive.

```
01.<provider class="bucket.user.providers.CachingCredentialsProvider">
02.<property name="chain.classname"
>com.atlassian.confluence.user.providers.jira.JiraJdbcCredentialsProvider</property>
03.                     <property name="chain.datasource">java:comp/env/jdbc/JiraDS</property>
04.          </provider>
05.          <provider class="bucket.user.providers.CachingAccessProvider">
06.                     <property name="chain.classname"
>com.atlassian.confluence.user.providers.jira.JiraJdbcAccessProvider</property>
07.                     <property name="chain.datasource">java:comp/env/jdbc/JiraDS</property>
08.          </provider>
09.          <provider class="bucket.user.providers.CachingProfileProvider">
10.                     <property name="chain.classname"
>com.atlassian.confluence.user.providers.jira.JiraJdbcProfileProvider</property>
11.                     <property name="chain.datasource">java:comp/env/jdbc/JiraDS</property>
12.                     <property name="chain.configuration.provider.class"
>bucket.user.BucketHibernateConfigProvider</property>
13.          </provider>
```

Simply add the **<property name="chain.PROPERTY_NAME_HERE">NEW_VALUE</property>** element, to override a property (see below) with a new value.

Name Value Pairs for JiraJdbcAccessProvider, JiraJdbcProfileProvider and JiraJdbcCredentialsProvider

Property	Default value
userTable	userbase
userName	userName
userPassword	password_hash
groupTable	groupbase
groupName	groupname
membershipTable	membershipbase
membershipUserName	user_name
membershipGroupName	group_name
userId	id
membershipId	userId

Setting Up Trusted Communication between JIRA and Confluence

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

When JIRA is configured to trust Confluence in this way, we call Confluence the '**trusted application**' and JIRA the '**trusting application**'.

Trusted communication is used when embedding information from one application (e.g. a list of JIRA issues) into another application (e.g. a Confluence page). Currently only JIRA can be configured to trust Confluence, and only the following two macros have been enhanced to use trusted communication:

- JIRA Issues macro
- JIRA Portlet macro

Further implementations will follow, especially as we roll out the tight integration required between Atlassian products for JIRA Studio.

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

On this page:

- Prerequisites
- Why do we need Trusted Communication?
- Overview
- Configuring JIRA to Trust Confluence
- Configuring the Macro Plugin in Confluence
- Adding the Macro to a Confluence Page
- Viewing the Confluence Page
- Security Risks
- Troubleshooting
- Technical Overview of the Trusted Applications Authentication (TAA) Protocol

Prerequisites

- JIRA 3.12.0 or later.
- Confluence 2.7.0 or later.
- In order to authenticate successfully against JIRA, the Confluence user must also be registered as a JIRA user with the same username.

Common user base recommended
It is **highly recommended** that your JIRA and Confluence instances share a common user base, rather than two separate user bases with duplicated usernames. You will receive an error if Confluence passes JIRA a username which JIRA cannot recognise. Also, with separate user bases you run the risk that the same username may be used by two different people. The trusted application does not supply the user's password, so the trusting application will assume the username belongs to the user registered in the trusting application's own user base.

Tip: Try Atlassian Crowd for a tidy user management solution.

Why do we need Trusted Communication?

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

The reasons we require the user credentials are:

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

Overview

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

- A JIRA System Administrator configures JIRA to trust Confluence.
- A Confluence System Administrator configures the macro plugin to use (or not use) trusted communication.
- A Confluence user adds one of the macros to a Confluence page.

- A Confluence user or anonymous user views the Confluence page.

Configuring JIRA to Trust Confluence

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

Using the JIRA Administration Console, the JIRA System Administrator defines Confluence as a trusted application by specifying the Confluence instance's URL and other information. Refer to the JIRA documentation for details.

Configuring the Macro Plugin in Confluence

By default, Confluence ships with trusted communication enabled for the following macros:

- JIRA Issues macro
- JIRA Portlet macro

A Confluence System Administrator can decide on the level of trusted communication used by the macros. The different levels are:

- Ignore trusted communications altogether. Trusted communication is turned off at the global level.
- Perform trusted communications whenever the macro is used on a Confluence page, but do not show certain warning messages.
- Perform trusted communications whenever the macro is used on a Confluence page, and show all warning messages. **This is the default configuration.**

To change the default trusted communication level for the JIRA Macros plugin,

1. Go to the Confluence '**Administration Console**'. To do this:
    - Open the '**Browse**' menu and select '**Confluence Admin**'. The 'Administration Console' view will open.
2. Select '**Plugins**' in the left-hand panel.
3. The '**Plugin Manager**' screen appears, showing a list of installed plugins. Scroll down and click the '**JIRA Macros**' link.
4. The '**JIRA Macros**' panel appears in the top middle of the screen, as shown below. Click '**Enable**' or '**Disable**' next to the following options:
    - '**JIRA application trust support**' – With this option enabled, Confluence will attempt trusted communication with JIRA whenever a user views a page containing the JIRA Issues or Portlet macro, provided criteria are met as described below. With this option disabled, Confluence will never attempt trusted communication with JIRA for these macros.

      Disable the above option if you do not intend to configure trusted communication between JIRA and Confluence.
    - '**JIRA application trust warnings**' – With this option enabled, Confluence will display all error and warning messages that may arise from a problem during trusted communication (assuming that trusted communication is enabled). With this option disabled, Confluence will suppress certain warnings. See troubleshooting below.

      Disable the above option if you have a large number of existing JIRA macros already on your Confluence instance, pointing at a diverse range of JIRA servers. Some of those JIRA servers may have a trusted communication link established (requiring the functionality to be enabled) while other JIRA servers may have no trusted communication link. In this case, you may want to turn off the warning messages so they do not appear on your Confluence pages where the JIRA macros point to non-trusting JIRA servers.

Screenshot: JIRA Macros panel in Plugin Manager



Adding the Macro to a Confluence Page

The Confluence user can add and edit the macros as described on the following pages:

- Using the JIRA Issues macro
- Using the JIRA Portlet macro

> **Remove the username and password from your macro markup code**
> Prior to Confluence 2.7, you needed to include a username and password in the macro markup code if you wanted to display JIRA issues which had restricted viewing. Once your administrator has set up trusted communication between Confluence and JIRA, you no longer need to include a username and password in the markup code for your JIRA macros.

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

What you want to do	Macro parameter	URL parameter	Comments
Display the JIRA issues which the logged-in user is authorised to see. And if the user is not logged in, display only issues which allow unrestricted viewing.			Do not specify any authentication parameters. In this case, the behaviour depends on the way your administrator has set up trusted communication between JIRA and Confluence. Here is a summary of the behaviour. If trusted communication is **enabled**, the authorisation will work seamlessly. When a logged-in user views your page, they will see only the JIRA issues they are allowed to see. And if they are not logged in, they will see only the issues which allow unrestricted viewing. If trusted communication is **disabled**, the Confluence page will show only the JIRA issues which allow unrestricted viewing.

Ensure that Confluence will display only the JIRA issues which allow unrestricted viewing.	`anonymous`		Regardless of who the user is (logged in or not), the Confluence page will show only anonymously-visible issues. Confluence will not attempt to set up a trusted communication link with JIRA in this case.
Use a pre-determined username and password to access the JIRA issues.		`&os_username=MYNAME&os_password=MYPASSWORD`	**Not recommended**. Prior to Confluence 2.7, this was the only way of displaying issues with restricted viewing. For Confluence 2.7 and later, this method will still work. Confluence will not attempt to set up a trusted communication link with JIRA in this case.

Refer to the section below for details of what happens when a user views a Confluence page containing a JIRA macro.

Viewing the Confluence Page

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

- If the macro markup contains an explicit username and password in the URL parameter, Confluence will not request trusted communication with JIRA. Confluence will retrieve the JIRA issues which the specified username is authorised to see. This behaviour is the same as Confluence versions prior to 2.7.
- If the macro markup contains the `anonymous` parameter, Confluence will retrieve only the JIRA issues which allow unrestricted viewing. Confluence will not attempt to set up a trusted communication link with JIRA in this case.
- If the user is anonymous (not logged in), Confluence will retrieve only the JIRA issues which allow unrestricted viewing. Confluence will not attempt to set up a trusted communication link with JIRA in this case.
- If trusted communication is disabled via the Plugin Manager in Confluence, then Confluence will not request trusted communication with JIRA. So if there is no explicit username and password in the markup code, Confluence will retrieve only the JIRA issues which allow unrestricted viewing. This behaviour is the same as Confluence versions prior to 2.7.
- If trusted communication is enabled via the Plugin Manager in Confluence:
  - If the user is logged in, then Confluence attempts trusted communication with JIRA. Confluence sends the username to JIRA. JIRA returns a set of issues which that username is authorised to access, based on the JIRA user base and the JIRA groups and permissions. Confluence displays those issues on the page.
  - If JIRA or Confluence encounters a problem during the trusted communication process, an error message may appear on the Confluence page above the macro output – see troubleshooting below.

Security Risks

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

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

Troubleshooting

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

Warning Message	Cause	Solution	W M C T C
`javax.net.ssl.SSLHandshakeException: sun.security.validator.ValidatorException: PKIX path building failed: sun.security.provider.certpath.SunCertPathBuilderException: unable to find valid certification path to requested target`	JIRA is running over SSL	Add JIRA's SSL Certificate to the Java Keystore	N

`The JIRA server does not recognise your user name. Issues have been retrieved anonymously.`	The logged-in Confluence user is not registered in the JIRA user base.	Add the username to your JIRA user base. It is **highly recommended** that your JIRA and Confluence instances share a common user base.	N
`The JIRA server does not trust this Confluence instance for user authentication. Issues have been retrieved anonymously. You can set the macro to always use an anonymous request by setting the 'anonymous' parameter to 'true'.`	Your JIRA instance has not been configured to trust your Confluence instance.	One of the following solutions: • Configure JIRA to trust Confluence. • Disable trusted communications for the JIRA macros in Confluence. • Use the `anonymous` parameter in all your JIRA Issues and JIRA Portlet macros.	Y
`The JIRA server does not support trust requests. Issues have been retrieved anonymously. You can set the macro to always use an anonymous request by setting the 'anonymous' parameter to 'true'.`	Your JIRA instance is not able to handle trusted communications (i.e. the JIRA version is earlier than 3.12.0).	One of the following solutions: • Download the latest version of JIRA and then configure JIRA to trust Confluence. • Disable trusted communications for the JIRA macros in Confluence. • Use the `anonymous` parameter in all your JIRA Issues and JIRA Portlet macros.	Y
`Failed to login trusted application: confluence:14159892 due to: com.atlassian.security.auth.trustedapps.CertificateTooOldException: OLD_CERT; Certificate too old.`	There is a date/time difference between the JIRA server and Confluence server.	• Certificate Too Old KnowledgeBase Entry	-

Consult Troubleshooting the JIRA Issues Macro and Trusted Applications for further troubleshooting.

Technical Overview of the Trusted Applications Authentication (TAA) Protocol

Read this section if you want a bit more information on the technical side of things.

Atlassian has developed its own protocol to set up trust between JIRA and Confluence. Below is a technical overview of the process.

Configuring JIRA to trust Confluence:

1. When the JIRA System Administrator provides the base URL of the Confluence instance, JIRA requests a trusted application authentication certificate from Confluence. The certificate contains Confluence's trusted application ID and public key (generated specifically for use with the TAA protocol).
2. JIRA validates the certificate and asks the System Administrator for a few extra details about the trust relationship, such as a name for the Confluence instance, timeout, allowed IP addresses and allowed request URLs.
3. JIRA stores all this information in the database.

Making a trusted request from Confluence to JIRA:

1. Confluence sends a web request to JIRA, appending additional headers to the request, including:
    - Timestamp (nonce) of the request + user name of the currently logged-in Confluence user, encrypted with a symmetric key (generated on the fly).
    - The symmetric key, encrypted with Confluence's private key.
    - Confluence's application ID (as displayed when trusted communication was established).
2. JIRA attempts to decode the encrypted headers, using the stored information about the relationship. It conducts the following checks to validate the request:
    - The trusted application ID refers to a valid trusted application.
    - The given username exists in the JIRA user base.
    - The agreed timeout has not expired.
    - The request originated from a trusted IP address.
    - The resource being requested matches those specified in the URL match list.
3. If any of these checks fails, a response is sent to Confluence indicating the reason for failure. Otherwise, JIRA will authenticate the specified user for the duration of the single request, and respond with the resources (i.e. the JIRA issues).

RELATED TOPICS

JIRA Issues Macro
JIRA Portlet Macro
Connect to LDAP, JIRA or Other Services Via SSL
Single Sign-on Integration with JIRA and Confluence
Troubleshooting the JIRA Issues Macro and Trusted Applications

Confluence Clustering Overview

It is possible to run Confluence in a clustered environment instead of on a single server. This means that you can run multiple copies of Confluence in a cluster, so that clients (such as a browser) can connect to any copy and see the same information.



Consider your options carefully before deciding on a clustered installation
While we have tried to make clustering Confluence as easy and administrator-friendly as possible, it is a major architectural change and requires extra planning for deployment and upgrades. Please consider the information on the Cluster Checklist and then consult Atlassian support before making your final decision.

This page gives an overview and links to further pages with information on installing, configuring and administering a Confluence cluster.

Before Deciding to Run a Confluence Cluster

1. Read and consider the details on the Cluster Checklist.
2. Consider the difference between clustering for scalability and clustering for high availability (HA).
3. Contact Atlassian support for further information and advice.

Technical Overview

Read a technical overview of clustering in Confluence.

Server and Network Requirements

- Server hardware requirements
- Technical overview of Confluence clustering
- Diagram of recommended network topology

Installation and Upgrading

There are two methods of installing Confluence in a cluster, depending on whether you have existing data:

- Fresh installation
- Existing data

If you are upgrading an existing Confluence cluster to a new version of Confluence, refer to the cluster upgrade guide.

Configuration and Administration

- Cluster Administration page in the **Administration Console**
- Changing datasources in clusters

Troubleshooting

- Cluster troubleshooting

RELATED TOPICS

Operating Large or Mission-Critical Confluence Installations
Performance Tuning
Requesting Performance Support
Confluence Administrator's Guide
Confluence Configuration Guide

Technical Overview of Clustering in Confluence



Overview of clustering documentation
Refer to the overview of Confluence clustering in the *Administrators' Guide*.

Introduction

From version 2.3, Confluence has had the ability to configure and run multiple copies of itself in a cluster, so that clients can connect to any copy and see the same information. In effect, a Confluence cluster behaves as a single, powerful Confluence installation. While we have tried to make clustering Confluence as easy and administrator-friendly as possible, it is a major architectural change from earlier versions (or non-clustered installations) and consequently, requires extra planning for deployment and upgrades.

This document will give a technical overview of clustering in Confluence, primarily for those users and developers who will be installing and configuring Confluence in a cluster. A separate overview is available for Confluence plugin developers.

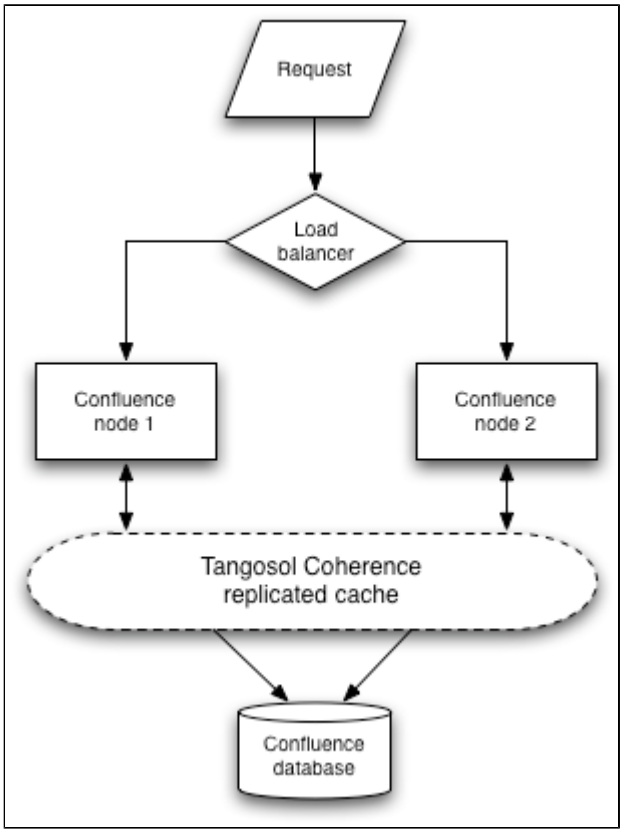
Cluster topology

A simple description of the cluster topology for Confluence would be **multiple applications, shared data source**. A cluster of Confluence consists of:

- multiple homogeneous installations of Confluence (called *nodes* below)
    - a Confluence home directory for each installation.
- a distributed Oracle Coherence cache (formerly known as Tangosol Coherence), which all nodes use via a multicast group - see networking summary below
- a single database, which all nodes connect to

The user is responsible for configuring an appropriate HTTP load balancer in front of the clustered installations. Typically this means using mod_jk or another application server load-balancing technology. The load balancer must be configured to support **session affinity**.

Communication between clustered nodes is minimised by using a distributed cache which propagates updates to all other nodes automatically. Where necessary, Coherence provides a locking mechanism for synchronising jobs and a RMI interface for more complex communication.



Confluence cluster topology (simplified)

LAN Clustering Only

Atlassian **only supports clustering over a local area network**. While it is theoretically possible to configure Confluence to cluster across a WAN, the latency involved is likely to kill performance of the cluster. We can't stop you trying, of course, but you're going to have to work out how to configure Coherence yourself, and we're not going to support the resulting mess.

Homogeneous Confluence installations

All the Confluence installations must be running exactly the same application, down to the lowest level. Items that must be the same include:

- Confluence version
- Application server version
- JDK version
- Libraries and plugins in the Confluence classpath, WEB-INF/lib
- Libraries in the application server classpath

The installation section has more information how to ensure homogeneous node installations.

Creating a Confluence cluster

When installing Confluence in a clustered setup, you will be responsible for configuring your web server and load balancer to distribute traffic between each node. No additional software is required as Coherence is bundled with Confluence.

Here is an overview of the process:

1. Obtain a clustered licence key from Atlassian for each node
2. Upgrade a single node to the clustered licence
3. Start the cluster from that node's administration menu, specifying a name and optionally a preferred network interface
4. Restart the single node and test it
5. Copy the Confluence application and Confluence home directory to the second node
6. Bring up the second node and it will automatically join the cluster.

Copying the Confluence application and home directory helps ensure that the installations are homogeneous.

An alternative to this method is to copy the Confluence web application, but not the Confluence home directory. In this case, the installation wizard will require your cluster name to connect to the other nodes, and it will automatically configure itself. You will need to rebuild the index manually after this installation, however.

There is now full documentation for a Confluence Cluster Installation.

Upgrade process

Another consequence of the homogeneous requirement is that upgrades must be done by following a strict process.

1. All cluster nodes are brought down
2. Upgrade a single node to the latest Confluence version
3. Start the single node so it can upgrade the database
4. Upgrade subsequent nodes and start them one-by-one.

This is the only safe method of upgrading a Confluence cluster.

Single database

The Confluence database in a cluster is shared by all nodes. This means that the database must be able to scale to service *all* the Confluence nodes, which will probably mean implementing some kind of database cluster and JDBC-level load balancing. We can not offer support with scaling or tuning your database, you will need to talk to your DBA or database vendor.

For obvious reasons, you must have an external database to run Massive - you can not cluster Confluence when using the embedded HSQL database.

The most important requirement for the cluster database is that it have sufficient connections available to support the expected number of application nodes. For example, if each Confluence instance has a connection pool of 20 connections and you expect to run a cluster with four nodes, your database server must allow at least 80 connections to the Confluence database. In practice, you may require more than the minimum for debugging or administrative purposes.

In a cluster, attachments must be stored in the database. Configuring a cluster in an existing installation will automatically migrate your attachments to the database. Non-clustered installations still have the option of using the Confluence home directory for storing attachments.

While attachments are stored in the database, they are temporarily written to the cluster node's local filesystem, designated `<confluence-home>/temp` folder, when being streamed to users (so Confluence doesn't have to hold open database connections unnecessarily). For this reason, Confluence will still need enough temporary disk space to hold any attachments currently in transit.

Distributed cache

In a normal configuration, Confluence uses many caches to reduce the number of database queries required for common operations. Viewing a page might require dozens of permissions checks, and it would be very slow if Confluence queried the database for this information with every page view. However, caches must be carefully maintained so they are consistent with the application data. If the page permissions change, the old invalid data needs to be removed from the cache so it can be replaced with a fresh correct copy.

To preserve consistent caches across a cluster, Confluence uses a distributed cache called Oracle Coherence, which manages replicating cache updates transparently across all nodes. The network requirements of the distributed cache are quite simple, but must be preserved if the cluster is to work properly.

To discover other nodes in the cluster, Confluence broadcasts a join request on a multicast network address. Confluence must be able to open a UDP port on this multicast address, or it will not be able to find the other cluster nodes.

Once the nodes are discovered, each responds with a unicast (normal) IP address and port where it can be contacted for cache updates. Confluence must be able to open a UDP port for regular communication with the other nodes.

Because the Coherence network requirements are different to those required by the Confluence database connection, the situation can arise where Confluence can use the database but not talk to the other nodes in the cluster via Coherence. When Confluence detects this, it will shut itself down in a cluster panic.

For more details on the network configuration of the distributed cache, see the networking summary

Home directory

Confluence's home directory has a much-reduced role in a cluster. Because the application data must be shared between all nodes for consistency, the only information stored in the Confluence home directory is either node-specific, or needed to start Confluence. This includes information related to:

- database connection
- license

- cluster connection

The only application data stored in the Confluence home directory is the **Lucene search index**. Confluence synchronises this data itself by keeping track of indexing tasks in the database.

This is also why we recommend copying the Confluence home directory from the first node when setting up subsequent nodes. If you did not copy the Confluence home directory, you would need to rebuild the search index from scratch on the subsequent nodes after installation.

Event handling

Broadcasting events to all nodes in a cluster is supported in Confluence, but not recommended. The cluster topology uses a shared data store so that application state does not need to be synchronised by events.

The event broadcasting is done only for certain events, like installing a plugin. When a plugin is installed in one node, Confluence puts the plugin data in the database, and notifies the other nodes that they need to load the plugin into memory.

Indexing

Confluence maintains a copy of its Lucene search index on each node of the cluster. This index is used for many things beside full-text searches, including RSS feeds and lists of recently updated content. Indexing in a cluster works like this:

1. Node 1 gets a request to save some page update
2. After saving the page in the database, Node 1 adds a "page-updated" index entry to the queue, which is in the database
3. Periodically, each node picks up the "latest entries" from the queue, where what is latest is determined from a timestamp on a file in the Confluence home directory which indicates when the queue was last inspected. This process is called "flushing the index queue".
4. Each node independently updates its local Lucene index. The "page-updated" index entry is internally changed into a delete-document task and an add-document task to apply the changes to Lucene.
5. Each node updates the timestamp on its index-queue-timestamp file to reflect the most recent processing or "flushing" of the index queue.

Because of step #3, if the timing of the nodes is not synchronised or changes sporadically (due to a virtualisation environment, typically), index changes will not be correctly synchronised in the cluster. This is the most common cause of index sync problems in clusters.

If a node is disconnected from the cluster for a short amount of time (less than three hours), it will be able to bring its copy of the index up-to-date when it rejoins the cluster. If a node is down for a long amount of time and its lucene index has become stale as a result, you may want to avoid the expensive operation of rebuilding the index. To do that, you must copy a "live" version of the Lucene index from an active node. Simply replace the contents of the `Confluence Home]/index` directory with those from an active node before bringing the stale node back up.

Job synchronisation

For tasks such as sending the daily report emails, it is important that only one node in the cluster does this. Otherwise you would get multiple emails from Confluence every day.

Confluence uses locks in the Coherence distributed cache to ensure only one node can be running certain jobs at a time. This ensures email notifications will only be sent once.

Activity tracking

Activity tracking does not work in a cluster, and will be disabled for clustered deployments. We're working on making the activity tracker clusterable in a future release. You can follow this issue. You can try some other options for tracking usage.

Cluster panic

In some situations, there can be a network issue or firewall that prevents the distributed cache from communicating but still allows Confluence to update the database. This is a dangerous situation because when the caches on the detached nodes become inconsistent, users on different nodes will see different information and updates can be lost.

Confluence can detect this problem by checking a database value against a cached value, and if they differ, all the clustered nodes will be shut down with a 'Cluster panic' message. This is considered a fatal error because the consequences can cause damage to your data. For those administrators that like to live on the edge, there is a system property to prevent cluster panic and allow data corruption. For more information, see Cluster safety mechanism.

If a cluster panic does occur, you need to ensure proper network connectivity between the clustered nodes. Most likely multicast traffic is being blocked or not routed correctly. See the networking summary below.

Summary of network requirements

In addition to normal connectivity with its database, all clustered Confluence instances require access to a multicast group and the ability to open a UDP unicast port.

By default, the multicast address is automatically generated from the cluster name you provide when starting the cluster and the multicast port is fixed. During cluster setup, Confluence will prompt for the unicast IP address to use if the server has multiple network interfaces, and by default the unicast port is fixed. The cluster multicast group will be joined on the same network interface as the bound unicast IP address.

For any settings which are not configurable through the Confluence web interface, they can be configured via an XML file in the Confluence

home directory for more exotic networking requirements.

Scaling Confluence On A Single Server

Since the maximum addressable memory on a 32 bit JVM is 4GB, some large servers may scale Java applications by running JVM instances concurrently. This would be implemented as separate, clustered Confluence nodes running on a single server and communicating internally. Because each JVM replicates the cache entirely, it may be useful to test a single, massive instance running a 64 bit JVM as an alternative. This configuration may result in superior performance than an internal cluster.

Geographically Distributed Clusters

Collocating nodes is strongly recommended as high latency will almost certainly degrade performance due to the overhead of cache replication. Cluster nodes will provide the best performance if servers are physically adjacent. However, as long as all nodes share a LAN, users may wish to test alternative configurations to see how performance is affected.

RELATED TOPICS

Server Hardware Requirements Guide
Overview of Confluence Clusters

Cluster safety mechanism

Introduction

A mechanism was added in Confluence 2.3 and above to ensure database consistency when running multiple cluster nodes against the same database. This is called the *cluster safety mechanism*, and is designed to ensure that your wiki cannot become inconsistent because updates by one user are not visible to another. A failure of this mechanism is a fatal error in Confluence and is called *cluster panic*.

Because the cluster safety mechanism helps prevents data inconsistency whenever any two copies of Confluence running against the same database, it is enabled in *all* instances of Confluence, not just clusters.

How cluster safety works

A scheduled task, ClusterSafetyJob, runs every 30 seconds in Confluence. In a cluster, this job is run only on one of the nodes. The scheduled task operates on a *safety number* – a randomly generated number that is stored both in the database and in the distributed cache used across a cluster. It does the following:

1. **Generate** a new random number
2. **Compare the existing safety numbers**, if there is already a safety number in both the database and the cache.
3. **If the numbers differ, publish a ClusterPanicEvent**. Currently in Confluence, this causes the following to happen:
    - disable all access to the application
    - disable all scheduled tasks
    - update the database safety number to a new value, which will cause all nodes accessing the database to fail.
4. If the numbers are the same or aren't set yet, **update the safety numbers**:
    - set the safety number in the database to the new random number
    - set the safety number in the cache to the new random number.

How to fix it

This page has been migrated to Database is being updated by an instance which is not part of the current cluster.

Technical details

The cluster safety number in the database is stored in the `CLUSTERSAFETY` table. This table has just one row: the current safety number.

Changing Datasources Manually in a Cluster



The recommended way of changing database connections is to shut down the whole cluster, install Confluence into new and empty directories and use the Setup Wizard to configure all new database connection settings.

However, if you wish to manually change your settings, you may proceed as described below.

 It is **strongly recommended** that you test all of the following in a staging or test instance of Confluence before performing these steps in your production environment.

Step 1: Prepare

- Locate the `confluence-cfg.xml` file in the Confluence home directory.
- Make a backup copy of that file.
- Prepare the necessary changes to that file.

Step 2: Shut Down Confluence

You need to shut down all the nodes in the cluster, not just one.

Step 3: Apply your Changes

Apply your configuration changes to the required node.

Step 4: Restart the Changed Node

It is crucial that you bring up the node on which you applied the changes **first**. Otherwise you will get an error message, and have to shut down all instances again.

Step 5: Restart all Other Nodes

Done.

RELATED PAGES

Overview of Confluence Clusters

Cluster Troubleshooting

Clustering is vastly improved in Confluence 3.0. See the Confluence 3 Performance improvements page for more information. If you're having cluster performance problems in a version prior to Confluence 3, we suggest an upgrade!

Overview of clustering documentation
Refer to the overview of Confluence clustering.

This page covers troubleshooting for the Clustered Edition. If you're experiencing Cluster Panic messages in a Standard Edition, visit the Knowledge Base article Database is being updated by an instance which is not part of the current cluster.

On this page:

- Symptoms
- Confluence cluster debugging tools
  - Multicast
  - Mapping interface to IP address.
  - Debugging tools
  - Add multicast route
  - Check firewall
  - Prefer IPv4
  - Change multicast interface
  - Increase multicast TTL
  - Check intermediate routers
  - Advanced Tangosol configuration
- Didn't find a solution?
- Related

Symptoms

Below is a list of potential problems with a Confluence cluster, and their likely solutions. The solutions are listed below.

Problem	Likely solutions
`Database is being updated by an instance which is not part of the current cluster` errors on a stand-alone	Database is being updated by an instance which is not part of the current cluster
`Database is being updated by an instance which is not part of the current cluster` errors on a cluster	Add multicast route, Check firewall
`Cannot assign requested address` on startup, featuring an IPv6 address	Prefer IPv4
Error in log: `The interface is not suitable for multicast communication`	Change multicast interface, Add multicast route
Multicast being sent, but not received (detectable with Multicast Test)	Check firewall, Check intermediate routers, Increase multicast TTL

Any issue not covered here	Contact support

Confluence cluster debugging tools

There is an umbrella issue opened for all cluster debugging tools here

It includes the tools listed below.

Multicast

- Which multicast address?

The multicast address and port used by Confluence can be found on the Cluster Administration page, or in `confluence.cfg.xml` in the Confluence home directory.

- Multicast address generation.

Confluence uses a hashing algorithm to take the inputted name during setup and it is then turned into a multicast address stored in the config file. Thus, once the initial setup is completed, Confluence will use the address this is the reason why user can change the address if needed, without actually changing the name. Consequently the additional nodes using the same multicast address specified in the config file are able to join the cluster.

Each node has a multicast address configured in the `confluence-cfg.xml` file

```
1.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.
Multicast Test	Coherence tool for testing multicast traffic from one node to another.
`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 the Multicast Test tool shows that multicast traffic can't be sent or received correctly, adding a route for multicast traffic on the correct interface will often fix the problem. The example below is for a Ubuntu Linux system:

```
route add -net 224.0.0.0 netmask 240.0.0.0 dev eth0
```

To support multiple applications using multicast on different interfaces, you may need to specify a route specific to the Confluence multicast address.

Check firewall

Ensure your firewall allows UDP traffic on the multicast address and port used by Confluence.

Prefer IPv4

There's a known issue with IPv6, especially on Linux.

The fix is to add `-Djava.net.preferIPv4Stack=true` to `JAVA_OPTS`. This tells the JVM to try binding an IPv4 address first, and resort to IPv6 only if that fails.
Note: *A more radical approach is to add NETWORKING_IPV6=no to /etc/sysconfig/network, yet probably should be left for a later consideration on a production machine.*

Change multicast interface

Confluence might have selected the incorrect interface for multicast traffic, which means it cannot connect to other nodes in the cluster. To override the interface used for multicast traffic after initial setup, edit `confluence.cfg.xml` in the Confluence home directory and add a property (or change the existing one) to select your desired network interface. For example to tell Confluence to use `eth1`:

```
1.<property name="confluence.cluster.interface">eth1</property>
```

Increase multicast TTL

The multicast time-to-live (TTL) specifies how many *hops* a multicast packet should be allowed to travel before it is discarded by a router. It should be set to the number of routers in between your clustered nodes: 0 if both are on the same machine, 1 if on two different machines linked by a switch or cable, 2 if on two different machines with one intermediate router, and so on.

Create a file in the Confluence home directory called `tangosol-coherence-override.xml`. Add the following to it, setting the TTL value appropriately (1 is the default):

```
1.<?xml version='1.0'?>
2.<coherence>
3.<cluster-config>
4.     <multicast-listener>
5.          <time-to-live system-property='tangosol.coherence.ttl'>1</time-to-live>
6.     </multicast-listener>
7.</cluster-config>
8.</coherence>
```

Alternatively, simply start Confluence with the system property: `-Dtangosol.coherence.ttl=1`. Again, 1 is the default value, and you should change it to something appropriate to your network topology.

Check intermediate routers

Advanced switches and routers have the ability to understand multicast traffic, and route it appropriately. Unfortunately sometimes this functionality doesn't work correctly with the multicast management information (IGMP) published by the operating system running Confluence.

If multicast traffic is problematic, try disabling advanced multicast features on switches and routers in between the clustered nodes. These features can prevent multicast traffic being transmitted by certain operating systems.

For best results, use the simplest network topology possible for the cluster traffic between the nodes. For two nodes, that means a single network cable. For larger numbers, try using a single high-quality switch.

Advanced Tangosol configuration

If the solution to your problem involves changes to the Tangosol configuration, these changes should **not** be made to the Confluence configuration in `confluence/WEB-INF/classes/`. Instead, to ensure your configuration survives upgrades, make your changes via:

- Tangosol system properties
- creating a `tangosol-coherence-override.xml` file in the Confluence home directory.

Examples of making these changes are shown in the increasing the TTL section.

Didn't find a solution?

Check Related Articles from the Confluence Knowledge Base

- Changing Datasources Manually in a Cluster
- Cluster safety mechanism
- Cluster Troubleshooting
- Apache and Tomcat load balancing
- Recommended network topology
- Cluster Administration page
- Cluster Panic triggers
- Technical Overview of Clustering in Confluence
- How do I supress cluster warning message in confluence?
- Confluence Cluster Installation
- Upgrading a Confluence Cluster
- Confluence Clustering Overview
- Viewing and Editing License Details

Open JIRA Features and Bug Reports

JIRA Issues (45 issues)

Type	Key	Summary	Assignee	Reporter	Priority	Status	Resolution	Created	Update
	CONF-8959	Attachment migration does not happen when upgrading to a clustered license	Unassigned	Nicholas Ilacqua [Atlassian]		Open	Unresolved	Jul 19, 2007	Jul 02, 2009
	CONF-14120	Hibernates UpdateTimestampsCache doesn't handle concurrent writes	Unassigned	Chris Kiehl [Atlassian]		Open	Unresolved	Jan 05, 2009	May 05 2009
	CONF-12689	Support Confluence cluster upgrades without an outage	Unassigned	Igor Minar		Open	Unresolved	Aug 06, 2008	Sep 10, 2008
	CONF-12287	Coherence cache fails while retrieving profile picture metadata (dashboard or view page shows UnexpectedRollbackException)	Unassigned	Matt Ryall [Atlassian]		Open	Unresolved	Jul 01, 2008	Feb 15, 2010
	CONF-9297	Confluence should be able to automatically recover from cluster panics	Unassigned	Gary Weaver		Open	Unresolved	Aug 27, 2007	Mar 25, 2009
	CONF-10054	OSUser provider does not generate IDs correctly in clustered environment	Unassigned	Christopher Owen [Atlassian]		Open	Unresolved	Nov 26, 2007	Dec 01, 2008
	CONF-10980	Cluster debugging/troubleshooting tools	Unassigned	Ivan Benko [Atlassian]		Open	Unresolved	Mar 06, 2008	Mar 06, 2008
	CONF-12486	ClassNotFoundException logged on cluster node startup	Unassigned	Anatoli Kazatchkov [Atlassian]		Open	Unresolved	Jul 17, 2008	Aug 25, 2009
	CONF-10325	Viewing the members of a group in a clustered environment works only on one node and not the other.	Unassigned	Partha Kamal [Atlassian]		Open	Unresolved	Dec 27, 2007	Jul 02, 2009
	CONF-9594	ConditionalPropertySet's cannot be cached breaking cluster installations that delegate user management to JIRA	Unassigned	Dave Loeng [Atlassian]		Open	Unresolved	Sep 28, 2007	Jul 02, 2009
	CONF-14657	Retrieving the global settings in a clustered environment causes a lot of contention	Unassigned	Chris Kiehl [Atlassian]		Open	Unresolved	Feb 21, 2009	Nov 08, 2009
	CONF-10868	Node that can not join cluster due to license restriction causes	Unassigned	Ivan Benko		Open	Unresolved	Feb 29,	Sep 03,

		cluster panic		[Atlassian]				2008	2008
	CONF-16419	Installing a font for PDF export in a cluster will not carry to cluster nodes that are down or unavailable.	Unassigned	Charles Miller [Atlassian]		Open	Unresolved	Jul 20, 2009	Aug 05, 2009
	CONF-14948	Support failover NICs for cluster configuration...	Unassigned	Tony Atkins [Atlassian]		Open	Unresolved	Mar 19, 2009	Mar 19, 2009
	CONF-10323	Coherence Lock being held when it appears no thread should have the lock. Causes ConcurrentModificationException	Unassigned	Paul Curren [Atlassian]		Open	Unresolved	Dec 26, 2007	Jan 14, 2010
	CONF-8716	Determine index mismatch in cluster and warn on cluster info page	Unassigned	Matt Ryall [Atlassian]		Open	Unresolved	Jun 17, 2007	Oct 29, 2007
	CONF-12345	Park issue :)	Unassigned	Dave Loeng [Atlassian]		Open	Unresolved	Jul 07, 2008	Jan 21, 2009
	CONF-9813	Disable attachments migration to Filesystem in Cluster	Unassigned	Gurleen Anand [Atlassian]		Open	Unresolved	Oct 24, 2007	Oct 25, 2007
	CONF-9040	Authenticator (subclass of DefaultAuthenticator) can be called twice at almost exactly same time by 2 or more clustered servers	Unassigned	Gary Weaver		Open	Unresolved	Jul 30, 2007	Nov 04, 2007
	CONF-14045	Cluster mysql tests corrupts the database	Unassigned	Brian Nguyen [Atlassian]		Open	Unresolved	Dec 21, 2008	Jan 08, 2009
	CONF-13421	Layout customisations are not propagated to other cluster nodes	Unassigned	Matt Ryall [Atlassian]		Open	Unresolved	Oct 16, 2008	Dec 09, 2008
	CONF-14076	SchedulerException when running cluster builds	Unassigned	Brian Nguyen [Atlassian]		Open	Unresolved	Dec 29, 2008	Jan 01, 2009
	CONF-12614	Intermittent ConcurrentModificationException in cluster	Unassigned	Anatoli Kazatchkov [Atlassian]		Open	Unresolved	Jul 29, 2008	Mar 31, 2009
	CONF-9335	In cluster, allow attachments to be stored on file system in network-shared directory	Unassigned	Jeremy Largman [Atlassian]		Open	Unresolved	Aug 29, 2007	Feb 18, 2010
	CONF-18241	Uninstalling a plugin from a cluster sometimes fails	Matthew Jensen [Atlassian]	Don Willis [Atlassian]		Open	Unresolved	Jan 13, 2010	Mar 27, 2010
	CONF-10953	Support unicast addressing in cluster when well-known-addresses WKA are defined	Unassigned	Ivan Benko [Atlassian]		Open	Unresolved	Mar 06, 2008	Aug 25, 2009
	CONF-10981	Check how many nodes/processes running in a cluster and their identity	Unassigned	Ivan Benko [Atlassian]		Open	Unresolved	Mar 06, 2008	Mar 06, 2008
	CONF-12421	Don't use distrbuted cache for storing Captcha's in a cluster	Unassigned	Paul Curren [Atlassian]		Open	Unresolved	Jul 13, 2008	Jul 13, 2008
	CONF-9749	Coherence does not allow the disabling of all JDK shutdown hooks	Unassigned	Christopher Owen [Atlassian]		Open	Unresolved	Oct 17, 2007	Jan 29, 2008
	CONF-10979	List confluence cluster interface	Unassigned	Ivan Benko [Atlassian]		Open	Unresolved	Mar 06, 2008	Mar 06, 2008
	CONF-13870	After a site Import into a cluster, admin console displays attachment storage as filesystem	Unassigned	Agnes Ro [Atlassian]		Open	Unresolved	Nov 27, 2008	Dec 03, 2008
	CONF-10977	Generate new Multicast address from a "new" cluster name	Unassigned	Ivan Benko [Atlassian]		Open	Unresolved	Mar 06, 2008	Sep 11, 2008
	CONF-9020	Cluster nodes do not get notified of Layout changes	Unassigned	Roberto Dominguez		Open	Unresolved	Jul 27, 2007	Oct 08, 2007
	CONF-9059	Cluster setup network interface selection shows loopback interface	Unassigned	Matt Ryall [Atlassian]		Open	Unresolved	Aug 01, 2007	Aug 01, 2007

	CONF-10635	Database logging of clustersafety access	Unassigned	James Fleming [Atlassian]		Open	Unresolved	Feb 03, 2008	Feb 03, 2008
	CONF-9281	Plugin's I18n properties not loaded in other cluster nodes unless restarted	Unassigned	Roberto Dominguez		Open	Unresolved	Aug 26, 2007	Feb 15, 2010
	CONF-13698	Changing custom html on one node of a cluster is not imideatly reflected on the other node.	Unassigned	Anatoli Kazatchkov [Atlassian]		Open	Unresolved	Nov 12, 2008	Nov 13, 2008
	CONF-7368	Confluence Cluster setup dies horribly when DNS fails	Unassigned	Don Willis [Atlassian]		Open	Unresolved	Nov 23, 2006	Jan 19, 2009
	CONF-14338	Specify an arbitrary multicast port for a cluster	Unassigned	James Fleming [Atlassian]		Open	Unresolved	Jan 30, 2009	Jan 30, 2009
	CONF-17040	Cannot build milestones outside Atlassian due to coherence	Unassigned	Jonathan Gilbert [Atlassian]		Open	Unresolved	Sep 25, 2009	Dec 08, 2009
	CONF-14088	Locking on cache keys needs to check if the lock was actually aquired	Unassigned	Chris Kiehl [Atlassian]		Open	Unresolved	Dec 30, 2008	Nov 22, 2009
	CONF-9712	Plugins which don't work in a cluster shouldn't look like an error	Unassigned	Gary Weaver		Open	Unresolved	Oct 15, 2007	Oct 16, 2007
	CONF-9846	Tangosol configuration: the (optional) cluster-name element is in the wrong place	Unassigned	Don Willis [Atlassian]		Open	Unresolved	Oct 30, 2007	Dec 20, 2007
	CONF-8300	Cannot override TTL configuration through tangosol coherence properties	Unassigned	Matthew Jensen [Atlassian]		Needs Verification	Unresolved	Apr 20, 2007	Nov 10, 2009
	CONF-10330	Certain Multicast IP addresses sending CHANGE_TO_EXCLUDE_MODE causing IGMP traffic to be blocked	Unassigned	Partha Kamal [Atlassian]		Needs Verification	Unresolved	Dec 28, 2007	Nov 26, 2008

Contact Atlassian support

We have dedicated staff on hand to support your installation of Confluence. Please follow the instructions for raising a support request and mention that you're having trouble setting up your Confluence cluster.

Related

Cluster Safety Mechanism

Cluster Panic triggers

This page has been migrated to Database is being updated by an instance which is not part of the current cluster.

Multicast Test

This page describes the **Multicast Test**, a Coherence tool for testing multicast traffic from one node to another. You may find this useful when troubleshooting a clustered installation of Confluence.

In order to run the Multicast test, you need to first download the attached Coherence zip file.

The Multicast Test comes as a script called `multicast-test`, which you will find located in the `bin` folder in the above zip file.

Instructions on how to run this script file can be found in the Coherence documentation. You may like to go straight to the subheading called '**Example**' in the guide, where there is an example on how to use the `multicast-test` script.

RELATED TOPICS

Cluster Troubleshooting
Confluence Clustering Overview

Clustering for Scalability vs Clustering for High Availability (HA)

People occasionally enquire about setting up High-Availability (HA) Confluence clusters. Confluence's clustering is designed to solve a different problem, that of scaling under high load. This page explains the difference.

On this page:

- What is High Availability (HA)?
- What does Confluence's clustering do, then?
- So what kind of resilience can I build into a Confluence installation?
- What's the difference between load balancing and failover?
- What do you mean by 'session affinity'?

What is High Availability (HA)?

HA means that your application *will be* available, without interruption. It's a very difficult thing to achieve, and is typically what people are talking about when they refer to five-nines availability.

In the context of application clustering, it means that any given node (or combination of nodes) can be shut down, blown up, or simply disconnected from the network unexpectedly, and the rest of the cluster will continue operating cleanly as long as at least one node remains. It requires that nodes can be upgraded individually while the rest of the cluster operates, and that no disruption will result when a noes rejoins the cluster. It typically also requires that nodes be installed in geographically separate locations.

What does Confluence's clustering do, then?

Confluence's clustering system allows a single installation to serve a much greater number of concurrent requests than a single server. This is what we refer to as 'scaling under load'.

It does provide a certain amount of resilience, as the death of one node won't bring the other(s) down. However, it requires very low network latency, which rules out geographic separation of the servers, and upgrading can only be performed while the entire cluster is shut down. This doesn't mean that Confluence's clustering is buggy or broken. It simply reflects the difference between the two design aims.

So what kind of resilience can I build into a Confluence installation?

It's still entirely possible to build a resilient Confluence installation, using a 'cold-failover' approach in which two (or more) servers share a database and (normally) a network-mounted file system, where no more than one server is actually running at any given time.

Several different approaches are feasible, but the common elements are:

- a well-configured load balancer (session affinity is irrelevant in this case)
- a reliable monitoring system which can detect and shut down a misbehaving Confluence instance before starting the spare server
- startup scripts with added smarts to check for the presence of another running node before deciding whether to start up a server
- servers with the same view of both the database and the home directory.



It's vital to ensure that only one server is running at any one time, in this kind of setup. If a server starts while another is already running against the same database, the result will be a cluster panic that shuts down both servers.

A single database becomes the single point of failure in such a system. This can be alleviated by database clustering, or by replication from the 'active' database server to the standby server(s) if you wish to separate the failover systems while keeping database latency to a minimum.

In the same vein, the home directory can be hosted on a shared network system — SAN or NAS, preferably with its own replication/rapid recovery system — though there's a known issue to consider. Alternatively, to avoid the use of networked file systems, a utility such as `rsync` can be used to periodically bring the spare servers' home directories up to date, so long as you keep the period sufficiently short — probably between one and five minutes, depending on the rate of activity. This can be avoided altogether by keeping attachments in the database; it increases the demands on the bandwidth between the application and database servers, but guarantees that the system is in a consistent state at switchover. If the data is at all sensitive or confidential, it's advisable to run `rsync` over `ssh`, to minimise the opportunity for the data to be captured on its way across the network.

What's the difference between load balancing and failover?

Load balancing means that all servers are active, and new requests are distributed among them. Several strategies are available, but the most common are:

- round-robin — the first request goes to the first server, the second request goes to the second server, and so on. When you run out of servers, the next request goes to the first server, and around it goes again.
- percentage-based — if (for example) you have two servers, and one can handle twice the load of the other, you can tell the load balancer to send two requests to the stronger server for every request that goes to the weaker one.
- availability — the load balancer sends a test query to each of the servers every second or so, and directs each new request to the server that's currently responding the fastest.

Failover means that only one server is active at any given time, and normally involves two servers (any number of servers may be involved, depending on the system). If the active one stops responding, requests are directed to the other server — the system 'fails over' to the second one.

'Cold failover' means that the second server is only started up after the first one has been shut down. This is the case for non-clustered Confluence.

'Hot failover' or 'hot standby' means that all servers are running at all times, and that the load is directed entirely toward one server at any one time.

A load balancer can be used in both scenarios, especially if it's smart enough to keep track of which servers are currently running.

Failover can also be managed via DNS, in a sufficiently well-controlled environment.

What do you mean by 'session affinity'?

Sessions consist of several transmissions in each direction between the client (browser) and the server. Session affinity means that the load balancer keeps track of which server received the initial transmission from a given browser, and that it will then send any subsequent requests from that browser to the same server.

This is necessary with Confluence clustering, in particular, because sessions are not shared across cluster nodes. If you log into one node and then send a request to another, the other node will send you the login screen because it doesn't recognise your session cookie.

RELATED TOPICS

Confluence Clustering Overview

Recommended network topology

Atlassian recommends a network topology similar to the one shown below, to get the best results from a Confluence Clustered deployment.

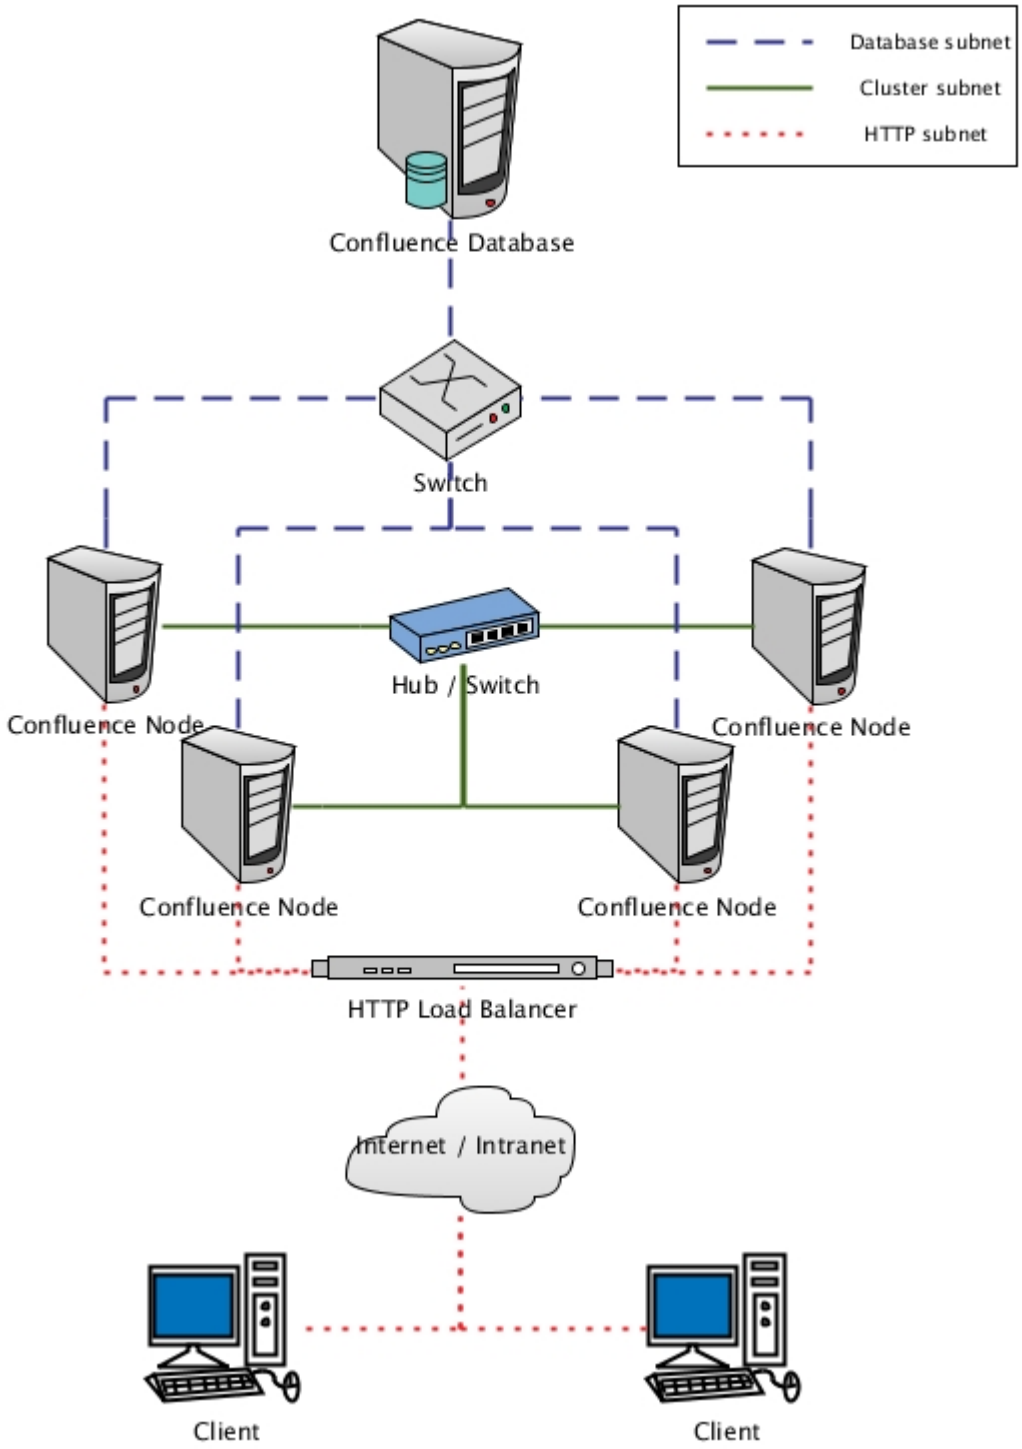
The number of Confluence nodes in the deployment is adjustable — select the number which suits your own requirements.

The most important aspect is that cluster, database and HTTP (client) traffic are all carried on separate subnets. It is possible, on a sufficiently fast network, to carry cluster and database traffic on the same subnet but we do strongly recommend that HTTP traffic be always confined to a separate subnet on production deployments.

Confluence Clustered does not support clustered communication over WAN, VLAN or VPN. All Confluence Clustered nodes must be on the same local subnet, ideally networked via an ethernet hub or simple switch. The cluster communication network must also support multicast IP networking.



Use this example as a basis for your own network diagram
When you are considering a Confluence Clustered deployment, you should prepare a network diagram like the one on this page. This will facilitate discussion with Atlassian Support and help with your own planning. Please refer to the cluster checklist for more guidance on planning your clustered deployment.



Cluster Administration page

Overview of clustering documentation
Refer to the overview of Confluence clustering.

Overview

Any instance of Confluence which uses a clustered license has a Cluster Configuration page which includes information about the active cluster.

To open the Cluster Administration page,

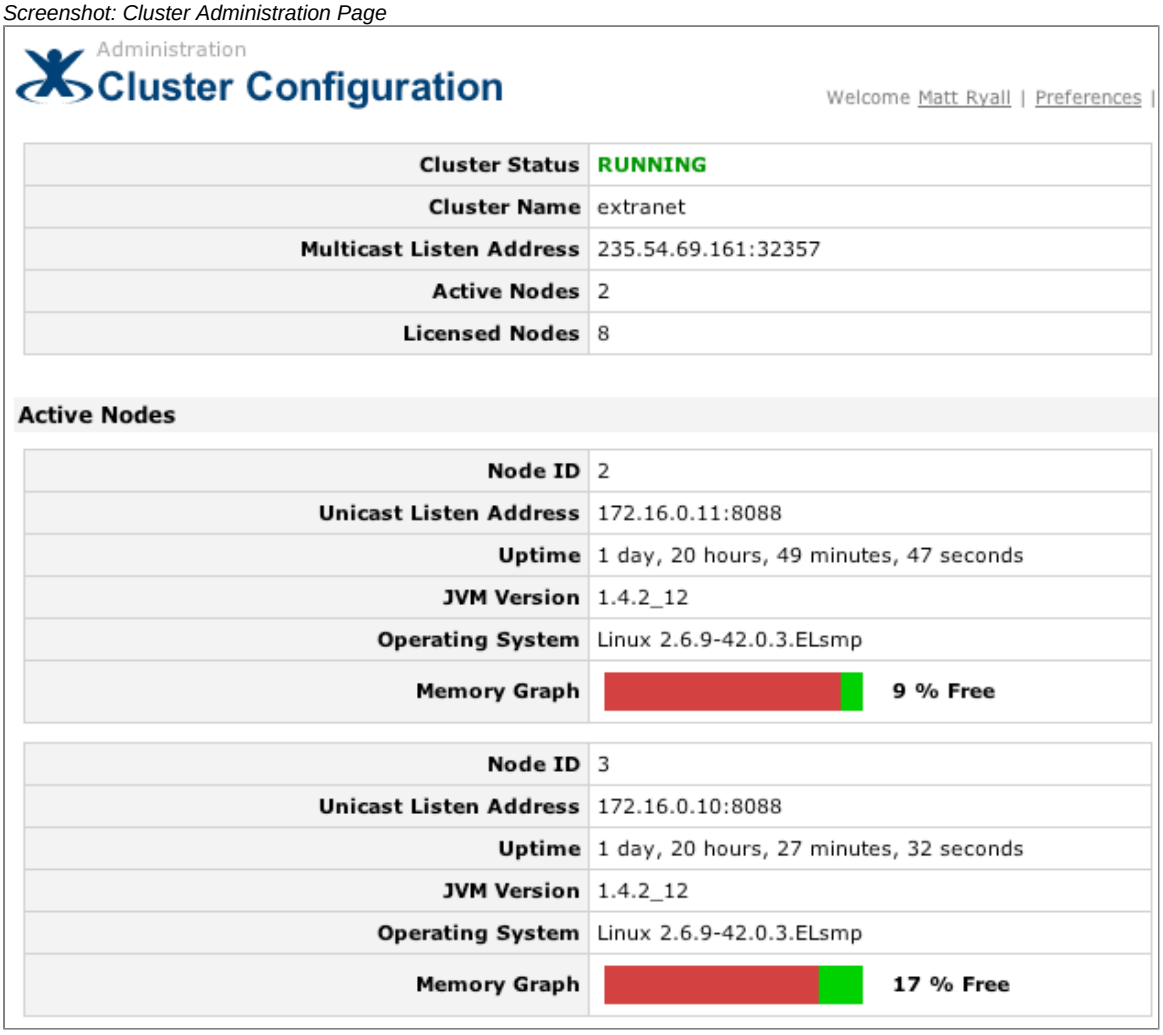
1. Go to the Confluence '**Administration Console**'. To do this:
    - Open the '**Browse**' menu and select '**Confluence Admin**'. The 'Administration Console' view will open.
2. Click '**Cluster Configuration**' in the left-hand menu, in the section called 'Clustering'.

Availability

To access this functionality, you must:

- Be a System Administrator (i.e. have global System Administrator permissions), and
- be using Confluence 2.3 or later, and
- be using a clustered Confluence license.

Screenshot: Cluster Administration Page



This page shows your cluster configuration, and allows you to start a new Confluence cluster using data from this instance.

Cluster Status indicates whether your cluster is currently running.

Licensed nodes is the maximum number of instances of Confluence your license allows in a cluster.

Active nodes lists the instances of Confluence currently participating in the cluster.

Starting a new cluster will perform the following changes:

- enable a clustered cache
- migrate attachments from file system to the database
- publish database connection information so other nodes can join the cluster.

All access to Confluence will be locked while this takes place, and you will be forced to restart Confluence afterwards.

Cluster name is a short name for identifying your cluster. Other Confluence instances can join the cluster using this name.

To join an existing cluster, start a clean copy of Confluence on this node and select 'Join Cluster' during the setup wizard.

Related documents

Overview of Confluence Clusters
Confluence Cluster Installation
Cluster Troubleshooting

Cluster Checklist

It is possible to run Confluence in a **clustered** environment instead of on a single server. This means that you can run multiple copies of Confluence in a cluster, so that clients (such as a browser) can connect to any copy and see the same information.

Refer to the clustering overview for more information and a list of related pages about clustering Confluence.

Consider your options carefully before deciding on a clustered installation
While we have tried to make clustering Confluence as easy and administrator-friendly as possible, it is a major architectural change and requires extra planning for deployment and upgrades. Please consider the information below and then consult Atlassian Sales before making your final decision.

Summary of the information on this page:

- Purpose of this Document
- Assumed Knowledge
- General Considerations
  - Confluence Clustered is designed to scale the number of simultaneously connected users at a much better performance than what a single node can achieve
  - Confluence Clustered will not improve performance in systems with few users.
  - Confluence Clustered is not a high availability solution.
  - Confluence Clustered is not for disaster recovery nor for transparent failover.
- Server Setup
  - The number of supported cluster nodes is limited to four.
  - All cluster nodes must have the same version of OS, application server, etc.
  - Use good and up-to-date hardware.
  - Confluence Clustered is not supported when run in VMware or other virtualisations.
  - Confluence should be the only application on the cluster servers.
  - Do not upgrade and switch to Confluence Clustered at the same time
- Database Setup
  - Run the database on its own physical server.
  - Attachments must be stored in a database and not the local file system
  - Make sure that you use a supported version of a database server to store Confluence's data.
  - Your database must be provisioned to store a large volume of binary data.
  - You need an experienced DBA available to troubleshoot database performance issues.
- Network Setup
  - We recommend hardware load balancers or putting a software loadbalancer onto its on server.
  - Use separate network adapters for communication between servers.
  - The switch connecting the Confluence cluster nodes must not be a 'smart switch'.
  - Cisco switches need additional configuration.
  - It is recommended that the database is on a different physical network from the Confluence server nodes.
  - Minimize the latency between the Confluence cluster nodes and the database.
  - Prepare a network diagram.
  - You need network support staff available to troubleshoot cluster communication issues.
- Staging Environment
  - You need a staging environment that is exactly the same as your production system.

Purpose of this Document

The purpose of this cluster checklist is to help you:

- Decide whether Confluence Clustered is the right solution for you.
- Create a plan for your clustered deployment.

As a service to our customers, we offer to review your deployment plan and make recommendations to help you avoid common pitfalls. To make use of this service, please consider all the information below carefully while planning your clustered deployment. Then contact Atlassian Pre-Sales for recommendations.

If you need to raise a support request with Atlassian during or after cluster deployment, we will need to ask you questions about your configuration. It will save crucial time if you can provide us with your deployment plan.

For more information about clustering Confluence, refer to the clustering overview.

Assumed Knowledge

In writing this document, we have assumed that our readers have an in-depth knowledge of the following technical areas:

- Database
- Networking
- Application servers
- Load balancers

Before starting a clustered deployment please read the information on this page carefully, as well as the linked documentation, to assess if you have the assumed knowledge.

General Considerations

What will Confluence Clustered do for you?
The points in this section of the page will help you evaluate your reasons for considering a clustered deployment, and then decide whether Confluence Clustered is the right solution for your environment.

Confluence Clustered is designed to scale the number of simultaneously connected users at a much better performance than what a single node can achieve

Confluence Clustered will not improve performance in systems with few users.

Clustering Confluence means that user requests can be served by independent machines. The performance gains are substantial, and have improved a lot further since Confluence 3.0. Clustering is especially great in dealing with spikes to the load, e.g. during certain hours of business. Just note that if rendering a complicated page (e.g. containing many macros or rendering many graphs) takes five seconds on an otherwise idle server, it will not be faster in a clustered environment. Also, the first step when you encounter performance issues is to tune your existing system, make sure you are using the right hardware and have looked at your database.

Confluence Clustered is not a high availability solution.

Confluence Clustered is not designed specifically to provide a high availability solution.

General availability is higher in a Confluence cluster than on a single installation, you can for example take one node down for minor maintenance tasks e.g. when adding a new CPU or adding RAM. But you still have to bring down all nodes at the same time for software upgrades. Also there are certain conditions, like loss of network connectivity between nodes ('split brain'), that will result in the cluster shutting itself down. Confluence Clustered offers higher reliability, but not high availability.

Confluence Clustered is not for disaster recovery nor for transparent failover.

If one node crashes, there is no transparent failover for the connected client. Also, our network requirements (see below) make Confluence unsuitable for deployment to different cities or even to different buildings.

Server Setup

The number of supported cluster nodes is limited to four.

Not supported. In theory, you can connect more than four nodes — but that is not covered by Atlassian Support.

All cluster nodes must have the same version of OS, application server, etc.

Confluence requires a homogeneous environment. All Confluence cluster nodes must have the same version of the following:

- Operating system
- CPU
- Installed memory
- Java
- Application server

Note that 'same version' means '**same to the last digit**'. For example, Java v1.4.2_16 is not the same as v1.4.2_15

We strongly recommend user to have the same memory configuration (both the JVM and the physical memory) because a cluster uses a replicated cache. A replicated cache requires the same amount of memory on each node in the operating cluster. The memory allocations must be equal.

Use good and up-to-date hardware.

While the details are up to you, we strongly suggest that your servers have at least 4GB of physical RAM. A high number of concurrent users means that a lot of RAM will be consumed. You usually don't need to assign more than 4GB per JVM process, and most of the time even just 1GB or 2GB will be fine, you should just be prepared to fine tune the settings.

Confluence Clustered is not supported when run in VMware or other virtualisations.

Not supported. We strongly discourage you to deploy a production environment of Confluence to virtual servers, and we will not be able to support you when problems arise.

When running a Confluence cluster your goal is high capacity and performance, so you should not risk lower performance by virtualising it and sharing a computer with other processes.

Many customers who are running Confluence on VMware, or similar virtualisation solutions, experience major performance problems that are extremely hard to pinpoint. Since the problems are not related to Confluence itself, we will not be able to help you.

Confluence should be the only application on the cluster servers.

No additional applications (other than core operating system services) should be running on the same servers as Confluence.

Since your goal should be increased capacity and performance, you should not risk this by running any other process on the machine with a Confluence Clustered node. While it may be fine to run JIRA, Confluence and Bamboo on a dedicated Atlassian software server for small installations, it is strongly discouraged for clustering Confluence.

Do not upgrade and switch to Confluence Clustered at the same time

If you plan to migrate to a clustered solution, make sure you are migrating within the same version of Confluence. If you plan to upgrade to a higher version of Confluence, do this **before** the migration to the clustered version.
For example, if you are currently running Confluence 2.9.2 standalone, and want to roll out the clustered version of Confluence 3.0, you must first upgrade to Confluence 3.0 standalone and check that everything works fine (e.g. by running and monitoring your production system for a week). Then you are in a good position to migrate to the clustered version.

Database Setup

Run the database on its own physical server.

You are optimising for performance, so you don't want the database to slow down your application servers, or vice versa. In high load scenarios, the database may need to have better hardware than the application servers to be able to handle all requests. You should find out by performing loadtesting.

Attachments must be stored in a database and not the local file system

Storing attachments in the database is the only supported attachment storage configuration for clustering Confluence.

Make sure that you use a supported version of a database server to store Confluence's data.

Please check that your intended database is officially supported by Atlassian Confluence. The load on an average cluster solution is higher than on a single box installation, and it is therefore even more crucial to use the right database vendor and version.

Your database must be provisioned to store a large volume of binary data.

Note that Confluence clustered stores file attachments in the database, and you need an experienced DBA who can monitor and manage the data growth.

You need an experienced DBA available to troubleshoot database performance issues.

Not having an experienced full-time DBA at hand at short notice when entering the realm of high load is dangerous. While small installations of Confluence basically work 'out of the box', anything that involves high load and a lot of database space requires continual monitoring, optimising and fine tuning of the Confluence database. When we ramp up the load on our loadtesting environment, we see that database usage goes up as well. Having powerful hardware in place helps, but if there are queries that become inefficient with you particular load pattern, you need an expert to tune it. As an example, we have seen PostgresSQL switch its internal caching mechanism when a particular table reached a certain size, which resulted in a drop of performance by about 200ms per request. This happened from one second to the other. Being able to troubleshoot and then fix issues like these is important in any enterprise system, but it is even more in a high load scenario.

Network Setup

We recommend hardware load balancers or putting a software loadbalancer onto its on server.

If you use a software load balancer (which is fine except for really extreme installations), it must be deployed on a machine of its own. Running a software load balancer on a cluster node is not supported. If a node unexpectedly got overwhelmed by a spike in load, a load balancer on that node would turn unresponsive. As a result, your whole cluster would be inaccessible even though the other nodes would be available. So using a different server is common practice and common sense.

Use separate network adapters for communication between servers.

The Confluence cluster nodes should have a separate physical network (i.e. separate NICs) for inter-server communication.

This is the best way of getting the cluster to run fast and reliably. Performance problems are likely to occur if you connect cluster nodes via a network that has lots of other data streaming through it.

The switch connecting the Confluence cluster nodes must not be a 'smart switch'.

Not supported. Smart switches are not covered by Atlassian Support for Confluence Clustered.

Do not use smart switches between cluster nodes. Many problems have been reported and attributed to smart switches. They have a

tendency to interrupt broadcast or multicast traffic, thus reliably killing a cluster after a certain amount of time has passed. This makes troubleshooting especially complex and tedious.

Cisco switches need additional configuration.

If the switch connecting the Confluence cluster nodes is a Cisco switch then it might need additional configuration to support Confluence clustering.

Please make sure you find out all the details about your switches before you start the deployment.

It is recommended that the database is on a different physical network from the Confluence server nodes.

Since you want to increase your capacity and performance for high loads, it is recommended to have your database on a different network. Please refer to the recommended topology diagram for more information.

Minimize the latency between the Confluence cluster nodes and the database.

Even though having the nodes and the database on the same physical network usually suffices, you should take the time to explicitly measure network latency, and make sure it is as close to zero as possible.

Prepare a network diagram.

To facilitate discussion and to ease planning, you should prepare a network diagram like this example of recommended network topology.

If you request support with Confluence Clustered, we may ask for your network diagram. We recommend that you create one similar to our example before you proceed with the installation.

You need network support staff available to troubleshoot cluster communication issues.

Setting up a cluster is not trivial. Even small problems in network design will be expanded in a clustered installation. (This is true of any kind of software.)

It is absolutely vital that you have dedicated network staff available to track down problems when they arise. A cluster will usually be used by thousands of users, and you don't want to keep them waiting because a network card breaks, or because someone made an undocumented change to the network and you don't have an expert around who can figure it out.

Staging Environment

You need a staging environment that is exactly the same as your production system.

You must be able to test drive any change to the cluster (installing upgrades, installing plugins) and to perform other tests (checking connectivity, debugging problems) on a staging cluster.

The staging environment must be:

- On the same OS, database, and Java version as your production environment.
- Clustered.

If you require support, we may for example ask you to turn off certain third-party plugins. If you can't do this in your production environment and you don't have a staging environment for troubleshooting, we may not be able to help you.

Getting a license for your staging environment

Only a technical contact for your commercial/academic license is able to create a Developer license

Atlassian supplies 'developer' licenses which can be used by existing commercial license holders who wish to deploy non-production installations of our software to use in QA/staging environments. Developer licenses are free of charge to commercial license holders and, like our commercial offerings, they include 12 months of updates starting from the date of purchase of the commercial license.

If you hold a commercial license, you can obtain a free developer license by performing the following:

1. Log in to your Atlassian account.
2. Under the "Licenses" heading, all of your licenses will be displayed. Click the plus sign next to a license to view its details.
3. Click the **'View Developer License'** link in the bottom right corner of the license detail panel, below your commercial license key.

RELATED TOPICS

Changing Datasources Manually in a Cluster

Cluster safety mechanism

Cluster Troubleshooting
Apache and Tomcat load balancing
Recommended network topology
Cluster Administration page
Cluster Panic triggers
Technical Overview of Clustering in Confluence
How do I supress cluster warning message in confluence?
Confluence Cluster Installation
Upgrading a Confluence Cluster
Confluence Clustering Overview
Viewing and Editing License Details

Confluence Security

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.

> This document is for system administrators looking to evaluate the security of the Confluence web application. This section only addresses overall application security. It does not address Confluence's internal security model (user/group management and content permissions). For information about user management, groups and permissions, please refer to the internal security overview.

On this page:

- Application Security Overview
  - Password Storage
  - Buffer Overflows
  - SQL Injection
  - Script Injection
  - Cross-Site Scripting
  - Transport Layer Security
  - Session Management
  - Plugin Security
  - Administrator Trust Model
  - Stack Traces
- Finding and Reporting a Security Vulnerability
- Publication of Confluence Security Advisories
- Severity Levels
- Our Patch Policy
- Published Security Advisories

Application Security Overview

Password Storage

When Confluence's internal user management is used, passwords are hashed through SHA1 before being stored in the database. There is no mechanism within Confluence to retrieve a user's password – when password recovery is performed, a new random password is generated and mailed to the user's registered address.

When external user management is enabled, password storage is delegated to the external system.

Buffer Overflows

Confluence is a 100% pure Java application with no native components. As such it is highly resistant to buffer overflow vulnerabilities – possible buffer overruns are limited to those that are bugs in the Java Runtime Environment itself.

SQL Injection

Confluence interacts with the database through the Hibernate Object-Relational mapper. Database queries are generated using standard APIs for parameter replacement rather than string concatenation. As such, Confluence is highly resistant to SQL injection attacks.

Script Injection

Confluence is a self-contained Java application and does not launch external processes. As such, it is highly resistant to script injection attacks.

Cross-Site Scripting

As a content-management system that allows user-generated content to be posted on the web, precautions have been taken within the application to prevent cross-site scripting attacks:

- The wiki markup language in Confluence does not support dangerous HTML markup
- Macros allowing the insertion of raw HTML are disabled by default
- HTML uploaded as a file attachment is served with a content-type requesting the file be downloaded, rather than being displayed inline
- Only system administrators can make HTML-level customisations of the application

When cross-site scripting vulnerabilities are found in the Confluence web application, we endeavour to fix them as quickly as possible.

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: Adding SSL for Secure Logins and Page Security

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 Standalone. If you are deploying Confluence in some other application server, you should ensure that it is not vulnerable to session hijacking.

Plugin Security

Administrators install third party plugins **at their own risk**. Plugins run in the same virtual machine as the Confluence server, and have access to the Java runtime environment, and the Confluence server API.

Administrators should always be aware of the source of the plugins they are installing, and whether they trust those plugins.

Administrator Trust Model

Confluence is written under the assumption that anyone given System Administrator privileges is trusted. System administrators are able, either directly or by installing plugins, to perform any operation that the Confluence application is capable of.

As with any application, you should not run Confluence as the root/Administrator user. If you want Confluence to listen on a privileged network port, you should set up port forwarding or proxying rather than run Confluence with additional privileges. The extra-careful may consider running Confluence inside a `chroot` jail.

Stack Traces

To help debug support cases and provide legendary support, Confluence provides stack traces through the web interface when an error occurs. These stack traces include information about what Confluence was doing at the time, and some information about your deployment server.

Only non-personal information is supplied such as operating system and version and Java version. With proper network security, this is not enough information to be considered dangerous. No usernames or passwords are included.

Finding and Reporting a Security Vulnerability

Atlassian's approach to reporting security vulnerabilities is detailed in How to Report a Security Issue.

Publication of Confluence Security Advisories

Atlassian's approach to releasing security advisories is detailed in Security Advisory Publishing Policy.

Severity Levels

Atlassian's approach to ranking security issues is detailed in Severity Levels for Security Issues.

Our Patch Policy

Atlassian's approach to releasing patches for security issues is detailed in Security Patch Policy.

Published Security Advisories

- Confluence Community Security Advisory 2006-01-19
- Confluence Security Advisory 2005-02-09
- Confluence Security Advisory 2005-12-05
- Confluence Security Advisory 2006-01-20
- Confluence Security Advisory 2006-01-23
- Confluence Security Advisory 2006-06-14
- Confluence Security Advisory 2007-07-26
- Confluence Security Advisory 2007-08-08
- Confluence Security Advisory 2007-11-19
- Confluence Security Advisory 2007-11-27
- Confluence Security Advisory 2007-12-14
- Confluence Security Advisory 2008-01-24
- Confluence Security Advisory 2008-03-06
- Confluence Security Advisory 2008-03-19
- Confluence Security Advisory 2008-05-21
- Confluence Security Advisory 2008-07-03
- Confluence Security Advisory 2008-09-08
- Confluence Security Advisory 2008-10-14
- Confluence Security Advisory 2008-12-03
- Confluence Security Advisory 2009-01-07
- Confluence Security Advisory 2009-02-18
- Confluence Security Advisory 2009-04-15
- Confluence Security Advisory 2009-06-01
- Confluence Security Advisory 2009-06-16
- Confluence Security Advisory 2009-08-20
- Confluence Security Advisory 2009-10-06
- Confluence Security Advisory 2009-12-08

Related Server Security Pages

Adding SSL for Secure Logins and Page Security

Click to see pages related to user and group permissions.

Hiding the People Directory

Revoking Space Permissions

Security Overview

Confluence Security Advisory 2006-01-23

Confluence Cookies

Assigning Space Permissions

Confluence Security

Edit in Word Link Macro

Space Permissions Overview

View File Macro

HTML Macro

How to Hide the Referrer

Confluence Community Security Advisory 2006-01-19

This security advisory is not endorsed by Atlassian - this is a public service advisory from a member of the confluence community. **Please** remember to backup any modified files, and use these instructions at your own risk. While this information is based on Confluence v2.1.2, it may have uses with older affected versions of Confluence.

The official security advisory is located at Confluence Security Advisory 2006-01-20

Problem

There is a possibility of XSS exploitation of the Full Name user profile field when displayed.

Solution

The problem was unescaped outputting of the fullname - wrapping the output in $generalUtil.htmlEncode() resolve it. The vast majority of the problem can be resolved by changing `/confluence/template/includes/macros.vm` in the distribution on the following lines:

- 180
- 186
- 200
- 340
- 893

I have attached the modified macros.vm file here which you can copy into your distribution.

Scope

There are other places which are still affected which Atlassian have been made aware of, a complete resolution should be provided by Atlassian in their own offical advisory.

I hope this helps some of you!

Confluence Security Advisory 2005-02-09

A flaw has been found in Confluence by which attackers can bypass Confluence security and change content on the site. Atlassian STRONGLY recommends that all Confluence customers apply the fix described below immediately, or upgrade to Confluence 1.3.3

Vulnerability

By crafting custom URLs, any person with the ability to browse Confluence can modify content on the site, bypassing security settings. This vulnerability does not allow users to view content they would not normally be able to view, or escalate their privileges in other ways.

This flaw affects all versions of Confluence prior to 1.3.3, including the 1.4-DR development releases.

Fix

This vulnerability is fixed in Confluence 1.3.3 and later. Customers who do not wish to migrate to 1.3.3 can fix this bug using the procedure below:

1. Edit the file confluence/WEB-INF/classes/xwork.xml
2. Find the following section near the top of the file (around line 34):

```
01.<interceptor-stack name="defaultStack">
02.          <interceptor-ref name="profiling">
03.                    <param name="location">Before defaultStack</param>
04.          </interceptor-ref>
05.          <interceptor-ref name="transaction"/>
06.          <interceptor-ref name="authentication"/>
07.          <interceptor-ref name="requestParameterHack"/>
08.          <interceptor-ref name="eventnotifier"/>
09.          <interceptor-ref name="autowire"/>
10.          <interceptor-ref name="params"/>
11.          <interceptor-ref name="servlet"/>
12.          <interceptor-ref name="pageAware"/>
13.          <interceptor-ref name="permissions"/>
14.          <interceptor-ref name="profiling">
15.                    <param name="location">After defaultStack</param>
16.          </interceptor-ref>
17.</interceptor-stack>
```

3. Locate the "autowire" and "params" entries:

```
1.                                   <interceptor-ref name="eventnotifier"/>
2.-->           <interceptor-ref name="autowire"/>      <--
3.-->           <interceptor-ref name="params"/>        <--
4.                                   <interceptor-ref name="servlet"/>
```

4. Swap the two lines around. The whole stack should now look like this:

```
<interceptor-stack name="defaultStack">
        <interceptor-ref name="profiling">
                <param name="location">Before defaultStack</param>
        </interceptor-ref>
        <interceptor-ref name="transaction"/>
        <interceptor-ref name="authentication"/>
        <interceptor-ref name="requestParameterHack"/>
        <interceptor-ref name="eventnotifier"/>
        <interceptor-ref name="params"/>
        <interceptor-ref name="autowire"/>
        <interceptor-ref name="servlet"/>
        <interceptor-ref name="pageAware"/>
        <interceptor-ref name="permissions"/>
        <interceptor-ref name="profiling">
                <param name="location">After defaultStack</param>
        </interceptor-ref>
</interceptor-stack>
```

5. Restart Confluence.

Confluence Security Advisory 2005-12-05

A flaw has been found in Confluence by which attackers to inject malicious HTML code into Confluence. Atlassian STRONGLY recommends that all Confluence customers apply the fix described below immediately, or upgrade to Confluence 2.0.2

Vulnerability

By entering HTML code into the Confluence search input fields, attackers can cause arbitrary scripting code to be executed by the user's browser in the security context of the Confluence instance.

This flaw affects all versions of Confluence between 1.4-DR releases and 2.0.1.

(Atlassian was not informed of the problem before it was published by third-party security researchers. You can read the third-party security advisory here: http://secunia.com/advisories/17833/. The vulnerability was originally reported here.)

Fix

This vulnerability is fixed in Confluence 2.0.2 and later. Customers who do not wish to migrate to 2.0.2 can fix this bug using the procedure below:

1. Edit the confluence/decorators/components/searchresults.vmd
2. Replace the following reference (around line 48):

```
$action.getText("search.result", [$start, $end, $total, $queryString])
```

with

```
$action.getText("search.result", [$start, $end, $total,
$generalUtil.escapeXml($queryString)]).
```

3. Edit the confluence/search/searchsite-results.vm.
4. Replace the following reference (around line 11):

```
Searched for <b>$action.searchQuery.queryString</b>
```

with

```
Searched for <b>$generalUtil.escapeXml($action.searchQuery.queryString)</b>
```

5. Restart Confluence.

Alternatively, you can download the patched source files from CONF-4825. If you are patching a 2.0.x installation, then use the files with the .2.0 suffix. If you are patching a 1.4.x installation, then use the files with the .1.4 suffix.

Confluence Security Advisory 2006-01-20

A flaw has been found in Confluence by which attackers to inject malicious HTML code into Confluence. Atlassian STRONGLY recommends that all Confluence customers apply the fix described below immediately, or upgrade to Confluence 2.1.3.

Vulnerability

By entering HTML/JavaScript code into the full name of a user's profile, attackers can cause arbitrary scripting code to be executed by the user's browser in the security context of the Confluence instance.

This flaw affects all versions of Confluence between 1.4-DR releases and 2.1.2.

This issue was initally reported by Ricardo Sueiras and a fix was quickly documented by Dan Hardiker at the Confluence Community Security Advisory 2006-01-19 page. Our thanks to them for bringing this to our attention.

There is an issue in JIRA at CONF-5233.

Fix

This vulnerability is fixed in Confluence 2.1.3 and later. Customers who do not wish to migrate to 2.1.3 can fix this bug using the procedure below:

Steps to fix:

1. Copy macros.vm to your confluence/template/includes folder
2. Restart Confluence

Note: If you are using version 1.4.4, please download and copy this file instead. You will need to rename it back to `macros.vm`.

If you are not using any of the above versions, you will need to replace wrap calls to display full names of users in $generalUtil.htmlEncode(). Alternatively, send us an email. We do however encourage you to use the latest stable point release regardless of the version you are using.

Confluence Security Advisory 2006-01-23

A flaw has been found in Confluence by which the unrestricted content of a space can be revealed in search results.

Vulnerability

By entering in a space key and blank query string into the Search macro, pages from the specified space will be displayed, without filtering on page and space permissions. This can allow unpermitted users to view the excerpts of pages they don't have access to.

This flaw is confirmed to affect all releases from 1.4 to 2.1.2.

More information is available at CONF-5189.

Fix

This vulnerability is fixed in Confluence 2.1.3 and later. We strongly suggest that customers upgrade to this release to fix the vulnerability.

Customers who are using 1.4.x and do not wish to upgrade can download a patched class from CONF-5198.

Confluence Security Advisory 2006-06-14

Vulnerability

By crafting a custom HTTP request, an attacker can delete or modify global permissions settings on a Confluence site.

This flaw affects all Confluence versions between 1.4 and 2.2.2. 2.2.3 and later are not vulnerable.

Fix

This issue has been fixed in Confluence 2.2.3. Patches are also available for all versions of Confluence betwen 1.4 and 2.2.2. For more information, please see this issue report.

Atlassian STRONGLY recommends that all customers either upgrade to Confluence 2.2.3, or apply the patch.

Confluence Security Advisory 2007-07-26

In this advisory:

- Users with view permission in a space can copy and save a page
- Space name and key are not validated nor escaped

Users with view permission in a space can copy and save a page

Vulnerability

A user who has only view permissions in a space can copy a page and then save it in the space. In this way, users can create a page in a space where they have only view permission.

This flaw affects only Confluence version 2.5.4.

Fix

This issue has been fixed in Confluence 2.5.5. A patch is also available for Confluence 2.5.4. For more information, including instructions on applying the patch, please see this issue report.

If you are using Confluence 2.5.4, Atlassian **strongly** recommends that you upgrade to Confluence 2.5.5 or apply the patch.

Space name and key are not validated nor escaped

Vulnerability

The input for space name and key is not validated properly - any characters are allowed. This makes a Confluence instance vulnerable to an XSS attack.

Fix

This issue has been fixed in Confluence 2.5.5. For more information, please see this issue report.

Atlassian recommends that you upgrade to Confluence 2.5.5.

Confluence Security Advisory 2007-08-08

In this advisory:

- Input in the RSS Feed Builder is not validated
- Input when editing Space Permissions is not validated
- Number of labels that can be added to a page is not restricted
- Input when editing navigation themes is not validated
- Viewing of space content alphabetically is not validated
- Input when editing Space Name is not validated
- Input when viewing attachments by file-type is not validated

Input in the RSS Feed Builder is not validated

Vulnerability

The input for the RSS Feed Builder is not required to be escaped. This can make a Confluence instance vulnerable to an XSS attack.

Fix

This issue has been fixed in Confluence 2.5.6. For more information, please see CONF-8993.

Atlassian recommends that you upgrade to Confluence 2.5.6.

Input when editing Space Permissions is not validated

Vulnerability

The 'Grant permission to' field on the 'Edit Space Permissions' screen is not validated. This can make a Confluence instance vulnerable to an XSS or DoS attack.

Fix

This issue has been fixed in Confluence 2.5.6. For more information, please see CONF-8980 and CONF-8979.

Atlassian recommends that you upgrade to Confluence 2.5.6.

Number of labels that can be added to a page is not restricted

Vulnerability

There is no restriction on the number of labels that can be added to a page at a time. This can make a Confluence instance vulnerable to a DoS attack.

Fix

This issue has been fixed in Confluence 2.5.6. For more information, please see CONF-8978.

Atlassian recommends that you upgrade to Confluence 2.5.6.

Input when editing navigation themes is not validated

Vulnerability

The 'Navigation Page' specified in the 'Left Navigation Theme' configuration is not validated. This can make a Confluence instance vulnerable to a XSS attack.

Fix

This issue has been fixed in Confluence 2.5.6. For more information, please see CONF-8956.

Atlassian recommends that you upgrade to Confluence 2.5.6.

Viewing of space content alphabetically is not validated

Vulnerability

When viewing space content by alphabetic character, the input is not validated as being alphabetic. This can make a Confluence instance vulnerable to an XSS attack.

Fix

This issue has been fixed in Confluence 2.5.6. For more information, please see CONF-8952.

Atlassian recommends that you upgrade to Confluence 2.5.6.

Input when editing Space Name is not validated

Vulnerability

The 'Name' field on the 'Edit Space Details' screen is not validated. This can make a Confluence instance vulnerable to an XSS attack.

Fix

This issue has been fixed in Confluence 2.5.6. For more information, please see CONF-8951.

Atlassian recommends that you upgrade to Confluence 2.5.6.

Input when viewing attachments by file-type is not validated

Vulnerability

The 'Filter By Extension' field on the 'List Space Attachments' screen is not validated. This can make a Confluence instance vulnerable to an XSS attack.

Fix

This issue has been fixed in Confluence 2.5.6. For more information, please see CONF-8950.

Atlassian recommends that you upgrade to Confluence 2.5.6.

Confluence Security Advisory 2007-11-19

In this advisory:

- DWR debug mode enabled
- XSS vulnerability in exception error page
- XSS vulnerability in the URL destination for the print icon
- XSS vulnerability in wiki markup for images

Atlassian recommends that you upgrade to Confluence 2.6.1 to fix the vulnerabilities described below.

DWR debug mode enabled

Vulnerability

Debug mode was enabled by default on Direct Web Remoting (DWR). This made it easy for a potential attacker to find information about available AJAX request handlers in Confluence.

Fix

This issue has been fixed in Confluence 2.6.1. If you do not wish to upgrade at this time, you can fix the problem by editing your

`<confluence install>/confluence/WEB-INF/web.xml` file. For more information, please see CONF-9718.

XSS vulnerability in exception error page

Vulnerability

The attributes and parameters were not escaped on the Confluence exception error page. This is a potential vulnerability to a cross-site scripting attack.

Fix

This issue has been fixed in Confluence 2.6.1. For more information, please see CONF-9704 and CONF-9560.

XSS vulnerability in the URL destination for the print icon

Vulnerability

The print icon on the HTTP 404 error page uses the path of the requested URL, which potentially contains malicious JavaScript. The 404 page did not correctly escape it. This is a potential vulnerability to a cross-site scripting attack.

Fix

This issue has been fixed in Confluence 2.6.1. A patch is supplied for customers with **Confluence version 2.6** who do not wish to upgrade at this time. For more information, please see CONF-9456.

XSS vulnerability in wiki markup for images

Vulnerability

When using image URLs in wiki markup, quotes were not correctly escaped. This is a potential vulnerability to a cross-site scripting attack.

Fix

This issue has been fixed in Confluence 2.6.1. For customers with **Confluence 2.6** who do not with to upgrade at this time, the new `atlassian-renderer` JAR should resolve this issue. For more information, please see CONF-9209.

Confluence Security Advisory 2007-11-27

In this advisory:

- XSS Type 2 Vulnerabilities in Macros and Wiki Markup
  - Severity
  - Risk Assessment
  - Risk Mitigation
  - Vulnerability
  - Fix

XSS Type 2 Vulnerabilities in Macros and Wiki Markup

Severity

Atlassian rates this vulnerability as **high**, according to the scale published in Confluence Security. The scale allows us to rank a vulnerability as critical, high, moderate or low.

Risk Assessment

We have identified and fixed some security flaws which may affect Confluence instances in a public environment. These flaws are XSS (cross-site scripting) vulnerabilities in some of Confluence's macros and Wiki Markup, which potentially allow a malicious user (hacker) to insert their own HTML tags or script into a Confluence page.

- The hacker might take advantage of this flaw to steal other users' session cookies or other credentials, by sending the credentials back to the hacker's own web server.
- The hacker's text and script might be displayed to other people viewing the Confluence page. This is potentially damaging to your company's reputation.

Atlassian recommends that you upgrade to Confluence 2.6.2 to fix the vulnerabilities described below.

You can read more about XSS attacks at cgisecurity, CERT and other places on the web.

Risk Mitigation

If you judge it necessary, you can disable public access (e.g. anonymous access and public signon) to your wiki until you have applied the necessary patch or upgrade. For even tighter control, you could restrict access to trusted groups only.

Vulnerability

The following macros are affected:

- {color}
- {panel}
- {section}
- {column}
- {code}

The Wiki Markup for inserting images (e.g. `!myImage.png!`) is also vulnerable to XSS exploitation.

Fix

The fix is to escape all user input, so that no user input is interpreted as HTML or CSS. In some cases we also perform stricter validation on the range of values a user can supply in an attribute.

These issues have been fixed in Confluence 2.6.2. For more information, please see CONF-9350.

> Our thanks to **Igor Minar**, who reported this issue to Atlassian. We fully support the reporting of vulnerabilities and we appreciate his working with us towards identifying and solving the problem.

Please let us know what you think of the format of this security advisory and the information we have provided.

Confluence Security Advisory 2007-12-14

In this advisory:

- XSS Vulnerability in Configure RSS Feed Action
  - Severity
  - Risk Assessment
  - Risk Mitigation
  - Vulnerability
  - Fix

XSS Vulnerability in Configure RSS Feed Action

Severity

Atlassian rates this vulnerability as **high**, according to the scale published in Confluence Security. The scale allows us to rank a vulnerability as critical, high, moderate or low.

Risk Assessment

We have identified and fixed a security flaw which may affect Confluence instances in a public environment. This flaw is an XSS (cross-site scripting) vulnerability in a Confluence action, which potentially allows a malicious user (hacker) to embed their own JavaScript into a Confluence page.

- The hacker might take advantage of this flaw to steal other users' session cookies or other credentials, by sending the credentials back to the hacker's own web server.
- The hacker's text and script might be displayed to other people viewing the Confluence page. This is potentially damaging to your company's reputation.

To fix the vulnerabilities described below, Atlassian recommends that you take one of the following steps:

- Upgrade to Confluence 2.7, or
- Download and install the patch for Confluence 2.5.8 or Confluence 2.6.2 from our JIRA site – see issue CONF-10164.

You can read more about XSS attacks at cgisecurity, CERT and other places on the web.

Risk Mitigation

If you judge it necessary, you can disable public access (e.g. anonymous access and public signon) to your wiki until you have applied the necessary patch or upgrade. For even tighter control, you could restrict access to trusted groups only.

Vulnerability

A hacker can inject their own JavaScript into the following Confluence action:

```
1.http://www.anyhost.com/confluence/dashboard/configurerssfeed.action
```

The above Confluence action is used to build an RSS feed based on your Confluence pages and news items. The action is invoked when a selects '**Feed Builder**' from your Confluence Dashboard. It can also be invoked by simply entering the URL into the browser address bar.

Fix

These issues have been fixed in **Confluence 2.7**, which you can download from the download centre.

A patch is available for **Confluence 2.5.8** and **Confluence 2.6.2**. For more information, please see CONF-10164.

> Our thanks to **jeff peichel**, who reported this issue to Atlassian. We fully support the reporting of vulnerabilities and we appreciate his working with us towards identifying and solving the problem.

Please let us know what you think of the format of this security advisory and the information we have provided.

Confluence Security Advisory 2008-01-24

In this advisory:

- XSS Vulnerability in Dashboard Action
  - Severity
  - Risk Assessment
  - Risk Mitigation
  - Vulnerability
  - Fix

XSS Vulnerability in Dashboard Action

Severity

Atlassian rates this vulnerability as **high**, according to the scale published in Confluence Security. The scale allows us to rank a vulnerability as critical, high, moderate or low.

Risk Assessment

We have identified and fixed a security flaw which may affect Confluence instances in a public environment. This flaw is an XSS (cross-site scripting) vulnerability in a Confluence action, which potentially allows a malicious user (hacker) to embed their own JavaScript into a Confluence page.

- The hacker might take advantage of this flaw to steal other users' session cookies or other credentials, by sending the credentials back to the hacker's own web server.
- The hacker's text and script might be displayed to other people viewing the Confluence page. This is potentially damaging to your company's reputation.

To fix the vulnerabilities described below, Atlassian recommends that you take one of the following steps:

- Upgrade to Confluence 2.7.1, or
- Download and install the patch for Confluence 2.6.2 or Confluence 2.7.0 from our JIRA site – see issue CONF-10289.

You can read more about XSS attacks at cgisecurity, CERT and other places on the web.

Risk Mitigation

If you judge it necessary, you can disable public access (e.g. anonymous access and public signon) to your wiki until you have applied the necessary patch or upgrade. For even tighter control, you could restrict access to trusted groups only.

Vulnerability

A hacker can inject their own JavaScript into the following Confluence action:

```
1.http://confluence-location/dashboard.action?spacesSelectedTab
```

The above Confluence action is used to determine which spaces are listed on a user's Dashboard. For example, the following URL requests a list of team spaces only:

```
1.http://confluence-location/dashboard.action?spacesSelectedTab=team
```

The action is invoked when a user selects one of the 'Spaces' tabs on the Dashboard, such as the '**Team**' tab. It can also be invoked by simply entering the URL into the browser address bar.

Fix

These issues have been fixed in **Confluence 2.7.1** (see the release notes), which you can download from the download centre.

A patch is available for **Confluence 2.6.2** and **Confluence 2.7.0**. For more information, please see CONF-10289.

> Our thanks to **Mary Johnson**, who reported this issue to Atlassian. We fully support the reporting of vulnerabilities and we appreciate her working with us towards identifying and solving the problem.

Please let us know what you think of the format of this security advisory and the information we have provided.

Confluence Security Advisory 2008-03-06

In this advisory:

- Users with View-Only Permission can Delete (Purge) Pages
  - Severity
  - Risk Assessment
  - Risk Mitigation
  - Vulnerability
  - Fix

Users with View-Only Permission can Delete (Purge) Pages

Severity

Atlassian rates this vulnerability as **high**, according to the scale published in Confluence Security. The scale allows us to rank a vulnerability as critical, high, moderate or low.

More explanation of the ranking we chose:

- You might rank this vulnerability as **critical**, because in most installations the vulnerability will allow anonymous users to delete information.
- We have chosen a ranking of **high**, because the vulnerability does not allow privilege escalation i.e. it doesn't allow users to gain administration privileges.

Risk Assessment

We have identified and fixed a security flaw which allowed users who have 'View' permission (or higher) on a space to purge (delete) any page in that space.

The following Confluence versions are vulnerable: All versions from **1.3 to 2.7.1** inclusive.

To fix the vulnerabilities described below, Atlassian recommends that you take one of the following steps:

- Upgrade to Confluence 2.7.2, or
- Download and install the patch for Confluence 2.6.x or Confluence 2.7.x from our JIRA site – see issue CONF-10807.

Risk Mitigation

If you judge it necessary, you can disable public access (e.g. anonymous access and public signon) to your wiki until you have applied the necessary patch or upgrade. For even tighter control, you could restrict access to trusted groups only.

If it is not immediately feasible to upgrade to Confluence 2.7.2 or apply a patch, we recommend an alternative strategy:

- As a temporary measure, you can block the URL which allows someone to purge (delete) a page. Please ask your website administrator to block the URL described below.
- The impact is that Space Administrators will not be able to purge individual pages or news items. However, Space Administrators can still use the 'Purge All' link to clear the entire contents of Trash.

Vulnerability

Description:
A user can use the following Confluence action to permanently delete (purge) any Confluence page, provided that the user has 'View' permission (or higher) in the space to which the page belongs:

```
1.http://confluence-location/pages/purgetrashitem.action?key=XXX&contentId=XXX
```

The above action is invoked when a space administrator clicks the 'Purge' link on the space's 'Trash' page next to a wiki page which has already been deleted.

The action can also be invoked by simply entering the URL into the browser address bar. In this way, it is possible for a user with 'View' permission (or higher) to remove a page via the 'Purge' action, even if the page has not been deleted.

Fix

These issues have been fixed in **Confluence 2.7.2** (see the release notes), which you can download from the download centre.

A patch is available for **Confluence 2.6.x**, **Confluence 2.7.0** and **Confluence 2.7.1**. For more information, please see CONF-10807.

Our thanks to **Neeraj Jhanji**, who reported this issue to Atlassian. We fully support the reporting of vulnerabilities and we appreciate his working with us towards identifying and solving the problem.

Confluence Security Advisory 2008-03-19

In this advisory:

- XSS Vulnerabilities in Various Confluence Actions
  - Severity
  - Risk Assessment
  - Risk Mitigation
  - Vulnerability
  - Fix

XSS Vulnerabilities in Various Confluence Actions

Severity

Atlassian rates these vulnerabilities as **high**, according to the scale published in Confluence Security. The scale allows us to rank a vulnerability as critical, high, moderate or low.

Risk Assessment

We have identified and fixed a number of security flaws which may affect Confluence instances in a public environment. The flaws are all XSS (cross-site scripting) vulnerabilities in various Confluence actions. Each vulnerability potentially allows a malicious user (hacker) to embed their own JavaScript into a Confluence page.

- The hacker might take advantage of the flaw to steal other users' session cookies or other credentials, by sending the credentials back to the hacker's own web server.
- The hacker's text and script might be displayed to other people viewing the Confluence page. This is potentially damaging to your company's reputation.

To fix the vulnerabilities described below, Atlassian recommends that you take one of the following steps:

- Upgrade to Confluence 2.7.3, or
- Download and install the patches for Confluence **2.6.x** from our JIRA site — refer to the list of issues below.

You can read more about XSS attacks at cgisecurity, CERT and other places on the web.

Risk Mitigation

If you judge it necessary, you can disable public access (e.g. anonymous access and public signon) to your wiki until you have applied the necessary patch or upgrade. For even tighter control, you could restrict access to trusted groups only.

Vulnerability

A hacker can inject their own JavaScript into the Confluence actions listed in the table below. Each of the actions is invoked when a user performs a specific function in Confluence, such as clicking a link or a button. The actions can also be invoked by simply entering the URL into the browser address bar.

For more details please refer to the related JIRA issue, also shown in the table below.

Confluence Actions	Affected Confluence Versions	More Details	Reporter (If Not Atlassian)
Create, edit or copy a page or news item	From **2.2 to 2.7.2** inclusive	CONF-11027	
Add a comment	From **2.2 to 2.7.2** inclusive	CONF-11027	
Create a space	From **2.2 to 2.7.2** inclusive	CONF-11042	Wyatt Crossin
Sign up for an account	From **2.2 to 2.7.2** inclusive	CONF-11005	
Choose a page (page picker)	From **2.2 to 2.7.2** inclusive	CONF-11137	
View a user	From **2.2 to 2.7.2** inclusive	CONF-11002	
Insert an image or link	From **2.2 to 2.7.2** inclusive	CONF-11141	

Choose a user or group (user picker and group picker)	From **2.2 to 2.7.2** inclusive	CONF-11040	Jean Marois
Add a user to favourites	From **2.0 to 2.7.2** inclusive	CONF-11026	
HTTP 500 error page	From **1.3 to 2.7.2** inclusive	CONF-11019	
Add bookmark	All Confluence instances that have the Social Bookmarking plugin. Note that the plugin is bundled with Confluence since version **2.6**, so Confluence 2.6.x and 2.7.x are vulnerable even if you don't use social bookmarking. Patches are supplied for Confluence 2.6.x and 2.7.x.	CONF-11153	

Fix

These issues have been fixed in **Confluence 2.7.3** (see the release notes), which you can download from the download centre.

Patches are available for **Confluence 2.6.x**. For more information, please refer to the specific JIRA issues shown in the table of vulnerabilities above.

> Our thanks to the people who reported some of the vulnerabilities listed above. We fully support the reporting of vulnerabilities and we appreciate their working with us towards identifying and solving the problem.

Confluence Security Advisory 2008-05-21

In this advisory:

- Users can Move Attachments to Any Page Regardless of Permissions
  - Severity
  - Risk Assessment
  - Risk Mitigation
  - Vulnerability
  - Fix
- XSS Vulnerability in Page Information View
  - Severity
  - Risk Assessment
  - Risk Mitigation
  - Vulnerability
  - Fix

Users can Move Attachments to Any Page Regardless of Permissions

Severity

Atlassian rates this vulnerability as **high**, according to the scale published in Confluence Security. The scale allows us to rank a vulnerability as critical, high, moderate or low.

Risk Assessment

We have identified and fixed a security flaw which allows users who have 'Create Page' permission in a space to move an attachment from a page in that space to any other page in the Confluence site, regardless of the user's permissions in the destination space.

The following Confluence versions are vulnerable: All versions from **1.0 to 2.8.0**.

Risk Mitigation

This security flaw grants extra powers only to users who already have 'Create Page' permissions in one of the spaces on the Confluence site. In most installations, this will be a trusted group of users.

If your Confluence instance allows a less trusted group of users to create and edit pages in one space, while restricting access to other spaces, you may judge it necessary to disable public access (e.g. anonymous access and public signon) to your wiki until you have applied the necessary patch or upgrade. For even tighter control, you could restrict access to trusted groups only.

Vulnerability

Any user who has 'Create Page' permission in a Confluence space can move an attachment from a page in that space to any other page in the Confluence site, regardless of the user's permissions in the destination space.

Note: If a user has permission to create a space, they will also have 'Create Page' permission in any space they create, including a personal space. Such users could upload an attachment onto the space they have created and then move the attachment to any page in the Confluence site.

Fix

This issue has been fixed in Confluence 2.8.1 (see the release notes), which you can download from the download centre.

Alternatively, you can download and install the patch for Confluence 2.7.x or Confluence 2.8.0 from our JIRA site – see issue CONF-11452.

> Our thanks to **Stafford Vaughan** from CustomWare, who reported this issue to Atlassian. We fully support the reporting of vulnerabilities and we appreciate it when people work with us towards identifying and solving a problem.

XSS Vulnerability in Page Information View

Severity

Atlassian rates this vulnerability as **high**, according to the scale published in Confluence Security. The scale allows us to rank a vulnerability as critical, high, moderate or low.

Risk Assessment

We have identified and fixed a security flaw which may affect Confluence instances in a public environment. This flaw is an XSS (cross-site scripting) vulnerability in a Confluence action, which potentially allows a malicious user (hacker) to embed their own JavaScript into a Confluence page.

- The hacker might take advantage of this flaw to steal other users' session cookies or other credentials, by sending the credentials back to the hacker's own web server.
- The hacker's text and script might be displayed to other people viewing the Confluence page. This is potentially damaging to your company's reputation.

You can read more about XSS attacks at cgisecurity, CERT and other places on the web.

The following Confluence versions are vulnerable: All versions from **1.3 to 2.8.0** inclusive.

Risk Mitigation

If you judge it necessary, you can hide referrers on page information views by disabling this functionality.

Vulnerability

A hacker can inject their own JavaScript into the referrer URLs which are displayed on the 'Info' view of a wiki page. The rogue JavaScript will be executed when a user opens the 'Info' view.

Fix

This issue has been fixed in Confluence 2.8.1 (see the release notes), which you can download from the download centre.

Alternatively, you can download and install the patch for Confluence 2.7.x or Confluence 2.8.0 from our JIRA site – see issue CONF-11524.

Confluence Security Advisory 2008-07-03

In this advisory:

- XSS Vulnerability in Various Confluence Actions
  - Severity
  - Risk Assessment
  - Risk Mitigation
  - Vulnerability
  - Fix

XSS Vulnerability in Various Confluence Actions

Severity

Atlassian rates these vulnerabilities as **high**, according to the scale published in Confluence Security. The scale allows us to rank a vulnerability as critical, high, moderate or low.

Risk Assessment

We have identified and fixed a number of security flaws which may affect Confluence instances in a public environment. The flaws are all XSS (cross-site scripting) vulnerabilities in various Confluence actions. Each vulnerability potentially allows a malicious user (hacker) to embed their own JavaScript into a Confluence page.

- The hacker might take advantage of the flaw to steal other users' session cookies or other credentials, by sending the credentials back to the hacker's own web server.
- The hacker's text and script might be displayed to other people viewing the Confluence page. This is potentially damaging to your

company's reputation.

You can read more about XSS attacks at cgisecurity, CERT and other places on the web.

Risk Mitigation

If you judge it necessary, you can disable public access (e.g. anonymous access and public signon) to your wiki until you have applied the necessary patch or upgrade. For even tighter control, you could restrict access to trusted groups only.

Vulnerability

A hacker can inject their own JavaScript into the Confluence actions listed in the table below. Each of the actions is invoked when a user performs a specific function in Confluence, such as clicking a link or a button. The actions can also be invoked by simply entering the URL into the browser address bar. The rogue JavaScript will be executed when a user invokes the URL.

For more details please refer to the related JIRA issue, also shown in the table below.

Confluence Actions	Affected Confluence Versions	More Details	Reporter (If Not Atlassian)
Create, edit or copy a page or news item	2.8.0 and 2.8.1	CONF-11985	James Rinker
Page picker and space picker	2.2.0 to 2.8.1 inclusive	CONF-11137	

Fix

These issues have been fixed in Confluence 2.8.2 (see the release notes), which you can download from the download centre.

Alternatively, you can download and install the patches provided on our JIRA site. For more information, please refer to the specific JIRA issues shown in the table of vulnerabilities above.

> Our thanks to **James Rinker** who reported some of the vulnerabilities listed above. We fully support the reporting of vulnerabilities and we appreciate his working with us towards identifying and solving the problem.

Confluence Security Advisory 2008-09-08

In this advisory:

- XSS Bug: Usernames Not HTML-Encoded in All Places
  - Severity
  - Risk Assessment
  - Risk Mitigation
  - Vulnerability
  - Fix
- Inherited Page Restrictions Are Not Applied After 2.9 Upgrade
  - Severity
  - Risk Assessment
  - Risk Mitigation
  - Vulnerability
  - Fix
- Access Vulnerability in View Wiki Markup Function
  - Severity
  - Risk Assessment
  - Risk Mitigation
  - Vulnerability
  - Fix
- Access Vulnerability in Copy Page Function
  - Severity
  - Risk Assessment
  - Risk Mitigation
  - Vulnerability
  - Fix
- Access Vulnerability in Diff Page Function
  - Severity
  - Risk Assessment
  - Risk Mitigation
  - Vulnerability
  - Fix

XSS Bug: Usernames Not HTML-Encoded in All Places

Severity

Atlassian rates this vulnerability as **HIGH**, according to the scale published in Confluence Security. The scale allows us to rank a vulnerability as critical, high, moderate or low.

Risk Assessment

We have identified and fixed a security flaw which allowed certain users to circumvent Confluence's security measures, by including HTML markup in their own username. This could allow a malicious user to execute Javascript on another user's authenticated session.

The following Confluence versions are vulnerable: All versions from **1.0 to 2.9**.

Risk Mitigation

If the user specified a username that included HTML markup (which could include Javascript), in some places Confluence would not correctly escape this source before displaying it. This could result in Javascript being executed in another user's authenticated session. To address the issue, you should update your Confluence instance as soon as possible (or follow the patch instructions on the issue).

Vulnerability

This is a classic Cross-Site Scripting issue where usernames could include malicious Javascript.

Fix

This issue has been fixed in Confluence 2.9.1 (see the release notes), which you can download from the download centre.

For more information, see issue CONF-7615 which has instructions on how to patch the affected velocity template.

Inherited Page Restrictions Are Not Applied After 2.9 Upgrade

Severity

Atlassian rates this vulnerability as **HIGH**, according to the scale published in Confluence Security. The scale allows us to rank a vulnerability as critical, high, moderate or low.

Risk Assessment

We have identified and fixed a security flaw that caused any content permission inherited by a page to be lost during the upgrade process to Confluence 2.9.

The following Confluence versions are vulnerable: Version **2.9**; specifically instances of Confluence that were *upgraded to version 2.9* (from an earlier version) only.

Risk Mitigation

This issue can be resolved by following the steps under **Fix**, or upgrading to Confluence 2.9.1. If this cannot be done immediately, it may be prudent to manually apply restrictions to each page that is normally protected by inherited restrictions (that is, all child pages residing under a restricted page). Enacting the fix is trivial and should take around ten minutes for a typical Confluence instance.

Vulnerability

If you had given a parent page restrictions prior to the 2.9 upgrade, then any child pages that should be inheriting these restrictions are no longer restricted. This potentially renders these child pages viewable and editable by Confluence users who should not have these rights. However you should note that any space level restrictions are still respected so these affected pages are only opened as far as the space level security allows for your site. Note for individual pages where you have manually set the permissions, those pages are not at risk — just the pages underneath them using inherited permissions.

Fix

This issue has been fixed in Confluence 2.9.1 (see the release notes), which you can download from the download centre.

Alternatively, you can apply the manual fix, which involves a simple series of actions in the Confluence administration screens.

For more information see issue CONF-12911.

Access Vulnerability in View Wiki Markup Function

Severity

Atlassian rates this vulnerability as **HIGH**, according to the scale published in Confluence Security. The scale allows us to rank a vulnerability as critical, high, moderate or low.

Risk Assessment

We have identified and fixed a security flaw which allows users who don't have the correct '**View Page**' permission in a space to view the Wiki Markup source of the page content.

The following Confluence versions are vulnerable: Version **2.9** only.

Risk Mitigation

If a user knows the URL to view the source of a page they will be able to bypass Confluence's security checks. This will allow the user to view the contents of a page they aren't meant to see.
To prevent unauthorised access, you may want to use your web server to reject all requests to URLs containing this string: `/pages/viewpagesrc.action`. You may judge it necessary to disable public access.

Vulnerability

If a user knows the ID of a page that they do not have **'View Page'** permission for they can use the view source URL to view the Wiki Markup of a page. This will allow them to copy and paste the contents of the page to another location, or simply read the markup and deduce its final content.

Note: the user will need to know the page ID of a page. Confluence will not provide any links to the restricted page through a search or other navigation.

Fix

This issue has been fixed in Confluence 2.9.1 (see the release notes), which you can download from the download centre.

For more information see issue CONF-12845.

Access Vulnerability in Copy Page Function

Severity

Atlassian rates this vulnerability as **HIGH**, according to the scale published in Confluence Security. The scale allows us to rank a vulnerability as critical, high, moderate or low.

Risk Assessment

We have identified and fixed a security flaw which allows users who don't have the correct 'View Page' permission in a space to copy a page and therefore see its content.

The following Confluence versions are vulnerable: All versions from **1.0 to 2.9**.

Risk Mitigation

If a user knows the URL to copy a page they will be able to bypass Confluence's security checks. This will allow the user to view the contents of a page they aren't meant to see.
To prevent unauthorised access, you may want to use your web server to reject all requests to URLs containing this string: `/pages/copypage.action`. You may judge it necessary to disable public access.

Vulnerability

If a user knows the ID of a page they do not have permissions for, they can use the copy page URL to copy the page to a space where they do have permission. This will allow them to create a new page based on the content of a page they aren't meant to see.

Fix

This issue has been fixed in Confluence 2.9.1 (see the release notes), which you can download from the download centre.

Alternatively, you can download and install the patch for Confluence 2.7.3 or 2.8.2 from our JIRA site – see issue CONF-12859.

Instruction on installing the patch can be found here.

Access Vulnerability in Diff Page Function

Severity

Atlassian rates this vulnerability as **HIGH**, according to the scale published in Confluence Security. The scale allows us to rank a vulnerability as critical, high, moderate or low.

Risk Assessment

We have identified and fixed a security flaw which allows users who don't have the correct '**View Page**' permission in a space to create a diff of a page (a comparison of its contents with another page) and therefore see its content.

The following Confluence versions are vulnerable: All versions from **1.0 to 2.9**.

Risk Mitigation

If a user knows the URL to perform a diff of a page they will be able to bypass Confluence's security checks. This will allow the user to view the contents of a page they aren't meant to see.
To prevent unauthorised access, you may want to use your web server to reject all requests to URLs containing this string: `/pages/diffpages.action`. You may judge it necessary to disable public access.

Vulnerability

If a user knows the ID of a page they do not have permissions for, they can use the 'Diff Page' URL to compare the contents of that page with one where they do. This will allow them to deduce the contents of a page they don't have access to.

Fix

This issue has been fixed in Confluence 2.9.1 (see the release notes), which you can download from the download centre.

Alternatively, you can download and install the patch for Confluence 2.7.3 or 2.8.2 from our JIRA site – see issue CONF-12860.

Instruction on installing the patch can be found here.

> Our thanks to **Neeraj Jhanji** from Atlassian Partner ImaHima, who reported the copy and diff page issues to Atlassian. We fully support the reporting of vulnerabilities and we appreciate it when people work with us towards identifying and solving a problem.

Confluence Security Advisory 2008-10-14

In this advisory:

- Parameter Injection Vulnerability in Confluence
  - Severity
  - Risk Assessment
  - Risk Mitigation
  - Vulnerability
  - Fix
- XSS Vulnerability in Various Confluence Actions and Plugins
  - Severity
  - Risk Assessment
  - Risk Mitigation
  - Vulnerability
  - Fix
- Privilege Escalation Vulnerability in Confluence Watches
  - Severity
  - Risk Assessment
  - Risk Mitigation
  - Vulnerability
  - Fix
- Privilege Escalation Vulnerability in Confluence Favourites
  - Severity
  - Risk Assessment
  - Risk Mitigation
  - Vulnerability
  - Fix

Parameter Injection Vulnerability in Confluence

Severity

Atlassian rates this vulnerability as **critical**, according to the scale published in Confluence Security. The scale allows us to rank a vulnerability as critical, high, moderate or low.

Risk Assessment

We have identified and fixed a flaw which would allow a malicious user (hacker) to inject their own values into a Confluence request by adding parameters to the URL string. This would allow a hacker to bypass Confluence's security checks and perform actions that they are not authorised to perform.

Risk Mitigation

To address the issue, you should upgrade Confluence as soon as possible or follow the patch instructions below. If you judge it necessary,

you can block all untrusted IP addresses from accessing Confluence.

Vulnerability

A hacker can design a URL string containing parameters which perform specific actions on the Confluence server, bypassing Confluence's security checks. This is because Confluence does not adequately sanitise user input before applying it as an action on the server.

Exploiting this issue could allow an attacker to access or modify data and compromise the Confluence application.

The following Confluence versions are vulnerable: All versions from **1.3 to 2.9.1**.

Fix

This issue has been fixed in Confluence 2.9.2 (see the release notes), which you can download from the download centre.

If you do not wish to upgrade to Confluence 2.9.2, a patch is available that will work with any affected version of Confluence. You can download and install the patch from on our JIRA site. For more information, please refer to CONF-13092.

XSS Vulnerability in Various Confluence Actions and Plugins

Severity

Atlassian rates these vulnerabilities as **high**, according to the scale published in Confluence Security. The scale allows us to rank a vulnerability as critical, high, moderate or low.

Risk Assessment

We have identified and fixed a number of security flaws which may affect Confluence instances in a public environment. The flaws are all XSS (cross-site scripting) vulnerabilities in various Confluence actions. Each vulnerability potentially allows a malicious user (hacker) to embed their own JavaScript into a Confluence page.

- The hacker might take advantage of the flaw to steal other users' session cookies or other credentials, by sending the credentials back to the hacker's own web server.
- The hacker's text and script might be displayed to other people viewing the Confluence page. This is potentially damaging to your company's reputation.

You can read more about XSS attacks at cgisecurity, CERT and other places on the web.

Risk Mitigation

If you judge it necessary, you can disable public access (e.g. anonymous access and public signon) to your wiki until you have applied the necessary patch or upgrade. For even tighter control, you could restrict access to trusted groups.

Vulnerability

A hacker can inject their own JavaScript into the Confluence actions listed in the table below. Each of the actions is invoked when a user performs a specific function in Confluence, such as clicking a link or a button. The actions can also be invoked by simply entering the URL into the browser address bar. The rogue JavaScript will be executed when a user invokes the URL.

For more details please refer to the related JIRA issue, also shown in the table below.

Confluence Actions	Affected Confluence Versions	More Details	Reporter (If Not Atlassian)
View children via the Pagetree plugin (bundled with Confluence)	2.8.0 to 2.9.1 inclusive	CONF-13043	Thomas Jaehnel
Update bookmark via the Social Bookmarking plugin (bundled with Confluence)	2.6.0 to 2.9.1 inclusive	CONF-13041	Thomas Jaehnel
Build RSS feed	2.0 to 2.9.1 inclusive	CONF-13042	Thomas Jaehnel
Search via Search macro	All versions from 1.0 to 2.9.1 inclusive	CONF-13040	Thomas Jaehnel
Search	All versions from 1.0 to 2.9.1 inclusive	CONF-12944	

Fix

These issues have been fixed in Confluence 2.9.2 (see the release notes), which you can download from the download centre.

If you do not wish to upgrade to Confluence 2.9.2, you can download and install the patches provided on our JIRA site. For more information, please refer to the specific JIRA issues shown in the table of vulnerabilities above.

Our thanks to **Thomas Jaehnel** of **OPTIMAbit**, who reported most of the XSS vulnerabilities listed above. We fully support the reporting of vulnerabilities and we appreciate it when people work with us to identify and solve the problem.

Privilege Escalation Vulnerability in Confluence Watches

Severity

Atlassian rates this vulnerability as **high**, according to the scale published in Confluence Security. The scale allows us to rank a vulnerability as critical, high, moderate or low.

Risk Assessment

We have identified and fixed a flaw which would allow an unauthorised user to add a Confluence page to the list of pages they are watching, even if the user does not have permission to view that page. Under some circumstances, the unauthorised user may thus have access to information they are not authorised to see.

Risk Mitigation

This flaw does not allow the unauthorised user to update the page, but it may give the user access to information that they do not have permission to see.

Vulnerability

An unauthorised user can manipulate the HTTP request, so that it adds a watch to a page which the user does not have permission to view. The page then appears in the user's list of watched pages, displaying the page title and the corresponding space name. In this way, the user can bypass Confluence's permission checks and gain access to information they are not authorised to see.

The following Confluence versions are vulnerable: All versions from **1.0 to 2.9.1**.

Fix

This issue has been fixed in Confluence 2.9.2 (see the release notes), which you can download from the download centre.

If you do not wish to upgrade to Confluence 2.9.2, you can download and install the patches provided on our JIRA site. For more information, please refer to CONF-13039.

Our thanks to **Thomas Jaehnel** of **OPTIMAbit**, who reported the vulnerability listed above. We fully support the reporting of vulnerabilities and we appreciate it when people work with us to identify and solve the problem.

Privilege Escalation Vulnerability in Confluence Favourites

Severity

Atlassian rates this vulnerability as **moderate**, according to the scale published in Confluence Security. The scale allows us to rank a vulnerability as critical, high, moderate or low.

Risk Assessment

We have identified and fixed a flaw which would allow an unauthorised user to add a Confluence page to their list of favourites, even if the user does not have permission to view that page. Under some circumstances, the unauthorised user may thus have access to information they are not authorised to see.

Risk Mitigation

This flaw does not allow the unauthorised user to update the page, and it gives the user only very limited access to the information they do not have permission to see.

Vulnerability

An unauthorised user can manipulate the HTTP request, so that it marks as 'favourite' a page which the user does not have permission to view. The page is then added to the number of favourites for the user. The user cannot see the page title or content, but can see that the favourite count has been incremented.

The following Confluence versions are vulnerable: All versions from **1.0 to 2.9.1**.

Fix

This issue has been fixed in Confluence 2.9.2 (see the release notes), which you can download from the download centre.

If you do not wish to upgrade to Confluence 2.9.2, you can download and install the patches provided on our JIRA site. For more information,

please refer to CONF-13044.

> Our thanks to **Thomas Jaehnel** of **OPTIMAbit**, who reported the vulnerability listed above. We fully support the reporting of vulnerabilities and we appreciate it when people work with us to identify and solve the problem.

Confluence Security Advisory 2008-12-03

In this advisory:

- XSS Vulnerability in Various Confluence Actions
  - Severity
  - Risk Assessment
  - Risk Mitigation
  - Vulnerability
  - Fix
- Users can View a List of All Attachments by Supplying an Edited URL
  - Severity
  - Risk Assessment
  - Risk Mitigation
  - Vulnerability
  - Fix

XSS Vulnerability in Various Confluence Actions

Severity

Atlassian rates these vulnerabilities as **high**, according to the scale published in Confluence Security. The scale allows us to rank a vulnerability as critical, high, moderate or low.

Risk Assessment

We have identified and fixed a number of security flaws which may affect Confluence instances in a public environment. The flaws are all XSS (cross-site scripting) vulnerabilities in various Confluence actions. Each vulnerability potentially allows a malicious user (hacker) to embed their own JavaScript into a Confluence page.

- The hacker might take advantage of the flaw to steal other users' session cookies or other credentials, by sending the credentials back to the hacker's own web server.
- The hacker's text and script might be displayed to other people viewing the Confluence page. This is potentially damaging to your company's reputation.

You can read more about XSS attacks at cgisecurity, CERT and other places on the web.

Risk Mitigation

If you judge it necessary, you can disable public access (e.g. anonymous access and public signon) to your wiki until you have applied the necessary patch or upgrade. For even tighter control, you could restrict access to trusted groups.

Vulnerability

A hacker can inject their own JavaScript into various Confluence URLs — see the table below for the affected functional areas. A URL may be invoked when a user performs a specific function in Confluence, such as clicking a link or a button. The URL can also be invoked by simply entering it into the browser address bar. If rogue JavaScript is injected into such a URL, the JavaScript will be executed when a user invokes the URL.

For more details please refer to the related JIRA issue, also shown in the table below.

Affected Confluence Functionality	Affected Confluence Versions	Fix Availability	More Details	Reporter (If Not Atlassian)
Handling of error messages. (Vulnerability in the DWR code library used by Confluence.)	2.7.3 to 2.9.2 inclusive	2.9.2 and 2.10	CONF-11808	Bjoern Froebe
Attachments macro.	2.8 to 2.9.2 inclusive	2.8.2, 2.9.2 and 2.10**	CONF-13713	
Uploading of attachments.	2.6 to 2.9.2 inclusive	2.8.2, 2.9.2 and 2.10	CONF-13717	
Inserting images as thumbnails.	2.8 to 2.9.2 inclusive	2.8.2, 2.9.2 and 2.10	CONF-13625	
Log events listed in the Confluence 500 error page.	2.9 to 2.9.2 inclusive	2.10 only	CONF-13584	

Wiki Markup link rendering.	2.7 to 2.9.2 inclusive	2.7.x, 2.8.x, 2.9.x, 2.10	CONF-13451	

** The patch for CONF-13717 also addresses the bug in CONF-13736.*
*** To fix this issue, please upgrade your Attachments plugin to the latest version. This plugin is available for Confluence 2.8.2, 2.9.2 and 2.10, via the Confluence Plugin Repository.*

Fix

These issues have been fixed in Confluence 2.10 (see the release notes), which you can download from the download centre.

If you do not wish to upgrade to Confluence 2.10, you can download and install the patches provided on our JIRA site. You will need to upgrade to the latest point release for the major version of Confluence that you are running (e.g. if you are running Confluence 2.8, you will need to upgrade to version 2.8.2) and then apply the patches. For more information, please refer to the specific JIRA issues shown in the table of vulnerabilities above.

Please note that one of the issues can **only be fixed by upgrading to Confluence 2.10**. Please see the table above for details.

> Our thanks to **Bjoern Froebe**, who reported one of the XSS vulnerabilities listed above. We fully support the reporting of vulnerabilities and we appreciate it when people work with us to identify and solve the problem.

Users can View a List of All Attachments by Supplying an Edited URL

Severity

Atlassian rates this vulnerability as **medium**, according to the scale published in Confluence Security. The scale allows us to rank a vulnerability as critical, high, moderate or low.

Risk Assessment

We have identified and fixed a security flaw which allows a user to view the list of all attachments for all pages in a Confluence instance, regardless of space-level or page-level permissions.

While the user cannot open the files, a range of metadata is available for viewing, including file name, the page that the file is attached to, the creator, and the creation and last-modified date of the attachment.

Risk Mitigation

If you judge it necessary, you can disable anonymous access to your wiki until you have applied the necessary patch or upgrade.

Vulnerability

If a user removes the space key from the URL while viewing attachments for a space, Confluence will display the full list of all attachments for all spaces. For more details, please refer to CONF-13874.

Fix

These issues have been fixed in Confluence 2.10 (see the release notes), which you can download from the download centre.

If you do not wish to upgrade to Confluence 2.10, you can download and install the patches provided in the JIRA issue, CONF-13874. You will need to upgrade to the latest point release for the major version of Confluence that you are running (e.g. if you are running Confluence 2.8, you will need to upgrade to version 2.8.2) and then apply the patch.

> Our thanks to **Matthew Goonan**, who reported this vulnerability. We fully support the reporting of vulnerabilities and we appreciate it when people work with us to identify and solve the problem.

Confluence Security Advisory 2009-01-07

In this advisory:

- Content Overwrite Vulnerability in the Office Connector Plugin
  - Severity
  - Risk Assessment
  - Risk Mitigation
  - Vulnerability
  - Fix

Content Overwrite Vulnerability in the Office Connector Plugin

Severity

Atlassian rates this vulnerability as **high**, according to the scale published in Confluence Security. The scale allows us to rank a vulnerability as critical, high, moderate or low.

Risk Assessment

We have identified a risk that makes it possible for users with read-only access to a Confluence wiki space to modify its contents via the document import feature of the Office Connector plugin. This issue, however, does not expose restricted content on a Confluence wiki space to unauthorised users.

Risk Mitigation

Please see the 'Fix' section below. If you cannot apply the fix immediately, you can consider taking one or more of the following steps:

- Disable the whole Office Connector plugin, as explained here.
- If you judge it necessary, you can disable public access (e.g. anonymous access and public signon) to your wiki until you have applied the necessary patch or upgrade.
- For even tighter control, you could restrict access to trusted groups.

Vulnerability

The Office Connector plugin was first bundled in Confluence version 2.10.0. Hence, this vulnerability affects Confluence **2.10.0** where the Office Connector Plugin is enabled. Additionally, this plugin is compatible with all versions of Confluence **from 2.3.0** onwards. Hence, if you have installed the plugin, this vulnerability will affect your Confluence instance.

Fix

Please download and install the latest version of the Office Connector plugin via the Confluence Plugin Repository (instructions here). If you wish to install this plugin manually, you can download it from here.

Alternatively, install or upgrade to Confluence version 2.10.1. (See the release notes.) The Confluence 2.10.1 installation files can be downloaded from the download centre.

For more information, please refer to CONF-14014.

> Our thanks to **Justin Wong**, who reported this vulnerability. We fully support the reporting of vulnerabilities and we appreciate it when people work with us to identify and solve the problem.

Confluence Security Advisory 2009-02-18

In this advisory:

- HTTP Header Injection Flaw
  - Severity
  - Risk Assessment
  - Risk Mitigation
  - Vulnerability
  - Fix

HTTP Header Injection Flaw

Severity

Atlassian rates this vulnerability as **high**, according to the scale published in Confluence Security. The scale allows us to rank a vulnerability as critical, high, moderate or low.

An Advanced Warning of this Security Advisory published last week on http://forums.atlassian.com, stated the severity of this vulnerability as **critical**. After further assessing the likelihood of attack, however, we have amended this to **high**.

Risk Assessment

We have identified and fixed a security flaw which may affect Confluence instances in a public environment. This flaw is an HTTP header injection vulnerability in the Seraph web framework that is used by Confluence. This potentially allows a malicious user (attacker) to modify the HTTP response to insert malicious code. An attacker could present a modified URL to users (e.g. disguised in an email message). If any user clicks the URL, the malicious code would be executed in the user's session.

- The attacker may take advantage of this flaw to steal other users' session cookies or other credentials, by sending the credentials back to the attacker's own web server.
- The attacker could also gain control over the underlying system, based on the privileges of the user whose session cookie has been stolen.
- The attacker could redirect the user to undesirable web sites. This is potentially damaging to your company's reputation.

Atlassian recommends that you upgrade to Confluence 2.10.2 to fix the vulnerabilities described below.

Risk Mitigation

We strongly recommend either patching or upgrading your Confluence installation to fix this vulnerability. Please see the 'Fix' section below.

Alternatively, you may consider taking the following step, although the time required to fix this vulnerability and the extent of its effectiveness will depend on your application server running Confluence and its configuration:

- Consult the vendor of your application server to see whether your application server is immune to header injection vulnerabilities or has configuration options to prevent such attacks. For example, the Coyote (HTTP) connector in Tomcat version 5.5 and later is immune to header injection attacks, as acknowledged in this reference.
  Technical note: In your application server, header injection vulnerabilities can be mitigated if the setHeader(), addHeader(), and sendRedirect() methods in the HttpServletResponse class have their parameters properly checked for header termination characters.
  You may wish to forward this technical note to the vendor of your application server to help them assess the vulnerability of your application server to header injection attacks.

Vulnerability

All versions of Confluence prior to 2.10.2 are vulnerable to this security flaw.

Fix

The fix updates the Seraph framework to a version which correctly encodes and validates redirect URLs before sending them back to the user.

To patch your existing installation of Confluence, please refer to CONF-14275. This JIRA issue contains the downloadable patch file and instructions on how to patch your existing Confluence installation.

Alternatively, install or upgrade to Confluence version 2.10.2. (See the release notes.) The Confluence 2.10.2 installation files can be downloaded from the download centre.

For more information, please refer to CONF-14275.

Confluence Security Advisory 2009-04-15

In this advisory:

- XSS Vulnerability in Various Confluence Macros
  - Severity
  - Risk Assessment
  - Risk Mitigation
  - Vulnerability
  - Fix
- HTTP Header Injection Flaw with Attachment Filenames
  - Severity
  - Risk Assessment
  - Risk Mitigation
  - Vulnerability
  - Fix

XSS Vulnerability in Various Confluence Macros

Severity

Atlassian rates this vulnerability as **high**, according to the scale published in Confluence Security. The scale allows us to rank a vulnerability as critical, high, moderate or low.

Risk Assessment

We have identified and fixed two security flaws which may affect Confluence instances in a public environment. These flaws are all cross-site scripting (XSS) vulnerabilities in Confluence's Index and Widget Macros. Each vulnerability potentially allows a malicious user (attacker) to embed their own JavaScript into a Confluence page, which will be executed when the page is rendered.

- The hacker might take advantage of the flaw to steal other users' session cookies or other credentials, by sending the credentials back to the hacker's own web server.
- The hacker's text and script might be displayed to other people viewing the Confluence page. This is potentially damaging to your company's reputation.

You can read more about XSS attacks at cgisecurity, CERT and other places on the web.

Risk Mitigation

We recommend either patching or upgrading your Confluence installation to fix this vulnerability. Please see the 'Fix' section below.

Alternatively if you are not in a position to undertake this immediately and you judge it necessary, you can disable public access (e.g. anonymous access and public sign-on) to your wiki until you have applied the necessary patch or upgrade. For even tighter control, you could restrict access to trusted groups.

You could also temporarily disable the **Widget Connector** plugin and the **Index Macro module of the Confluence Advanced Macros** plugin until you have applied the necessary patch or upgrade. Be aware, however, that this will cause any occurrence of these macros on existing pages or blogs in your Confluence site to render with 'Unknown Macro' indications.

Vulnerability

All versions of Confluence prior to 2.10.3 are vulnerable to this security flaw.

Fix

The fixes include an update to the Index Macro, such that it correctly renders content on the page and an update to the Widget Macro, such that it correctly encodes all parameters passed to it.

To patch your existing installation of Confluence, please refer to CONF-14753 for the Index Macro and CONF-14337 for the Widget Macro. These JIRA issues contain the downloadable patch files and instructions on how to patch your existing Confluence installation.

Alternatively, install or upgrade to Confluence version 2.10.3. (See the release notes.) The Confluence 2.10.3 installation files can be downloaded from the download centre.

For more information, please refer to CONF-14753 and CONF-14337.

> Our thanks to **Igor Minar**, who reported one of the XSS vulnerabilities listed above. We fully support the reporting of vulnerabilities and we appreciate it when people work with us to identify and solve the problem.

HTTP Header Injection Flaw with Attachment Filenames

Severity

Atlassian rates this vulnerability as **high**, according to the scale published in Confluence Security. The scale allows us to rank a vulnerability as critical, high, moderate or low.

Risk Assessment

We have identified and fixed a security flaw with attachment filenames. This vulnerability could lead to an HTTP Header Injection attack through the upload of attachments with modified filenames designed to exploit this flaw. An attacker could insert malicious code into the HTTP response, which would be executed in the user's session.

- The attacker may take advantage of this flaw to steal other users' session cookies or other credentials, by sending the credentials back to the attacker's own web server.
- The attacker could also gain control over the underlying system, based on the privileges of the user whose session cookie has been stolen.
- The attacker could redirect the user to undesirable web sites. This is potentially damaging to your company's reputation.

Risk Mitigation

We strongly recommend either patching or upgrading your Confluence installation to fix this vulnerability. Please see the 'Fix' section below.

If you judge it necessary, you can disable public access (e.g. anonymous access and public sign-on) to your wiki until you have applied the necessary patch or upgrade. For even tighter control, you could restrict access to trusted groups.

Alternatively, you may consider taking the following step, although the time required to fix this vulnerability and the extent of its effectiveness will depend on your application server running Confluence and its configuration:

- Consult the vendor of your application server to see whether your application server is immune to header injection vulnerabilities or has configuration options to prevent such attacks. For example, the Coyote (HTTP) connector in Tomcat version 5.5 and later is immune to header injection attacks, as acknowledged in this reference.
  Technical note: In your application server, header injection vulnerabilities can be mitigated if the setHeader(), addHeader(), and sendRedirect() methods in the HttpServletResponse class have their parameters properly checked for header termination characters.
  You may wish to forward this technical note to the vendor of your application server to help them assess the vulnerability of your application server to header injection attacks.

Vulnerability

All versions of Confluence prior to 2.10.3 are vulnerable to this security flaw.

Fix

The fix includes a new header-injection prevention filter in Confluence, which ensures attachment filenames or any other user-provided data is correctly encoded before being included in HTTP headers.

To patch your existing installation of Confluence, please refer to CONF-14704. This JIRA issue contains the downloadable patch files and instructions on how to patch your existing Confluence installation.

Alternatively, install or upgrade to Confluence version 2.10.3. (See the release notes.) The Confluence 2.10.3 installation files can be downloaded from the download centre.

For more information, please refer to CONF-14704.

Confluence Security Advisory 2009-06-01

In this advisory:

- XSS Vulnerability in Various Confluence Actions and Macros
  - Severity
  - Risk Assessment
  - Risk Mitigation
  - Vulnerability
  - Fix

XSS Vulnerability in Various Confluence Actions and Macros

Severity

Atlassian rates these vulnerabilities as **high**, according to the scale published in Confluence Security. The scale allows us to rank a vulnerability as critical, high, moderate or low.

Risk Assessment

We have identified and fixed a number of security flaws which may affect Confluence instances in a public environment. These are cross-site scripting (XSS) that affect various Confluence page/blog features and functions.

- The hacker might take advantage of the flaw to steal other users' session cookies or other credentials, by sending the credentials back to the hacker's own web server.
- The hacker's text and script might be displayed to other people viewing the Confluence page. This is potentially damaging to your company's reputation.

You can read more about XSS attacks at cgisecurity, CERT and other places on the web.

Risk Mitigation

We recommend either patching or upgrading your Confluence installation to fix these vulnerabilities. Please see the 'Fix' section below.

Alternatively, if you are not in a position to undertake this immediately and you judge it necessary, you can disable public access (e.g. anonymous access and public sign-on) to your wiki until you have applied the necessary patch or upgrade. For even tighter control, you could restrict access to trusted groups.

Vulnerability

A hacker can inject their own JavaScript into various Confluence URLs — see the table below for the affected functional areas. A URL may be invoked when a user performs a specific function in Confluence, such as clicking a link or a button. The URL can also be invoked by simply entering it into the browser address bar. If rogue JavaScript is injected into such a URL, the JavaScript will be executed when a user invokes the URL.

For more details please refer to the related JIRA issue, also shown in the table below.

Affected Confluence Functionality	Affected Confluence Versions	Fix Availability	More Details
Concurrent page edit message	All versions (1.0 to 2.10.3 inclusive)	2.9.2 and 2.10.3	CONF-15883
Gallery Macro (Confluence Advanced Macros Plugin)	All versions (1.0 to 2.10.3 inclusive)	2.10.3	CONF-15376
View File Macro (Office Connector Plugin)	2.10.0 to 2.10.3 inclusive *	2.10.3	CONF-15402
Instant Messenger Macro	All versions (1.0 to 2.10.3 inclusive)	2.8.2, 2.9.2 and 2.10.3	CONF-15397
Contributors Macro	2.3 to 2.10.3 inclusive	2.9.2 and 2.10.3	CONF-15399
JIRA Issues Macro	All versions (1.0 to 2.10.3 inclusive)	2.10.3	CONF-15754

* This vulnerability may be present in earlier Confluence versions with the Office Connector plugin installed.

Fix

These issues have been fixed in Confluence 3.0 (see the release notes), which you can download from the download centre.

If you do not wish to upgrade to Confluence 3.0, you can download and install the patches provided on our JIRA site. You will need to upgrade to the latest point release for the major version of Confluence that you are running (e.g. if you are running Confluence 2.9, you will need to upgrade to version 2.9.2) and then apply the patches. For more information, please refer to the specific JIRA issues shown in the table of vulnerabilities above.

Confluence Security Advisory 2009-06-16

In this advisory:

- Page Content Vulnerabilities
  - Severity
  - Risk Assessment
  - Risk Mitigation
  - Vulnerability
  - Fix

Page Content Vulnerabilities



If you have already upgraded to Confluence 3.0, then you are not affected by the vulnerabilities described on this page and there is no need to take any further action.

Severity

Atlassian rates these vulnerabilities as **high**, according to the scale published in Confluence Security. The scale allows us to rank a vulnerability as critical, high, moderate or low.

Risk Assessment

We have identified and fixed two security vulnerabilities which may affect Confluence instances in a public environment. Both of these fixes are associated with a tightening of user access restrictions when either viewing specific page content or adding new page content.

The first of these vulnerabilities allows a user without permission to view a given page, to view the contents of any files attached to that page using the view file macro. This assumes that the user has permission to edit or create another page within the Confluence site and knows the name of the file attached to the page they cannot view. For more information, please refer to the JIRA issue CONF-15809.

The second of these vulnerabilities allows users with space administrator permissions to import pages to a Confluence space. The security level of this function has been tightened to permit only users with the system administration permission to access it. For more information, please refer to CONF-15267.

Risk Mitigation

If you have not already upgraded to Confluence 3.0, then we recommend either patching or upgrading your Confluence installation to fix these vulnerabilities. Please see the 'Fix' section below.

Alternatively, if you are not in a position to undertake this immediately and you judge it necessary, you can disable public access (e.g. anonymous access and public sign-on) to your wiki until you have applied the necessary patch or upgrade. For even tighter control, you could restrict access to trusted groups.

Vulnerability

All versions of Confluence up to and including version 2.10.3 with the Office Connector plugin installed are affected by the first view file macro vulnerability.

All versions of Confluence 2.10.x are affected by the second page imports vulnerability.

Fix

These issues have been fixed in Confluence 3.0 (see the release notes), which you can download from the download centre.

If you do not wish to upgrade to Confluence 3.0, you can download and install the patches provided on our JIRA site. You will need to upgrade to the latest point release for the major version of Confluence that you are running (e.g. if you are running Confluence 2.10.0, you will need to upgrade to version 2.10.3) and then apply the patches. For more information, please refer to the specific JIRA issues shown below.

To download the patch to fix the first view file macro vulnerability, please refer to CONF-15809.

To download the patch to fix the second page import vulnerability, please refer to CONF-15267.

Confluence Security Advisory 2009-08-20

In this advisory:

- Privilege Escalation Vulnerability in Profile Picture Handling

        - Severity
        - Risk Assessment
        - Risk Mitigation
        - Vulnerability
        - Fix
    - XSS Vulnerability in Various Page and Blog Post Features and Functions
        - Severity
        - Risk Assessment
        - Risk Mitigation
        - Vulnerability
        - Fix

Privilege Escalation Vulnerability in Profile Picture Handling

Severity

Atlassian rates this vulnerability as **high**, according to the scale published in Confluence Security. The scale allows us to rank a vulnerability as critical, high, moderate or low.

Risk Assessment

We have identified a privilege escalation vulnerability, which could provide an attacker with access to administrative areas and functions of Confluence when specifying a profile picture. Under some circumstances, the attacker could gain access to Confluence administrative functions that they are not authorised to use.

Risk Mitigation

To address the issue, you should upgrade to Confluence 3.0.1 as soon as possible or follow the patch instructions in the Fix section below. If you judge it necessary, you can disable public signon to your wiki until you have applied the necessary patch or have performed the upgrade. For even tighter control, you could also restrict access to trusted groups or additionally, disable anonymous access until your system is patched or upgraded.

Vulnerability

The profile picture handling feature in all versions of Confluence up to 3.0.0 are affected by this issue. However, the Form Token Handling mechanism available in Confluence 3.0.0 and later means that the administrative areas in these versions of Confluence cannot be compromised by this vulnerability.

Fix

This issue has been fixed in Confluence 3.0.1 (see the release notes), which you can download from the download centre.

If you do not wish to upgrade to Confluence 3.0.1 and you are running Confluence 2.10.x, you can download and install the patches provided on our JIRA site. We strongly recommend that you upgrade to the latest point release (2.10.3) before applying the patch. For more information, please refer to CONF-16141.

> Our thanks to **Elliot Kendall** of **Emory University**, who reported this vulnerability. We fully support the reporting of vulnerabilities and we appreciate it when people work with us to identify and solve the problem.

XSS Vulnerability in Various Page and Blog Post Features and Functions

Severity

Atlassian rates these vulnerabilities as **high**, according to the scale published in Confluence Security. The scale allows us to rank a vulnerability as critical, high, moderate or low.

Risk Assessment

We have identified and fixed a number of XSS vulnerabilities in various Confluence page/blog features and functions, which may affect Confluence instances in a public environment.

XSS vulnerabilities potentially allow a malicious user (attacker) to embed their own JavaScript into a Confluence page.

- The attacker might take advantage of the vulnerability to steal other users' session cookies or other credentials, by sending the credentials back to the attacker's own web server.
- The attacker's text and script might be displayed to other people viewing the Confluence page. This is potentially damaging to your company's reputation.

You can read more about XSS attacks at cgisecurity, CERT and other places on the web.

Risk Mitigation

We recommend either patching or upgrading your Confluence installation to fix these vulnerabilities. Please see the 'Fix' section below.

Alternatively, if you are not in a position to undertake this immediately and you judge it necessary, you can disable public access (e.g. anonymous access and public sign-on) to your wiki until you have applied the necessary patch or upgrade. For even tighter control, you could restrict access to trusted groups.

Vulnerability

An attacker can inject their own JavaScript into the Confluence actions listed in the table below. Each of the actions is invoked when a user performs a specific function in Confluence, such as clicking a link or a button. The actions can also be invoked by simply entering the URL into the browser address bar. The rogue JavaScript will be executed when a user invokes the URL.

For more details please refer to the related JIRA issue, also shown in the table below.

Confluence action	Affected Confluence Versions	Fix Availability	More Details
Clicking a username link	3.0.0	3.0.0 and 3.0.1	CONF-15970
Moving pages between spaces	2.8 to 2.10.3 inclusive	2.10.x and 3.0.1	CONF-16019 * CONF-16135 *
Entering content into the WebDAV Configuration page	3.0.0 2.10.x with version 2.0 of the WebDAV plugin	2.10.x, 3.0.0 and 3.0.1	CONF-16136
Entering content into the PDF Export Stylesheet	3.0.0	3.0.0 and 3.0.1	CONF-16209

* Applying the patch for one of these issues fixes the other.

Fix

These issues have been fixed in Confluence 3.0.1 (see the release notes), which you can download from the download centre.

If you do not wish to upgrade to Confluence 3.0.1, you can patch your existing installation by downloading and installing the patched files provided on our JIRA site. For the WebDAV plugin vulnerability, this would involve upgrading the version of the plugin. We strongly recommend that you upgrade to the latest point release of the major version of Confluence that you are running before applying the patches. For example, if you are running Confluence 2.10.1, you should upgrade to version 2.10.3 and then apply the patches. For more information, please refer to the specific JIRA issues shown in the table of vulnerabilities above.

Confluence Security Advisory 2009-10-06

In this advisory:

- Session Fixation Vulnerability
  - Severity
  - Risk Assessment
  - Risk Mitigation
  - Vulnerability
  - Fix
- XSS Vulnerability in Various Confluence Macros
  - Severity
  - Risk Assessment
  - Risk Mitigation
  - Vulnerability
  - Fix

Session Fixation Vulnerability

Severity

Atlassian rates these vulnerabilities as **high**, according to the scale published in Confluence Security. The scale allows us to rank a vulnerability as critical, high, moderate or low.

Risk Assessment

We have identified and fixed a security vulnerability which may affect Confluence instances in a public environment. This vulnerability could lead to a session fixation attack, in which the malicious user (attacker) can gain access to a victim's Confluence resources whilst the victim is logged in to their Confluence user account.

The attacker does this by fixating (or setting) their session ID onto the victim's computer. While the victim is logged in, all the victim's privileges are associated with the attacker's session ID, effectively granting the attacker access to all of the Confluence data and resources accessible to the victim.

For more information about session fixation attacks, please refer to the following sources:

- Chris Shiflett's Security Corner article
- The Web Application Security Consortium's overview

Risk Mitigation

We recommend either patching or upgrading your Confluence installation to fix these vulnerabilities. Please see the 'Fix' section below.

Alternatively, if you are not in a position to undertake this immediately and you judge it necessary, you can disable public access (e.g. anonymous access and public sign-on) to your wiki until you have applied the necessary patch or upgrade. For even tighter control, you could restrict access to trusted groups.

Vulnerability

All versions of Confluence prior to 3.0.2 are vulnerable to this security issue.

Fix

These issues have been fixed in Confluence 3.0.2 (see the release notes), which you can download from the download centre.

If you do not wish to upgrade to Confluence 3.0.2 and you are currently running Confluence version 2.10.x or 3.0.x, you can patch your existing installation by downloading the appropriate patch file attached to JIRA issue CONF-15108 and installing the patch file using the instructions provided in this JIRA issue.

> Our thanks to **Ben L Broussard** who reported this vulnerability. We fully support the reporting of vulnerabilities and we appreciate it when people work with us to identify and solve the problem.

XSS Vulnerability in Various Confluence Macros

Severity

Atlassian rates these vulnerabilities as **high**, according to the scale published in Confluence Security. The scale allows us to rank a vulnerability as critical, high, moderate or low.

Risk Assessment

We have identified and fixed a number of security vulnerabilities which may affect Confluence instances in a public environment. These flaws are cross-site scripting (XSS) vulnerabilities in Confluence's pagetree, userlister and content by label macros. These XSS vulnerabilities potentially allow an attacker to embed their own JavaScript into a Confluence page.

- The attacker might take advantage of the vulnerability to steal other users' session cookies or other credentials, by sending the credentials back to the attacker's own web server.
- The attacker's text and script might be displayed to other people viewing the Confluence page. This is potentially damaging to your company's reputation.

You can read more about XSS attacks at cgisecurity, CERT and other places on the web.

Risk Mitigation

We recommend either patching or upgrading your Confluence installation to fix these vulnerabilities. Please see the 'Fix' section below.

Alternatively, if you are not in a position to undertake this immediately and you judge it necessary, you can disable public access (e.g. anonymous access and public sign-on) to your wiki until you have applied the necessary patch or upgrade. For even tighter control, you could restrict access to trusted groups.

Vulnerability

An attacker can inject their own JavaScript into the Confluence actions listed in the table below. Each of the actions is invoked when a user performs a specific function in Confluence, such as clicking a link or a button. The actions can also be invoked by simply entering the URL into the browser address bar. The rogue JavaScript will be executed when a user invokes the URL.

For more details please refer to the related JIRA issue, also shown in the table below.

Confluence action	Affected Confluence Versions	Fix Availability	More Details
Pagetree Macro	2.8.0 – 3.0.1	2.10.0 – 3.0.2 inclusive	CONF-16651
Userlister Macro	2.6.0 – 3.0.1	2.10.0 – 3.0.2 inclusive	CONF-16644
Content by Label Macro	2.10.0 – 3.0.1	2.10.0 – 3.0.2 inclusive	CONF-15440

Fix

These issues have been fixed in Confluence 3.0.2 (see the release notes), which you can download from the download centre.

If you do not wish to upgrade to Confluence 3.0.2, you can patch your existing installation by upgrading the plugins for these macros via the Confluence Plugin Repository to the version indicated in the JIRA issues listed in the vulnerability section (above).

Confluence Security Advisory 2009-12-08

In this advisory:

- XSS Vulnerability in Various Confluence Actions and Macros
  - Severity
  - Risk Assessment
  - Risk Mitigation
  - Vulnerability
  - Fix

XSS Vulnerability in Various Confluence Actions and Macros

Severity

Atlassian rates these vulnerabilities as **high**, according to the scale published in Severity Levels for Security Issues. The scale allows us to rank a vulnerability as critical, high, moderate or low.

Risk Assessment

We have identified and fixed a number of security vulnerabilities which may affect Confluence instances in a public environment. These flaws are cross-site scripting (XSS) vulnerabilities that could occur when creating a page or blog post in a personal space, using the `indexbrowser.jsp` form and when using the gallery macro.

- The attacker might take advantage of the vulnerability to steal other users' session cookies or other credentials, by sending the credentials back to the attacker's own web server.
- The attacker's text and script might be displayed to other people viewing the Confluence page. This is potentially damaging to your company's reputation.

You can read more about XSS attacks at cgisecurity, CERT and other places on the web.

Risk Mitigation

We recommend either patching or upgrading your Confluence installation to fix these vulnerabilities. Please see the 'Fix' section below.

Alternatively, if you are not in a position to undertake this immediately and you judge it necessary, you can disable public access (e.g. anonymous access and public sign-on) to your wiki until you have applied the necessary patch or upgrade. For even tighter control, you could restrict access to trusted groups.

Vulnerability

An attacker can inject their own JavaScript into the Confluence actions listed in the table below. Each of the actions is invoked when a user performs a specific function in Confluence, such as clicking a link or a button. The actions can also be invoked by simply entering the URL into the browser address bar. The rogue JavaScript will be executed when a user invokes the URL.

For more details please refer to the related JIRA issue, also shown in the table below.

Confluence action	Affected Confluence Versions	Fix Availability	More Details
Page or blog post creation in a personal space	2.10 – 3.0.2	3.0.0 – 3.1 inclusive	CONF-17031
Using the `indexbrowser.jsp` form	All versions prior to and including 3.0.2	3.0.0 – 3.1 inclusive	CONF-17165
Gallery macro	2.9 – 3.0.2	3.0.0 – 3.1 inclusive	CONF-17361
Page tree and page tree search macros	2.9 – 3.0.2	2.8 – 3.1 inclusive	CONF-17967
Status updates tab of the user profile area	3.0.0 – 3.0.2	3.0.0 – 3.1 inclusive	CONF-17933

Fix

These issues have been fixed in Confluence 3.1 (see the release notes), which you can download from the download centre.

If you do not wish to upgrade to Confluence 3.1, you can patch your existing installation by upgrading the plugins for these macros via the Confluence Plugin Repository to the version indicated in the JIRA issues listed in the vulnerability section (above).

Design and Layout

- Choosing a Default Language
- Custom Decorator Templates
- Customising Look and Feel Overview
    - Customising Colour Schemes
    - Customising Layouts
        - Adding a Navigation Sidebar
            - Adding an All Versions Section to your Navigation Bar
        - Upgrading Custom Layouts
    - Global Templates
    - Importing Templates
    - Modify Confluence Interface Text
    - Working With Decorator Macros
    - Customising a Specific Page
    - Customising PDF or HTML Content
    - Customising the Dashboard
    - Customising the eMail Templates
    - Customising the Login Page
- Themes Overview
    - Applying a Theme to a Site
    - Customising the Left Navigation Theme
    - Modifying Look and Feel (for themes)
        - Configuring the theme plugin
        - Including Cascading Stylesheets in Themes
    - Creating a Theme

RELATED TOPICS

Modifying Confluence Interface Text
Site Configuration

Choosing a Default Language

The default language in Confluence is applied to all spaces in your site. A user can choose to use another language in their user settings. This will only override the default language for their sessions, not for all users.

Confluence ships with a few language packs. You can install new language packs, as described on .

To change the default language,

1. From the '**Administration Console**' click on '**Languages**' under the heading '**Configuration**' in the left panel. The '**Language Configuration**' screen will display.
2. Select the language that you want to use as the default language for your Confluence instance by clicking either the flag image or the name of the language. The default language will be changed.

RELATED TOPICS

Configuring the Site Home Page
Configuring the Destination of View Space Links
Configuring the Server Base URL
Configuring the Site Support Address
Customising Default Space Content
Editing the Global Logo
Editing the Site Title
Editing the Site Welcome Message
Showing Link Icons

Editing User Settings

Administrators Guide Home Confluence Documentation Home

Custom Decorator Templates

About Decorators

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

What this means for Confluence is that you can customise the look and feel of almost all of your Confluence site through editing three decorators:

- The "Main" decorator defines the look and feel of most pages on the site
- The "Popup" decorator defines the look and feel of the popup windows such as the "Insert Link" and "History" pages.
- The "Printable" decorator defines the look and feel of the printable versions of pages (available through the  icon on each page)

You can view and edit these decorators from within Confluence: they are available from the "Layouts" option on the site's Administration menu. Changes to the decorators will affect all spaces hosted on that Confluence installation.

The decorator that is used to draw Confluence's administrative pages can not be edited from within Confluence. This means that if you make some editing mistake that renders the rest of the site unuseable, the administrative pages should still be available for you to fix the template.

Browsing the Default Decorators

At any time, you can browse the default decorators that come packaged with Confluence by following the "View Default" links on the "Site Layouts" page. The template browser also allows you to view the "#parsed" templates that are included within the template when it is compiled. While you can't edit these included templates, you will probably have to copy some or all of them into your custom template as you do your customisation.

Editing Custom Decorators: Add a Logo

To edit Confluence decorators, you should have a good knowledge of HTML, and some understanding of the Velocity templating language.

The first thing you will see when you choose to create a custom "Main" decorator is... there's not much to edit. By default, most of the content of this decorator is included from other files:

```
<html>
<head>
            <title>$title - Confluence</title>

            #standardHeader()
</head>
<body onload="placeFocus()">

<div id="Content">
            <table border="0" cellpadding="0" cellspacing="0" width="100%">
                        <tr>
                                    <td width="60%" rowspan=2 class="logocell"
>#pagetitle("spacenametitle")</td>

                                    <td width="40%" align="right" valign="top"
>#globalnavbar("table")</td>
                        </tr>
                        #if ($setup.isSetupComplete())
                        <tr>
                                    <td align=right valign="bottom">
                                                #usernavbar()
                                                #printableicon()
                                                #helpicon()
                                    </td>
                        </tr>
                        #end
            </table>

            #breadcrumbsAndSearch()

            <table border="0" cellpadding="5" cellspacing="0" width="100%"><tr><td>
<table border="0" cellpadding="0" cellspacing="0" width="100%"><tr>

            <td valign="top" class="pagebody">
## The "toolbar-style" page operations
## #if ($page.getProperty("page.operations"))
## <table align="right" class="toolbar"><tr><td>
## $page.getProperty("page.operations")
## </td></tr></table>
## #end

                        #if ($page.getProperty("page.surtitle"))
                                    $page.getProperty("page.surtitle")
                        #end

                        #if (!$page.getProperty("page.no-page-header"))
                                    <div class="pageheader">
                                                <span class="pagetitle">$title</span>
                                    </div>
                        #end

                        $body
            </td>

#parse ("/decorators/includes/complete_footer.vmd")
```

We can add our logo, changing the "logocell" table cell:

```
<td width="60%" rowspan=2 class="logocell">
<img align="right"
src=http://www.atlassian.com/images/atlassian_logo.gif
width="203" height="60">#pagetitle("spacenametitle")</td>
```

When you insert this into the right section of the template and hit save, visitors to the site will see the logo at the top of each page. Note, the administrative pages will be unaffected: you will have to go to the dashboard or to a space to see the changes you have made.

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.

If Something Goes Terribly Wrong

From the "Site Layouts" page in Confluence's administrative menu, you can delete your custom templates. When you do this, the default template will be restored, fixing anything that may have been broken.

Alternatively, the custom templates are stored inside the `confluence.home` directory you defined in `confluence-init.properties` when installing Confluence. If you have somehow managed to render Confluence completely unuseable through editing your templates, simply delete the `confluence.home/velocity` directory, and restart Confluence. The default templates will be restored.

WARNING: Only delete the `velocity` directory! Changing anything else inside your `confluence.home` is dangerous, and you could lose important data!

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 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.

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. Because we only officially support the modification of the three global decorator files, other changes may interact unpredictably with future versions of Confluence. When upgrading, you should always test your custom modifications thoroughly before deploying them on a live site.

Customising Look and Feel Overview

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

Any changes you make to the look and feel of the site at the global level will be applied as the default look and feel for all the spaces in the site. This means that any customisations will only be reflected in the "Default" theme. No other theme will have an impact from this change. An individual space can be configured to have its own look and feel through the space administration screens.

Here's how you can customise the look and feel of your site:

- **Colour Scheme** : Change the colour scheme of the user interface.
- **Layouts** : Edit how the controls are laid out in the site. This does not change the actual page layouts but the way the surrounding controls appear in the page.
- **Themes** : Use themes for advanced layout customisation.

RELATED TOPICS

- Adding a Site-Wide Banner
- Editing the Footer
- Upgrading Custom Layouts
- Customising a Specific Page
- Customising Colour Schemes
- Customising Layouts
- Customising Look and Feel Overview
- Global Templates
- Working With Decorator Macros

Administrators Guide Home Confluence Documentation Home

Customising Colour Schemes

A Confluence administrator can configure a new colour scheme for the site dynamically from the **Administration Console**.

The default colour scheme for the site will also become the default for all spaces within it. However, it is possible for space administrators to configure a different colour scheme for spaces from the space administration screens.

To change the site's colour scheme,

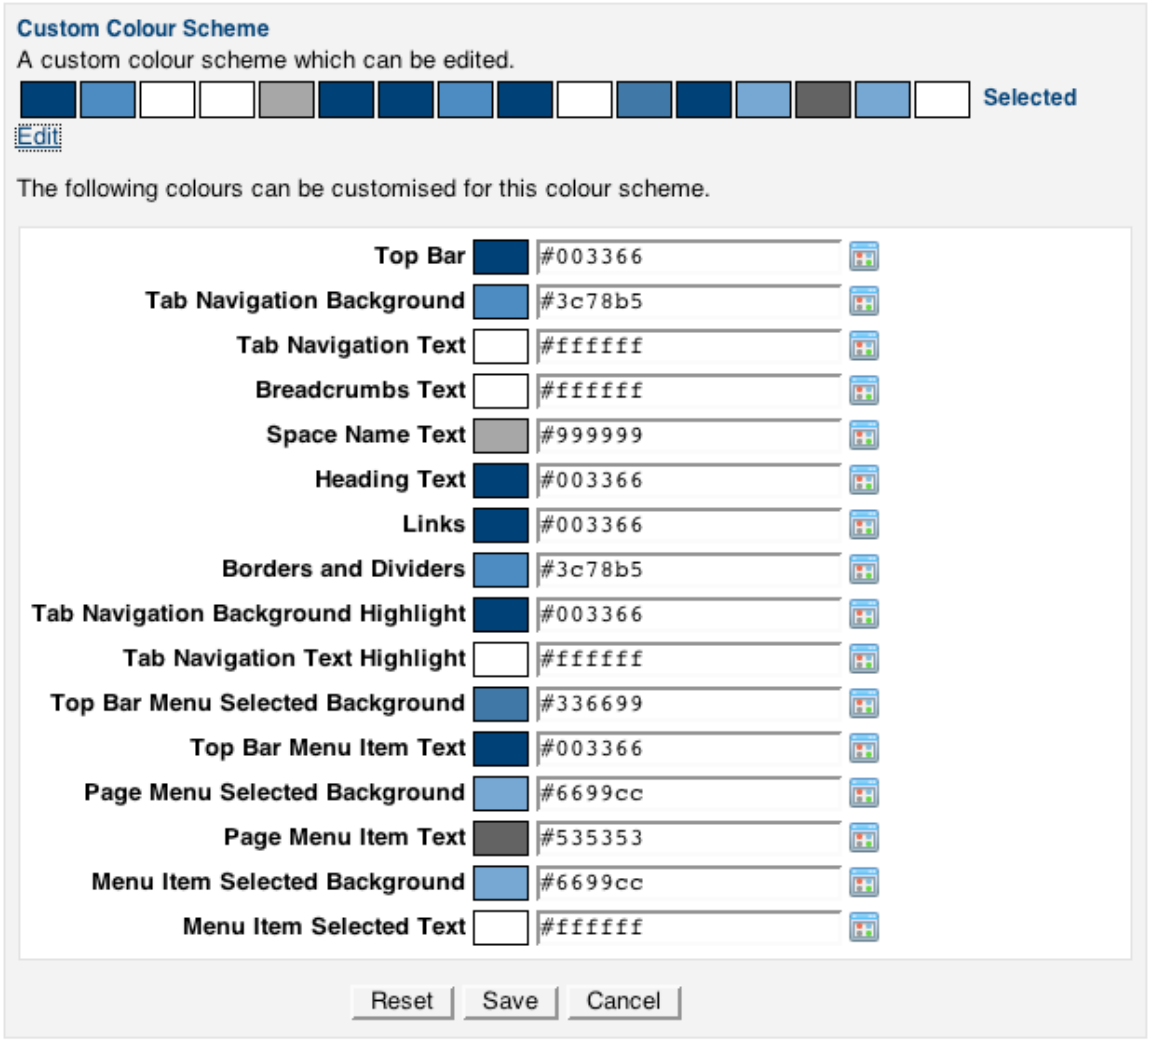
1. Go to the Confluence '**Administration Console**'. To do this:
    - Open the '**Browse**' menu and select '**Confluence Admin**'. The 'Administration Console' view will open.
2. Select '**Colour Scheme**' in the left-hand panel.
   This will bring up a new screen. See screenshot below.
3. Click '**Edit**'. Enter standard HTML/CSS2 colour codes, or use the colour-picker to choose a new colour from the palette provided. Any changes you make will immediately be reflected across the Confluence installation.

The colour scheme applies to the following UI elements:

- **Top Bar** - the bar across the top of the page that contains the breadcrumbs
- **Tab Navigation Background** - the background colour of the tab navigation menus
- **Tab Navigation Text** - the text of the tab navigation menus
- **Breadcrumbs Text** - the breadcrumbs text in the top bar of the page
- **Space Name Text** - the text of the current space name located above the page title
- **Heading Text** - all heading tags throughout the space.
- **Links** - all links throughout the space.
- **Borders and Dividers** - table borders and dividing lines.
- **Tab Navigation Background Highlight** - the background colour of the tab navigation menu when highlighted
- **Tab Navigation Text Highlight** - the text of the tab navigation menu when highlighted
- **Top Bar Menu Selected Background** - the background colour of the top bar drop down menu when selected
- **Top Bar Menu Item** - the text colour of the menu items in the top bar drop down menu
- **Page Menu Selected Background** - the background colour of the drop down page menu when selected
- **Page Menu Item Text** - the text of the menu items in the drop down page menu
- **Menu Item Selected Background** - the background colour of the menu item when selected (applies to both the top bar and page drop down menus)
- **Menu Item Selected Text** - the text colour of the menu item when selected (applies to both the top bar and page drop down menus)

Please note that some UI elements are specific to the default theme and may not take affect for other themes.

Screenshot : Editing a site's colour scheme



Handy Hint
If you mess things up, just click the '**Reset**' button and then try again.

RELATED TOPICS

- Adding a Site-Wide Banner
- Editing the Footer
- Upgrading Custom Layouts
- Customising a Specific Page
- Customising Colour Schemes
- Customising Layouts
- Customising Look and Feel Overview
- Global Templates
- Working With Decorator Macros

Administrators Guide Home Confluence Documentation Home

Customising Layouts

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

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

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

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

You require System Administrator permissions to perform these customisations.

You can also customise the layout files for a given space only. For more information, refer to Customising Layouts for a Space.

Space layout file customisations override the equivalent site layout file customisations.

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

Confluence is built on top of the open source SiteMesh library, a web-page layout system. Read more on the SiteMesh website. To edit the layout of Confluence, you will need to modify these decorator files. A decorator file is a .vmd file and is written in a very simple programming language called Velocity. You can learn more from the Velocity User Guide.

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

The decorator files are grouped into:

- **Site layouts** : These are used to define the controls that surround each page in the site. For example, the header and the footer.
- **Content layouts** : These control the appearance of content such as pages and news items: they don't change the way the pages themselves are displayed, but allow you to alter the way the surrounding comments or attachments are displayed.
- **Export Layouts**: These control the appearance of spaces and pages when they are exported to HTML. If you are using Confluence to generate a static website, for example, you will need to modify these layouts.

Editing a site decorator file

1. Go to the Confluence '**Administration Console**'. To do this:
    - Open the '**Browse**' menu and select '**Confluence Admin**'. The 'Administration Console' view will open.

2. Select '**Layouts**' under '**Look and Feel**' in the left-hand navigation panel. The decorators are grouped under **Site**, **Content** and **Export** layouts.
    - Click '**View Default**' to view the vmd file.
    - Click '**Create Custom**' to edit the default vmd file. This will open up the vmd file in edit mode.
3. Make changes and click '**Update**'.

If something goes wrong : Click '**Reset Default**' to revert to the original layouts.

Using Velocity macros

When editing Custom Decorator Templates, there are a number of macros available to define complex or variable parts of the page such as menus and breadcrumbs. You may insert these macros anywhere in your templates. More information on Working With Decorator Macros.

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 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.

In Confluence 2.6 and later, some Velocity files are located inside the Confluence JAR file that can be found at `confluence/WEB-INF/lib/confluence-x.x.x.jar`. To override files inside this JAR (which you can open with any ZIP tool like WinZip or 7-Zip), put your customised file in the same directory structure under `confluence/WEB-INF/classes/`.

For example, the file `templates/macros/alphaindex.vm` inside confluence.jar can be replace by putting your custom file in `WEB-INF/classes/templates/macros/alphaindex.vm`. You do not need to modify the file inside the JAR.

See also Editing Files within JAR Archives.

RELATED TOPICS

- Adding a Site-Wide Banner
- Editing the Footer
- Upgrading Custom Layouts
- Customising a Specific Page
- Customising Colour Schemes
- Customising Layouts
- Customising Look and Feel Overview
- Global Templates
- Working With Decorator Macros

Velocity Template Overview
Basic Introduction to Velocity

Administrators Guide Home Confluence Documentation Home

Adding a Navigation Sidebar

You can include a left-hand navigation sidebar (table of contents) in your Confluence space. There are two ways to do this:

- **Recommended: Use the Documentation Theme** – The Documentation theme provides the left-hand navigation sidebar that you see in this documentation. **Please go to the page that tells you how to configure the Documentation theme**.
- **Customise the Page Layouts** — This is an alternative method (documented below) that is more complex to set up than the Documentation theme and requires more maintenance with Confluence major release upgrades.

Every time you upgrade Confluence, you must re-apply the layout customisations described on this page

When you upgrade to a new major Confluence version (e.g. from Confluence 2.9.x to Confluence 2.10.x or from Confluence 3.0.x to Confluence 3.1.x), you will need to re-apply the layout customisation. See instructions below.

Customising your Layouts to Add a Navigation Sidebar

Screenshot: A left-hand navigation bar resulting from customising the page layouts



Follow the instructions below to add the navigation sidebar to your Confluence space.

Permissions required
To customise a space layout as described below, you must be a Space Administrator in the given space and you must be a System Administrator on the Confluence site.

Step 1. Create the TreeNavigation Page

First, you will create a Confluence page containing the `pagetree` macro. This is just a normal Confluence page. The only slight oddity is that it should reside at the root of your space, instead of under the space's home page.

Follow these instructions:

1. Go to the '**Space Pages**' view for the current space. To do this:
   - Go to a page in the space, open the '**Browse**' menu and select '**Pages**'. The 'Space Pages' view will open.

   You are now at the 'root' level of your space. The 'root' level contains pages that are added above the space's home page, not as children of the home page.
2. At the root level of the space, create a page named 'TreeNavigation'.
3. On the page, insert the following text:

```
1.{pagetree}
```

4. Now decide if you want to add extra functionality to your page tree. By default, using the code above, the page tree will use the home page of the space as its root. You can choose to:
   - Specify a different root for your page tree.
   - Add a search box at the top of the tree.
   - Allow the viewers to expand and collapse the whole tree.
   - Control other aspects of the display.

   For more information, read about the Pagetree macro.

Step 2. Change the Page Layout on your Space

Now you will change the page layout on your space, to include the above page on the left of every web page displayed.

1. Go to the '**Space Admin**' tab of the Browse Space view. To do this:
   - Go to a page in the space, open the '**Browse**' menu and select '**Space Admin**'. The 'Space Administration' view will open.
     '**Space Admin**' is only displayed if you are a space administrator for that space or you are a Confluence system administrator.
2. Make sure the Confluence Default theme is selected from the '**Themes**' menu.
3. Click '**Layout**' under the 'Look and Feel' section.
   'Layout' is only displayed if you are a Confluence Site Administrator.
4. Click '**Create Custom**' under the 'Page Layout' section.
5. In the layout, locate the 'VIEW' section, and find this code:

```
1.<div class="wiki-content">
2.$body
3.</div>
```

6. Replace the above code block with this code:

```
01.#if ($action.isPrintableVersion() == false)
02.<style>
03..spacetree * ul{
04.padding-left:0px;
05.margin-left: 0px;
06.}
07..spacetree * li{
08.margin-left: 5px;
09.padding-left:5px;
10.}
11.
12.</style>
13.
14.<table cellspacing="2" cellpadding="5">
15.<tr>
16.<td valign="top" align="left" width="22%" bgcolor="#F9F9F9" class="noprint">
17.<div class="tabletitle">Table of Contents</div>
18.<div class="spacetree">
19.#includePage($helper.spaceKey "TreeNavigation")
20.</div>
21.</td>
22.<td valign="top" align="left" width="78%" class="pagecontent">
23.<div class="wiki-content">
24.$body
25.</div>
26.</td>
27.</tr>
28.</table>
29.#else
30.<div class="wiki-content">
31.          $body
32.</div>
33.#end
```

7. If you want to, you can change the table title in the above code from 'Table of Contents' to something else. For example, it might say 'Confluence Documentation'.
8. Save the updated layout.

Re-Applying the Customisation on Upgrade

When you upgrade to a new major Confluence version (e.g. from Confluence 2.9.x to Confluence 2.10.x or from Confluence 3.0.x to Confluence 3.1.x), you will need to re-apply this customisation.

Reason:
The new Confluence version may contain updates to the page layouts. Because you have customised the page layouts, Confluence will not overwrite your customisation. So your space will not get the latest updates until you set the layout to default and then re-apply your changes.

Here's how to do it:

1. First make a copy of your customised code, if you have changed it from the code above:
   - Go to '**Space Admin**, click '**Layout**' and edit the customised page layout (as created above).
   - Copy the section of code that inserts the customised left-hand navigation panel.
   - Close the page layout.
2. Click '**Reset Default**' next to '**Page Layout**', to set the page layout back to default. This will bring in the new code for the upgraded

version of Confluence.
3. Create a custom page layout as described in step 2 above, and reinsert the custom left-hand navigation code.
4. Save the updated layout.

The 'All Versions' section in the navigation bar
A number of people have asked how we created the '**All Versions**' section at the top of our navigation side bar. Take a look at Adding an All Versions Section to your Navigation Bar.

RELATED TOPICS

Configuring the Documentation Theme
Customising Layouts
Upgrading Custom Layouts
Example Confluence Designs

Adding an All Versions Section to your Navigation Bar

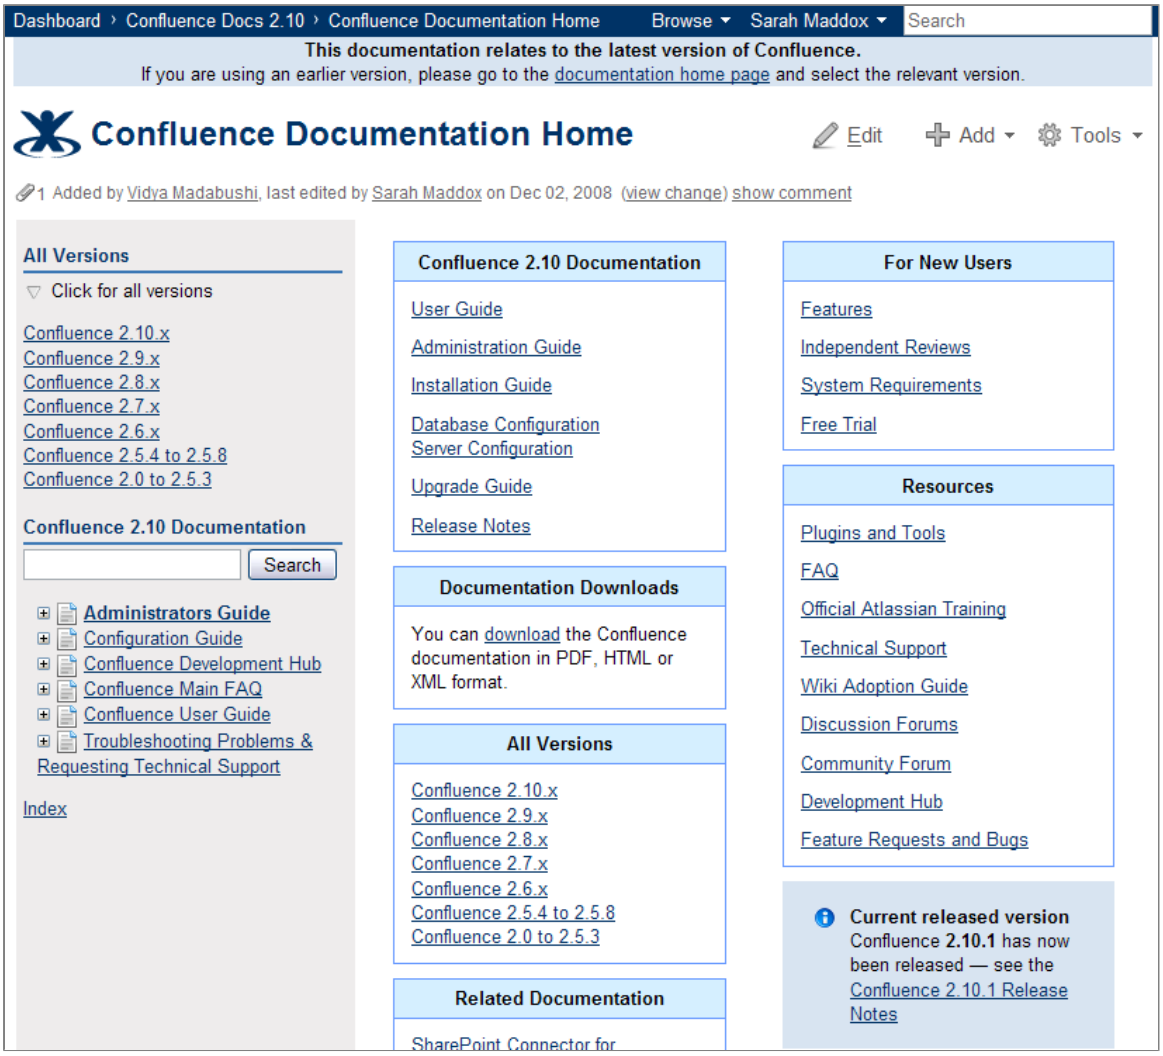
This page gives an example of how you might add an '**All Versions**' section to your navigation side bar, as currently used in the Confluence documentation, Crowd documentation and the other Atlassian product documentation spaces.

If you are viewing this page online on the Atlassian documentation wiki, you will be able to see the '**All Versions**' section at the top left of the navigation sidebar. Below is a screenshot.

A number of people have asked how we do it, so this page gives the answer. For details about creating the navigation side bar itself, please refer to Adding a Navigation Sidebar.

Hint: Viewing the Source Code of a Page
To see the Wiki Markup for one of the Atlassian documentation pages, open the '**Tools**' menu and select '**View Wiki Markup**'. You will see the macros and other markup used to create a page.

Screenshot: 'All Versions' section (expanded) at top left of navigation bar



Adding the Version Index to the Navigation Sidebar

This is how we added the '**All Versions**' section to the sidebar:

- For each product (Confluence, Crowd, Bamboo, etc) there is a page in the Inclusions Library of the ALLDOC space. The page lists all the versions of that product's documentation, linking to the relevant spaces. For example, here is the page for Confluence and the page for Crowd.
  We put the 'all versions' page in ALLDOC because the page is used in a number of different spaces, via the {include} macro. For example, the 'all versions' page may be included:
  - In every documentation space (each version) for the product concerned, such as DOC, CONF29, CONF28, CROWD, CROWD013, CROWD012, etc.
  - In the Enterprise Hosting doc space.
  - As a panel on the documentation home page, as shown in the 'All Versions' panel of the above screenshot, as well as in the left-hand navigation bar.
  - Any other places where useful.
- In each documentation space, there is a page called 'TreeNavigationVersions' like this one or this one, which copies in the content of the above 'all versions' page.
- For each documentation space, the space's page layout now includes two pages instead of just one:
  - The 'TreeNavigation' page, as already described on the page above.
  - The new 'TreeNavigationVersions' page.

Here's the relevant section of our page layout as it is currently for the Confluence documentation (DOC) space:

```
01.#if ($action.isPrintableVersion() == false)
02.<style>
03..spacetree * ul{
04.padding-left:0px;
05.margin-left: 0px;
06.}
07..spacetree * li{
08.margin-left: 5px;
09.padding-left:5px;
10.}
11.
12.</style>
13.
14.<table cellspacing="2" cellpadding="5">
15.<tr>
16.<td valign="top" align="left" width="30%" bgcolor="#eeecec" class="noprint">
17.<div class="tabletitle">All Versions</div>
18.<div class="spacetree">
19.#includePage($helper.spaceKey "TreeNavigationVersions")
20.</div>
21.<div class="tabletitle">Confluence 2.10 Documentation</div>
22.<div class="spacetree">
23.#includePage($helper.spaceKey "TreeNavigation")
24.</div>
25.</td>
26.<td valign="top" align="left" width="70%" class="pagecontent">
27.<div class="wiki-content">
28.$body
29.</div>
30.</td>
31.</tr>
32.</table>
33.#else
34.<div class="wiki-content">
35.          $body
36.</div>
37.#end
```

Adding the Expand/Collapse Functionality to the Version Index

Another question we are asked is how we group the content of the included page under a collapsible control or 'twisty'.

We use the {expand} macro. This is a 'user macro', which means that you can add it to your Confluence site by adding the code into the '**User Macros**' section of your Confluence Administration Console. The details are on the Expand macro's documentation page.

RELATED TOPICS

Adding a Navigation Sidebar

Upgrading Custom Layouts

As Confluence evolves, so do the default 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.



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

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

To reapply your custom layouts, you need to:

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

Step 1. Obtaining your Custom Layouts

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

If not, you should be able to find your customisations using the following method. This method extracts all site- and space-level layouts from your Confluence site as a single output. From this output, you should be able to identify your customisations.



This method is handy to use if you have:

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

Before Confluence 2.3, custom layouts are stored in the `velocity` directory within your Confluence home directory tree. You can open these files in any text editor.

In Confluence 2.3 and later, 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. Reapplying your Customisations

When you upgrade Confluence to another major release of Confluence, you will need to manually re-apply any customisations you made to any site-wide or space-specific layouts. Unless otherwise stated, you should not need to re-apply customisations after conducting a minor release upgrade of Confluence.



What are 'major' and 'minor' release upgrades?

Major release upgrades are ones where the 1st digit of Confluence's version number or the 1st digit after the 1st decimal place differ after the upgrade, for example, when upgrading from Confluence 3.0 to 3.1, or 2.8 to 3.0. Minor release upgrades are ones where the 1st digit of Confluence's version number and the 1st digit after the 1st decimal place remain the same after the upgrade, for example, when upgrading Confluence 3.0 to 3.0.1.

If you have made Confluence site-wide layout customisations:

1. Go to the Confluence '**Administration Console**'. To do this:
    - Open the '**Browse**' menu and select '**Confluence Admin**'. The 'Administration Console' view will open.
2. Select '**Layouts**' under '**Look and Feel**' in the left-hand navigation panel. The decorators are grouped under **Site**, **Content** and **Export** layouts.
3. Ensure you have all your customisations available (preferably in a form which can be copied and pasted).
4. Click '**Reset Default**' next to the layout whose customisations need to be reapplied.
5. Click '**Create Custom**' next to the same layout and reapply your customisations (by copying and pasting them) into the appropriate locations within the new default layout.
6. Click the '**Save**' button.
7. Repeat this procedure from step 4 for each layout whose customisations need to be reapplied.

If you have made space-specific layout customisations:

1. Visit any page in the relevant space.

2. Go to the '**Space Admin**' tab of the Browse Space view. To do this:
    - Go to a page in the space, open the '**Browse**' menu and select '**Space Admin**'. The 'Space Administration' view will open.
    '**Space Admin**' is only displayed if you are a space administrator for that space or you are a Confluence system administrator.
3. Select '**Layout**' under '**Look and Feel**' in the left-hand navigation panel. The decorators are grouped under **Site**, **Content** and **Export** layouts.
4. Ensure you have all your customisations available (preferably in a form which can be copied and pasted).
5. Click '**Reset Default**' next to the layout whose customisations need to be reapplied.
6. Click '**Create Custom**' next to the same layout and reapply your customisations (by copying and pasting them) into the appropriate locations within the new default layout.
7. Click the '**Save**' button.
8. Repeat this procedure from step 5 for each layout whose customisations need to be reapplied.

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.

For Confluence 2.6, the `velocity.properties` file is available in the `confluence-2.6.0.jar` file. 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-2.6.0.jar file.
6. Relocate the jar file to the appropriate directory.
7. Restart Confluence.

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.

Global Templates

A template is a pre-defined page that can be used as a prototype when creating new pages. Templates are useful for giving pages a common style or format.

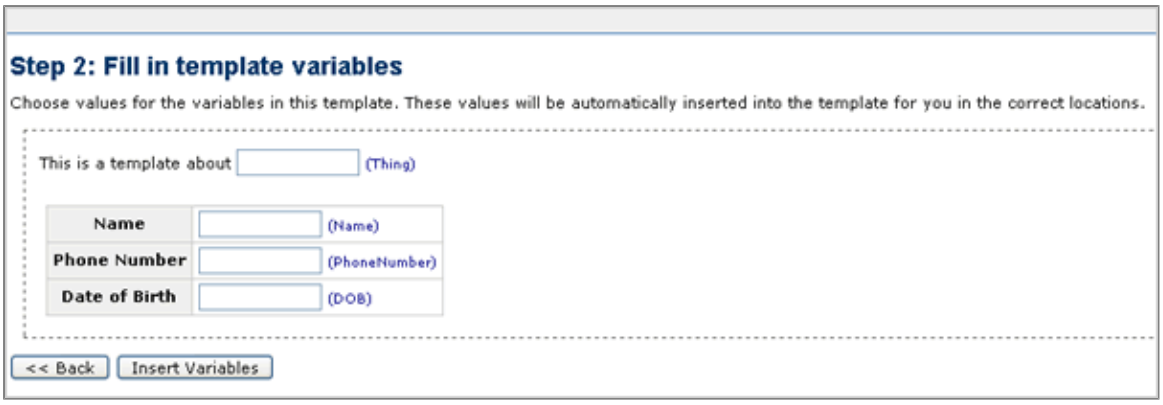
Templates are written in regular Confluence markup, using special markup to define form fields that need to be filled in.

Global Templates are defined by Confluence administrators and are available in every space across the site.

To add a global template,

1. Go to the Confluence '**Administration Console**'. To do this:
    - Open the '**Browse**' menu and select '**Confluence Admin**'. The 'Administration Console' view will open.
2. Select '**Global templates**' in the left navigation panel.
3. Click '**Add new global template**'.
4. Enter a name for your template in the '**Name**' text field and an optional description in the '**Description**' text field.
5. Using regular Confluence markup and form field markup (if you are using forms), enter content in the text-entry box as you would in any other Confluence page.
6. Click '**edit**' next to '**Labels**' if you want to use labels to categorise information. Add your labels. These labels will be included in all pages created using this template.
7. Preview and click '**Save**'.

Screenshot: A template as used to create a page



RELATED TOPICS

Creating a Page using a Template

Form Field Markup for Templates

Importing Templates

A template is a pre-defined page that can be used as a prototype when creating new pages. Templates are useful for giving pages a common style or format.

Templates are written in regular Confluence markup, using special markup to define form fields that need to be filled in.

Confluence ships with a number of templates, including the 'Charts', 'Document List' and 'Meeting Notes' templates. These templates are not available for use by default. However, if you have the appropriate permissions to access the Administration Console, you can import any of these templates to be used globally or within a specific space.

Known issue importing templates from multiple template bundles
There is a known issue preventing templates from being imported when multiple template bundles are available. Please read this KB article for further information.

Where can I find more template bundles?

- You can download template bundles from the Atlassian Plugin Exchange.
- You can also build your own custom template bundles. These are built as plugins and deployed to your Confluence instance. You can then import the templates from your custom template bundle, as described on this page. Read Creating A Template Bundle for instructions. Please note, you will need some programming knowledge to develop a custom template bundle.

Quick guide to importing a template

1. Go to the Confluence Administration Console and click '**Import Templates**'.
2. Select the templates that you want to import.
3. Choose which space to import the templates to, or whether to import them as global templates.
4. Click the '**Import**' button.

Importing a template

To import templates,

1. Log into Confluence as a System Administrator or Confluence Administrator.
2. Go to the Confluence '**Administration Console**'. To do this:
    - Open the '**Browse**' menu and select '**Confluence Admin**'. The 'Administration Console' view will open.
3. Select '**Import Templates**' in the left navigation panel. The '**Import Templates**' screen will display listing the template packages deployed to your Confluence instance (e.g. 'Default Templates Package') and the templates included in each package.
4. Select the templates to be imported by ticking the checkboxes next to the relevant template names.
   You can view a preview of the template by clicking the template name.
5. Select the import destination for the templates in the '**Import To**' dropdown. If you want the templates to only be available to a specific space, select the name of the space, otherwise select '**Global Templates**' to make the templates available to all spaces.
6. Click the '**Import**' button to import the selected templates.

If a template with the same name already exists on import, a duplicate template of the same name will be created. You will need to check each template and rename manually.

Removing the plugin that contains a template will not remove it from your Confluence instance if you have already imported it. You will need to remove it manually from the global templates/space.

Screenshots: Importing and Previewing a Template (click to view full-size images)





Related Topics

Global Templates

Working with Templates Overview

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 other languages than English.

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

```
\confluence\WEB-INF\lib\confluence-3.x.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-3.x.jar\com\atlassian\confluence\core\ConfluenceActionSupport.p`

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.
4. Restart Confluence

Common Modifications

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

Common Modifications

Task	Search For	Notes
Rename 'Dashboard'	Dashboard	The `dashboard.name` parameter has the name. To change 'Dashboard' to 'My Portal', change `dashboard.name=Dashboard` to `dashboard.name=My Portal` and update any other occurrences of the word 'Dashboard' in the instance
Modify login page text	login.	The `login.instructions` parameter has the "Enter your account details below to login to Confluence" text

Modify Keyboard Shortcuts

Confluence provides a set of keyboard shortcuts. You could customise the shortcuts by making modifications inside the `ConfluenceActionSupport.properties` file.

- To disable a particular shortcut, you can simply just comment out a respective line of code. One may like to disable the shortcut to one of the navigation links: *View, Edit, Attachments, Info* . For instance, to disable shortcut to *Attachments* one would comment out the following line:

```
1.#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.

Working With Decorator Macros

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

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

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

Macro	Usage
`#breadcrumbs()`	Draws the "You are here" breadcrumbs list, like the one found above the page name in the default template.
`#includePage(pageTitle)`	Includes a confluence page with the specified title. If you have 2 or more pages with the same title across multiple spaces, this macro will include the page belonging to the space you are currently viewing.
`#searchbox()`	Inserts a search box into the page, like the one to the far right of the breadcrumbs in the default template.
`#globalnavbar(type)`	Draws the global navigation bar, as found in the top right-hand corner of the default template. The navigation bar can be displayed in two modes:
`#globalnavbar("table")`	Displays the navigation bar in its default mode: drawn as a table of links with coloured backgrounds and mouse-over effects.
`#globalnavbar("text")`	Displays the navigation bar as series of text links separated by | characters.
`#usernavbar()`	Draws the user-specific navigation-bar. This bar contains the links to the user's profile and history, or to the login and signup pages if the user is not logged in.
`#helpicon()`	Draws the  help icon, and link to the Confluence help page.
`#printableicon()`	On pages where a printable version is available, draws the  printable page icon, linking to the printable version of the page. Otherwise, draws nothing
`#pagetitle(class)`	When you are viewing a page in a Confluence space, draws the name of the space that page is in. Otherwise, writes the word "CONFLUENCE".The "class" argument is the CSS class that the title should be drawn in. Unless you have customised your Confluence installation's CSS file, you should call this with "spacenametitle" as the class: `#pagetitle("spacenametitle")`

`#poweredby()`	Writes out the "Powered by Confluence" and Confluence version-number boilerplate found at the bottom of the default template.
`#bottomshadow()`	Draws the fading shadow-effect found at the bottom of the content area in the default template.
`#dashboardlink()`	Inserts a link to the dashboard page.

RELATED TOPICS

- Include Page Macro
- Enabling HTML macros
- Enabling the html-include Macro
- Adding, Editing and Removing User Macros

Customising a Specific Page

If you'd like to change the appearance of a specific page, you can modify the corresponding Velocity template. Here's how to find out which one:

1. Access the page. Note the name of the action. For example, the "Contact Administrators" page is `<baseUrl>/administrators.action`.
2. Browse to <confluence-install>/confluence/WEB-INF/lib/confluence-x.y.jar. Copy the file.
3. Unzip or unjar the file using a standard unzipper or the java jar utility.
4. Open xwork.xml. Search the file for the name of the action corresponding to the page you'd like to modify. You'll see an entry like:

```
1.<action name="administrators" class=
"com.atlassian.confluence.user.actions.AdministratorsAction">
2.          <interceptor-ref name="defaultStack"/>
3.          <result name="success" type="velocity">/administrators.vm</result>
4.</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.

RELATED CONTENT

- Adding a Site-Wide Banner
- Editing the Footer
- Upgrading Custom Layouts
- Customising a Specific Page
- Customising Colour Schemes
- Customising Layouts
- Customising Look and Feel Overview
- Global Templates
- Working With Decorator Macros

Customising PDF or HTML Content

To customise Confluence's PDF output, you can edit the CSS stylesheets used by the PDF exporter. See Editing the PDF Stylesheet.

To customise the HTML output, you will need to modify the file `confluence-x.y.z-jar/com/atlassian/confluence/pages/Page.htmlexport.vm`. See Editing Files within JAR Archives to learn how to repackage this file.

Customising the Dashboard

If you are a Confluence Administrator, you can customise the global dashboard, affecting the way all users will see the dashboard.

Confluence users can customise their view of the dashboard too. See the user's guide.

Sending Users to a Space Home Page instead of the Dashboard

See Configuring the Site Home Page.

Editing the Top Left-Hand Section of the Dashboard

See Editing the Site Welcome Message.

Modifying the Global Template or Layout

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

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

For example, search the Global Layout for these macros:

```
$helper.renderConfluenceMacro("{recently-updated-dashboard:dashboard|showProfilePic=true|types=page,b
$helper.renderConfluenceMacro("{favpages:maxResults=$maxFavouritePages}")
```

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

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

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

RELATED TOPICS

Customising your Personal Dashboard
Customising Look and Feel Overview

Customising the eMail Templates



Customisations to the Confluence email templates 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.

Only administrators with access to the server where Confluence is running can modify the Confluence email templates.

Process to change the email templates

1. Shut down your test instance of Confluence.
2. In the Confluence web application folder, find the file `/confluence/WEB-INF/lib/confluence-2.x.jar`.
3. Make a copy of this file as a backup.
4. Learn how to edit files within .jar archives.
5. Within the jar file, find the `/templates/email` folder. Find the appropriate file(s) within that folder.
6. Edit the file with a text editor to make the required changes. The content is mostly HTML, but has some Velocity template variables in it. See Velocity Template Overview for more information about how these work.
7. Again using the guide on editing files within .jar archives, either rejar the set of folders or drop the new files into the identical folder structure in the `WEB-INF/classes` directory.
8. Start Confluence up again and test your changes.
9. Apply the changes to your production Confluence instance.

The same process can be applied to modify most of the templates in the Confluence web application. For velocity files that are not in a jar file, you need not shut down and restart Confluence. 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.

RELATED TOPICS

- Velocity Template Overview
- Customising Layouts
- Customising Look and Feel Overview
- Modify Confluence Interface Text

Customising the Login Page

It's fairly straightforward to customise the Confluence login page, to add your own logo or custom text. This will not customise the login

process however, just what a user sees when she logs in.

> Customisations to the Confluence login page will need to be reapplied when you upgrade Confluence. Consider this before making drastic changes to the layout, and be sure to keep a list of what you have changed for your upgrade process later.

Only administrators with access to the server where Confluence is running can modify the Confluence login page.

Process to change the login page

1. Shut down your test instance of Confluence.

2. In the Confluence web application folder, 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 is mostly HTML, but has some Velocity template variables in it. See Velocity Template Overview for more information about how these work.

5. Start Confluence up again and test your changes.

6. Apply the changes to your production Confluence instance.

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.

Related topics

- Editing the Global Logo
- Velocity Template Overview
- Customising Layouts
- Customising Look and Feel Overview
- Modify Confluence Interface Text

Themes Overview

Themes are pre-defined style sets that can be applied to alter the appearance of your site. Themes allow you to personalise the 'look and feel' of Confluence. You can apply a theme to your entire Confluence site and to individual spaces. Choose a specific theme if you want to add new functionality or significantly alter the appearance of Confluence.

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

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

To look at the themes installed,

1. Go to the Confluence '**Administration Console**'. To do this:
    - Open the '**Browse**' menu and select '**Confluence Admin**'. The 'Administration Console' view will open.
2. Select '**Themes**' under 'Look and Feel' in the left-hand panel.
3. You will see a list of all installed themes.

RELATED TOPICS

- Including Cascading Stylesheets in Themes
- Applying a Theme to a Space
- Applying a Theme to a Site
- Creating a Theme
- Configuring the Documentation Theme
- Configuring the Easy Reader Theme

Administrators Guide Home Confluence Documentation Home

Applying a Theme to a Site

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

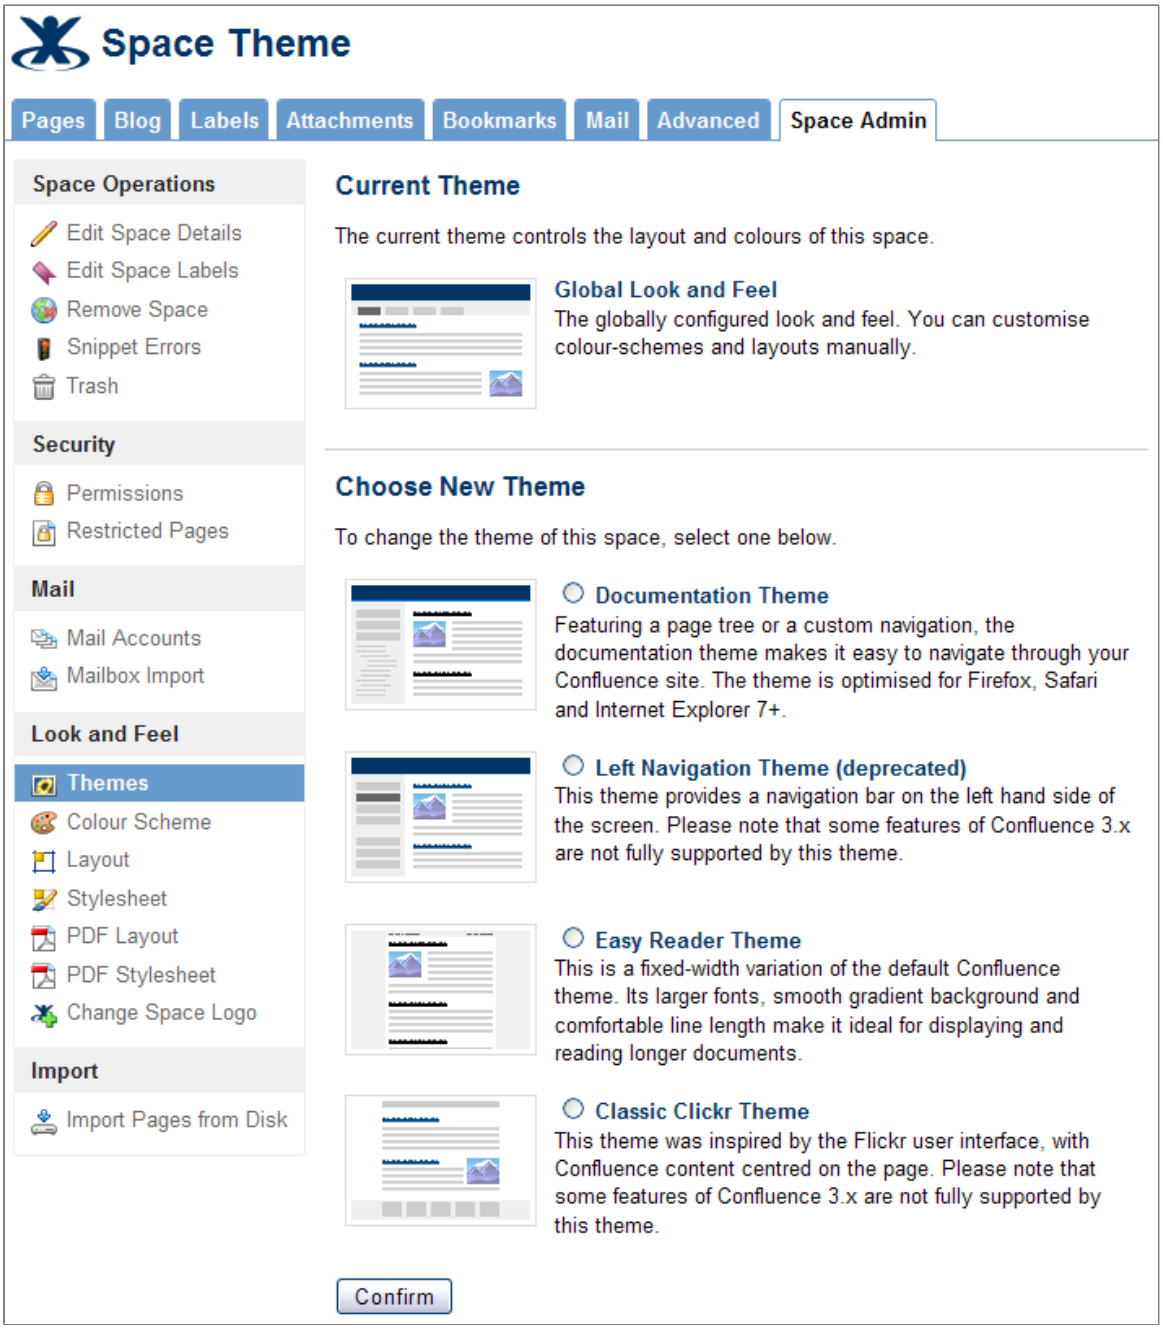
Confluence comes with a selection of themes. In addition, a site administrator can install new themes as plugins via the Confluence Administration Console. Provided that the theme is installed into your Confluence site, any space administrator can apply a theme to a space.

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

To apply a theme across the site,

1. Ensure that the theme you wish to apply has been installed as a plugin.
2. Go to the Confluence '**Administration Console**'. To do this:
    - Open the '**Browse**' menu and select '**Confluence Admin**'. The 'Administration Console' view will open.
3. Select '**Themes**' under '**Look and Feel**' in the left-hand panel.
4. The screen will display all available themes. Click a radio button to select a theme.
5. Click '**Confirm**'.

Screenshot : Applying a theme



RELATED TOPICS

- Including Cascading Stylesheets in Themes
- Applying a Theme to a Space
- Applying a Theme to a Site
- Creating a Theme
- Configuring the Documentation Theme
- Configuring the Easy Reader Theme

Administrators Guide Home Confluence Documentation Home

Customising the Left Navigation Theme

The Left Navigation theme is now deprecated
Some features of Confluence 3.x are not fully supported by this theme. Do not start using this theme now. If you are already using it, please consider moving to a different theme as soon as possible. We suggest the Documentation theme, as it provides a customisable left-hand navigation panel and additional configurable features.

Introduction

Confluence comes bundled with the *Left Navigation* theme. This theme has a navigation menu on the left-hand side of the screen, which can be customised to contain additional links, sections and even macros.

Creating a custom navigation page

By default, the left-navigation theme just displays the space icon (or profile icon for a personal space), and three menus: page operations, browse space, and add content.

To add your own content to the top, create a page in your space called 'Navigation'. Put content there that you want to appear on the left navigation menu.

A couple of tips:

- items in a bulleted list show up as normal menu items
- use 'h1' to add a section heading for your menu items.

Examples

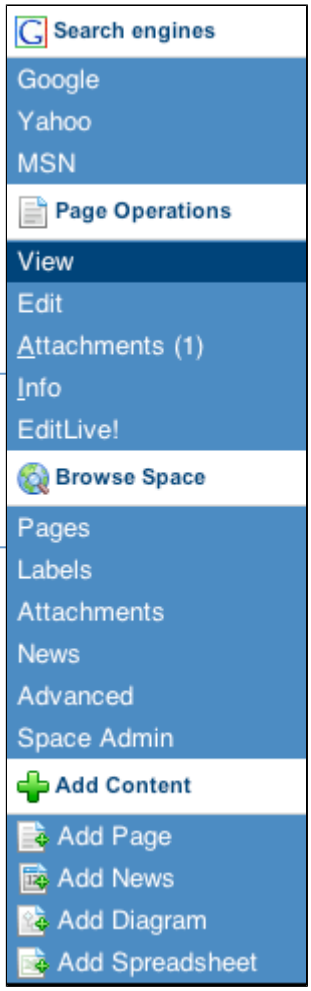
As an example, create a page called 'Navigation' with the following content:

```
h1. Search engines

* [Google|http://www.google.com]
* [Yahoo|http://www.yahoo.com]
* [MSN|http://search.msn.com]
```

This will give a left navigation menu like the image on the side of this page.

You can see another example of customised left-navigation theme on the Codegeist space with its associated Navigation page.



Custom left navigation

Alternative left-hand navigation
To insert an expandable/collapsible left-hand navigation menu, try the {pagetree} macro, supplied by the PageTree Plugin. You can follow the instructions to add the {pagetree} macro to your Confluence page layout. Note that this looks better if you use the Confluence Default theme rather than the Left Navigation theme.

RELATED TOPICS

Adding a Navigation Sidebar

Modifying Look and Feel (for themes)

Here's how you can define a new look and feel for Confluence in your theme:

1. **Layout** : Edit Confluence's layout by modifying the decorator files that are used to define it.
   - Working with Decorators
   - Velocity Template Overview
   - Configuring the atlassian.plugin.xml file to reference the decorators
2. **Colour schemes** : Configure a new colour scheme for your theme. **Optional**
   - Configuring a new colour scheme
   - Configuring the atlassian.plugin.xml file to include the new colour scheme
3. **Stylesheet** : Include a stylesheet to define your theme. **Optional**

Note that for every component you edit, you will need to configure the `atlassian-plugin.xml` which is the central configuration file for the plugin to override the default files with the new files you've created.

Layout: Working with decorators

What are decorators?
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 a very simple programming language called **Velocity**. Learn more about Velocity.

Confluence comes bundled with a set of decorator or VMD files that you can customize. Broadly these are categorised into **Site**, **Content** and **Export** decorators. These are further grouped into categories called *contexts* and under each context has various *modes* (ways of viewing the context).

To make editing easier, layout for similar screens (example: view and edit page screens) is configured through the same VMD file. So, if you want to customize how the Confluence **View Page Screen** or **Edit Page Screen** looks, you can make **both** of these changes inside one decorator file: **page.vmd**.

Decorator	Context	Mode	Comment
page.vmd	page	'view', 'edit', 'edit-preview', 'view-information', and 'view-attachments'	
blogpost.vmd	blogpost (news)	'view', 'edit', 'edit-preview', and 'remove'	We prefer to use 'news' as an end-user term; all templates and classes use 'blogpost' to indicate RSS related content
mail.vmd	mail	'view', 'view-thread' and 'remove'	
space.vmd	space-pages, space-mails, space-blogposts, space-templates, space-operations, space-administration	CONTEXT: "space-pages". MODES: "list-alphabetically", "list-recently-updated", "list-content-tree", "create-page". CONTEXT: "space-mail". MODES: "view-mail-archive". CONTEXT: "space-blogposts". MODES: "view-blogposts", "create-blogpost". CONTEXT: "space-templates". MODES: "view-templates". CONTEXT: "space-operations". MODES: "view-space-operations". CONTEXT: "space-administration". MODES: "view-space-administration", "list-permission-pages".	space.vmd handles a wide range of options, this context is accessed by clicking on 'browse space' in the default theme of Confluence (tabbed theme)

global.vmd	global	'dashboard', 'view-profile', 'edit-profile', 'change-password-profile', 'edit-notifications-profile'	
main.vmd	n/a (header and footer formatting)		main.vmd is used to control the header and footer of each page, not the page specific presentation logic

For example, if you wanted to remove the 'Attachments' tab on the view page screen, you would make this layout change in the page.vmd file - where the 'view' mode is handled (as shown below).

```
01.#*
02.     Display page based on mode: currently 'view', 'edit', 'preview-edit', 'info' and
'attachments.
03.     See the individual page templates (viewpage.vm, editpage.vm, etc.) for the setting of the
mode parameter.
04.*#
05.## VIEW
06.#if ($mode == "view")
07.
08.          <make layout modifications here>
09.
10.#elseif ...
```

Step One: Copying the decorators

The easiest way to begin configuring a new layout is by copying the default decorator files and editing them to suit your theme.

1. Go to the Confluence '**Administration Console**'. To do this:
    - Open the '**Browse**' menu and select '**Confluence Admin**'. The 'Administration Console' view will open.
2. Select **Layouts** in the left panel. This will display options to view and edit the default decorators.
3. Copy the files that you intend to modify and place them in a directory structure that makes sense to you. See example below.

Step Two: Creating a directory structure for the decorators:

You should place your decorators in a directory hierarchy which makes sense to you. We recommend that you place the atlassian-plugin.xml file at the top level of the directory structure, and then place the decorators in directories which make a meaningful division of what they do.

Here is an example:

```
1.atlassian-plugin.xml
2.com/atlassian/confluence/themes/mytheme/
3.com/atlassian/confluence/themes/mytheme/global.vmd
4.com/atlassian/confluence/themes/mytheme/space.vmd
5.com/atlassian/confluence/themes/mytheme/mail.vmd
6.com/atlassian/confluence/themes/mytheme/blogpost.vmd
7.com/atlassian/confluence/themes/mytheme/main.vmd
8.com/atlassian/confluence/themes/mytheme/page.vmd
```

Step Three: Editing the decorators

To edit the decorators, you will require knowledge of a very simple programming language called **Velocity**. Learn more about Velocity.

Decorator Macros

When editing the decorators, you will need to use **Decorator Macros** to draw complex or variable parts of the page such as menus and breadcrumbs. See Working With Decorator Macros

Theme Helper Object

When editing decorator files you will also come across a variable called **$helper** - this is the **theme helper object**.

The following table summarises what this object can do:

Behaviour	Explanation
$helper.domainName	displays the base URL of your Confluence instance on your page. This is useful for constructing links to your own Confluence pages.
$helper.spaceKey	returns the current space key or null if in a global context.

$helper.spaceName	returns the name of the current space
$helper.renderConfluenceMacro("{create-space-button}")	renders a call to a Confluence Macro for the velocity context
$helper.getText("key.key1")	looks up a key in a properties file matching `1.key.key1=A piece of text` and returns the matching value ("A piece of text")
$helper.action	returns the XWork action which processed the request for the current page.

If you are on a page or space screen you also have access to the actual page and space object by using `$helper.page` and `$helper.space` respectively.

If you want to deliver more into what other methods are available in this object, please see our API's for ThemeHelper.

Step Four: Configuring the central configuration file to reference the new decorators

How to do this is explained in Configuring the theme plugin

Working with colour schemes for themes

Configuring the colour scheme

The easiest way to configure a colour scheme is to do it dynamically from the **Administration Console** (as you would normally when you want to change the site's colour scheme online), and then express it as an xml file. This method makes it possible for you to experiment with different colours and test them out before including the colour scheme in your theme.

1. Go to the Confluence '**Administration Console**'. To do this:
    - Open the '**Browse**' menu and select '**Confluence Admin**'. The 'Administration Console' view will open.
2. Select '**Colour scheme**' in the left panel.
3. Use the colour picker to define the colours for the following UI elements:

- **Top Bar** - the bar across the top of the page that contains the breadcrumbs.
- **Space Name Text** - the text of the current space name located above the page title.
- **Heading Text** - all heading tags throughout the space.
- **Links** - all links throughout the space.
- **Borders and Dividers** - table borders and dividing lines.
- **Menu Bar Background** - background of top navigational buttons
- **Menu Bar Text** - text that appears on the menu bar
- **Menu Bar Background Highlight** - background colour of menu bar when highlighted.
- **Menu Bar Text Highlight** - menu bar text when highlighted

More information on customising colour schemes

Expressing the colour scheme as XML

Once, you have decided on the colours for the different UI elements, you will need to configure the atlassian.plugin.xml to include the new colour scheme. How to do this is explained in detail in Configuring the theme plugin.

RELATED TOPICS

Including Cascading Stylesheets in Themes

Applying a Theme to a Space

Applying a Theme to a Site

Creating a Theme

Configuring the Documentation Theme

Configuring the Easy Reader Theme

Administrators Guide Home Confluence Documentation Home

Configuring the theme plugin

Each plugin is described in its own atlassian-plugin.xml file, which specifies attributes of the plugin, including a description of each module it

contains. Once you have modified the different components to define a new look and feel for your theme, you will need to configure this file so Confluence knows where to look when overriding the default files.

The easiest way to begin is by copying the atlassian-plugin.xml from one of the default themes bundled with confluence and modifying it for your theme.

The structure of an atlassian-plugin.xml file is fairly self-explanatory. In the code segment below you will find a full example of an atlassian-plugin.xml:

```
<atlassian-plugin key="com.atlassian.confluence.themes.tabless" name="Plain Theme">
        <plugin-info>
                <description>This theme demonstrates a plain look and feel for Confluence. It
is useful as a building block for your own themes.</description>
                <version>1.0</version>
                <vendor name="Atlassian Software Systems Pty Ltd" url="
http://www.atlassian.com/"/>
        </plugin-info>

          <theme key="tabless" name="Tabless Theme" class=
"com.atlassian.confluence.themes.BasicTheme">
                    <description>plain Confluence theme.</description>
                    <layout key="com.atlassian.confluence.themes.tabless:main"/>
                    <layout key="com.atlassian.confluence.themes.tabless:global"/>
                    <layout key="com.atlassian.confluence.themes.tabless:space"/>
                    <layout key="com.atlassian.confluence.themes.tabless:page"/>
                    <layout key="com.atlassian.confluence.themes.tabless:blogpost"/>
                    <layout key="com.atlassian.confluence.themes.tabless:mail"/>
                    <colour-scheme key="com.atlassian.confluence.themes.tabless:earth-colours"/>
          </theme>

          <layout key="main" name="Main Decorator" class=
"com.atlassian.confluence.themes.VelocityDecorator"
                              overrides="/decorators/main.vmd">
                    <resource type="velocity" name="decorator"
                                              location=
"com/atlassian/confluence/themes/tabless/main.vmd"/>
          </layout>

          <layout key="global" name="Global Decorator" class=
"com.atlassian.confluence.themes.VelocityDecorator"
                              overrides="/decorators/global.vmd">
                    <resource type="velocity" name="decorator"
                                              location=
"com/atlassian/confluence/themes/tabless/global.vmd"/>
          </layout>

          <layout key="space" name="Space Decorator" class=
"com.atlassian.confluence.themes.VelocityDecorator"
                              overrides="/decorators/space.vmd">
                    <resource type="velocity" name="decorator"
                                              location=
"com/atlassian/confluence/themes/tabless/space.vmd"/>
          </layout>

          <layout key="page" name="Page Decorator" class=
"com.atlassian.confluence.themes.VelocityDecorator"
                              overrides="/decorators/page.vmd">
                    <resource type="velocity" name="decorator"
                                              location=
"com/atlassian/confluence/themes/tabless/page.vmd"/>
          </layout>

          <layout key="blogpost" name="Blogpost Decorator" class=
"com.atlassian.confluence.themes.VelocityDecorator"
                              overrides="/decorators/blogpost.vmd">
                    <resource type="velocity" name="decorator"
                                              location=
"com/atlassian/confluence/themes/tabless/blogpost.vmd"/>
          </layout>

          <layout key="mail" name="Mail Decorator" class=
"com.atlassian.confluence.themes.VelocityDecorator"
                              overrides="/decorators/mail.vmd">
                    <resource type="velocity" name="decorator"
                                              location=
"com/atlassian/confluence/themes/tabless/mail.vmd"/>
          </layout>


          <colour-scheme key="earth-colours" name="Brown and Red Earth Colours"
                                                 class=
"com.atlassian.confluence.themes.BaseColourScheme">
```

```
58.                         <colour key="topbar" value="#440000"/>
59.                         <colour key="spacename" value="#999999"/>
60.                         <colour key="headingtext" value="#663300"/>
61.                         <colour key="link" value="#663300"/>
62.                         <colour key="border" value="#440000"/>
63.                         <colour key="navbg" value="#663300"/>
64.                         <colour key="navtext" value="#ffffff"/>
65.                         <colour key="navselectedbg" value="#440000"/>
66.                         <colour key="navselectedtext" value="#ffffff"/>
67.           </colour-scheme>
68.
69.</atlassian-plugin>
```

Modifying the {{atlassian-plugin.xml}}file

We will configure this file section by section.

Plugin information

```
1.<atlassian-plugin key="com.atlassian.confluence.themes.tabless" name="Plain Theme">
2.        <plugin-info>
3.                  <description>This theme demonstrates a plain look and feel for Confluence. It
is useful as a building block for your own themes.</description>
4.                  <version>1.0</version>
5.                  <vendor name="Atlassian Software Systems Pty Ltd" url="http://www.atlassian.com/
"/>
6.        </plugin-info>
```

plugin key : Specify a key that uniquely identifies the plugin, eg. com.example.themes.dinosaur

name : Give the plugin a name.

description : Provide a short description of the plugin.

vendor : Replace the text with your information.

Theme information

```
1.<theme key="dinosaurs" name="Dinosaur Theme"
2.                  class="com.atlassian.confluence.themes.BasicTheme">
3.           <description>A nice theme for the kids</description>
4.           <colour-scheme key="com.example.themes.dinosaur:earth-colours"/>
5.           <layout key="com.example.themes.dinosaur:main"/>
6.           <layout key="com.example.themes.dinosaur:mail-template"/>
7.</theme>
```

theme key : Specify a key that uniquely identifies the theme.

class : The class of a theme must implement `com.atlassian.confluence.themes.Theme`. The `com.atlassian.confluence.themes.BasicTheme` class provided with Confluence gathers together all the resources listed within the module definition into a theme.

name : Give the theme a name. Make sure that you replace all instances of the theme name with this name.

description : Provide a short description of your theme

colour-scheme key : A theme can contain an optional `colour-scheme` element that defines which colour-scheme module this theme will use. If you are using a new colour scheme, enter its key.

layout key : A theme can contain any number of `layout` elements that define which layouts should be applied in this theme. Refer to these modules by their **module complete key** as shown above.

Referencing the decorators

You will need to add a layout entity as shown below for each of the decorators you are using. See working with decorators

```
1.<layout key="page" name="Page Decorator" class="com.atlassian.confluence.themes.VelocityDecorator"
2.                              overrides="/decorators/page.vmd">
3.                    <resource type="velocity" name="decorator"
4.                                              location=
"com/atlassian/confluence/themes/tabless/page.vmd"/>
5.</layout>
```

class : The class which each decorator, or layout, is mapped to must implement `com.atlassian.confluence.themes.VelocityDecorator`.

overrides : The layout entry must provide an `overrides` attribute which defines which decorator within Confluence is being overrridden by the theme.

Location : Specify the location of the new decorator file, so Confluence know where to look when overriding the default decorator.

It is possible for a theme to use modules that aren't in the same plugin as the theme. Just keep in mind that your theme will be messed up if the plugin that the theme depends on is removed.

Including the colour scheme

Colour schemes can be pre-configured for your theme dynamically from the **Administration Console**. See configuring colour schemes

To transport them within a theme however, they need to be expressed in the `atlassian-plugin.xml` file as shown above.

```
01.<colour-scheme key="earth-colours" name="Brown and Red Earth Colours"
02.                                   class="com.atlassian.confluence.themes.BaseColourScheme"
>
03.             <colour key="topbar" value="#440000"/>
04.             <colour key="spacename" value="#999999"/>
05.             <colour key="headingtext" value="#663300"/>
06.             <colour key="link" value="#663300"/>
07.             <colour key="border" value="#440000"/>
08.             <colour key="navbg" value="#663300"/>
09.             <colour key="navtext" value="#ffffff"/>
10.             <colour key="navselectedbg" value="#440000"/>
11.             <colour key="navselectedtext" value="#ffffff"/>
12.</colour-scheme>
```

colour-scheme key : Specify a key that uniquely identifies the colour scheme.

name : Give a name to the colour scheme.

class : The class of the colour scheme must implement `com.atlassian.confluence.themes.ColourScheme`. The `com.atlassian.confluence.themes.BaseColourScheme` class provided with Confluence sets the colours based on the module's configuration.

colour key: For each UI element, you will need to add its name and value.

see configuring colour scheme

RELATED TOPICS

- Including Cascading Stylesheets in Themes
- Applying a Theme to a Space
- Applying a Theme to a Site
- Creating a Theme
- Configuring the Documentation Theme
- Configuring the Easy Reader Theme

Administrators Guide Home Confluence Documentation Home

Including Cascading Stylesheets in Themes

Confluence allows you to integrate your own stylesheets within the theme plugin so you can have greater control over the appearance of your site. Confluence's main stylesheet is a useful reference when overriding styles and can be found in the Confluence install directory under `...confluence/styles/site-css.vm`.

CSS for Confluence 2.6
Please refer to the information about changes in Confluence 2.6.

Step One: Defining the stylesheet in the atlassian-plugin.xml

To make a stylesheet available to a decorator, you will need to reference it as a resource from within the central configuration file - `atlassian-plugin.xml`.

Here is an example where a stylesheet is being used to define the 'leftnavigation' theme:

```
1.<layout key="main" name="Main Decorator" class=
"com.atlassian.confluence.themes.VelocityDecorator"
2.overrides="/decorators/main.vmd">
3.<resource type="velocity" name="decorator"
4.location="templates/leftnavigation/main.vmd"/>
5.<resource type="stylesheet" name="leftnav.css"
6.location="templates/leftnavigation/leftnav-css.vm">
7.</resource>
8.</layout>
```

The resource parameter takes three arguments:

- **Type**: The type of resource-in this instance, 'stylesheet'.
- **Name**: The name of the stylesheet.
- **Location**: The location of the file represented in the jar archive you will use to bundle your theme.

Step Two: Using the stylesheet in the decorator

To reference the stylesheet in the decorator, you will need to use the #pluginStylesheet velocity macro.

For example, here's how you reference the leftnav.css file defined in the layout entry above:

```
1.#pluginStylesheet("com.atlassian.confluence.themes.leftnavigation:main" "leftnav.css")
```

The macro takes two arguments:

- **completePluginKey**: The complete plugin key which is constructed from the pluginkey and the layout key like this: {**pluginKey**}:{**layoutKey**}
  In the above example, `com.atlassian.confluence.themes.leftnavigation` is the key of the plugin, and `main` is the key of the layout.
- **stylesheetName**: the name of the stylesheet

If you place your stylesheet **after** the `#standardHeader` macro in the decorator, the contents of your custom stylesheet will override those in Confluence's default stylesheet.

If your stylesheet needs to reference the colour scheme, you need to use the space stylesheet macro instead:

```
1.#pluginSpaceStylesheet("com.atlassian.confluence.themes.leftnavigation:main" "leftnav.css"
$spaceKey)
```

You can then use colour scheme references in your stylesheet, similar to Confluence's stylesheets, and they will be replaced with the appropriate global or space-specific colour scheme:

```
1..navItemOver {
2.          color: $action.navSelectedTextColor;
3.}
```

RELATED TOPICS

Including Cascading Stylesheets in Themes (Confluence Docs 3.2)

Applying a Theme to a Space (Confluence Docs 3.2)

Applying a Theme to a Site (Confluence Docs 3.2)

Creating a Theme (Confluence Docs 3.2)

Configuring the Documentation Theme (Confluence Docs 3.2)

Configuring the Easy Reader Theme (Confluence Docs 3.2)

Administrators Guide Home Confluence Documentation Home

Creating a Theme

Unsure what a theme is? See the overview of themes.

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

- Get started with plugin development.
- Create a theme using the theme plugin module.

RELATED TOPICS

Including Cascading Stylesheets in Themes

Applying a Theme to a Space

Applying a Theme to a Site

Creating a Theme

Configuring the Documentation Theme

Configuring the Easy Reader Theme

Administrators Guide Home Confluence Documentation Home

Importing Data

- Importing Content from another Wiki
- Snip Snap Import

- Universal Wiki Converter
- Importing Content Into Confluence

Importing Content from another Wiki

The Universal Wiki Converter (UWC) allows you to import content from other wikis into Confluence. The Confluence Administration Console offers a link to the Universal Wiki Converter documentation and download sites.

> **You need to install and run the UWC separately from Confluence.**

The UWC is a standalone application that communicates with Confluence remotely. You cannot install the UWC directly into Confluence. Instead, download the UWC separately and run it according to the instructions below.

The UWC supports many wiki dialects. In addition, the UWC is an extensible framework, which means that developers can continue writing new conversion modules for other wikis. To see the latest list of conversions available, please refer to the UWC documentation.

- Download the latest version of the UWC.
- For information on installation and usage, see the UWC Quick Start Guide.
- For information on developing your own converter module, see the UWC Developer Documentation.
- For information about a specific wiki, including a list of currently supported wikis, see the UWC documentation.
- If you have questions or would like to share information about the UWC, please visit the UWC Forum.

Screenshot: Links from the Confluence Administration Console to the UWC



RELATED TOPICS

- Importing Content Into Confluence
- Importing Data

Snip Snap Import

The snipsnap importer allows you to import a Snip Snap XML backup file into a space in Confluence.

> **What is Snip Snap?**
> Snip Snap is a wiki used as a knowledge and content management tool. For more information, take a look at the Snip Snap home page and the Wikipedia page about Snip Snap.

Some limitations:

- Currently, attachments are not imported, and Confluence does not recognise duplicate users.
- You cannot import content into multiple spaces.

You need to have System Administrator permissions in order to perform this function.

To import a Snip Snap backup file into Confluence,

1. First, use Snip Snap to export a backup to an XML file. Now return to Confluence.
2. Go to the Confluence '**Administration Console**'. To do this:
    - Open the '**Browse**' menu and select '**Confluence Admin**'. The 'Administration Console' view will open.
3. Select '**SnipSnap Import**' in the left-hand panel.
4. Enter the location of the Snip Snap backup file in the input field displayed. You can also '**browse**' and locate the file.
5. Select a space to import the content into and click '**Save**'.

RELATED TOPICS

Importing Data

Administrators Guide Home Confluence Documentation Home

Installing Plugins and Macros

A **plugin** is an add-on to the core Confluence code, which can extend the Confluence functionality. Some plugins are shipped with Confluence, others are available for you to install yourself.

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

Installing Plugins

You need to have System Administrator permissions in order to install and configure plugins. This page introduces two methods of installing plugins:

- Via the Plugin Repository Client
- Manually

Installing and Configuring Plugins using the Plugin Repository Client

If the plugin you wish to install is listed in the Confluence Plugin Repository, you can use the Confluence Repository Client to install it. In the 'Administration' section of Confluence, click the Plugin Repository. Then find the plugin in the list and click the 'Install' link. There's more information in Installing and Configuring Plugins using the Plugin Repository Client.

Installing and Configuring Plugins Manually

If the plugin you wish to install is not listed in the Confluence Plugin Repository, you can still install it by uploading the plugin jar file to your Confluence site. In the 'Administration' section of Confluence, click the 'Plugins' link.

There's more information in Installing and Configuring Plugins Manually.

Troubleshooting

Problem Loading the Plugin Repository with IE7

There is a problem using the Plugin Repository with Internet Explorer 7. You may see an error message saying 'Object error'. In this case, the browser will just hang until you close the browser window or tab.

This problem is caused by a bug in IE7, which Microsoft have fixed in Windows Service Pack 3. Details are on our JIRA site at CONF-10837.

There are some workarounds for those who cannot install the Microsoft Service Pack. You can use a different browser instead of IE7, such as Firefox. Or you can upload the plugin jar into Confluence manually, via the 'Plugins' option in the Administration Console. See Installing and Configuring Plugins Manually.

RELATED TOPICS

- Installing and Configuring Plugins Manually
- Installing and Configuring Plugins using the Plugin Repository Client
- Plugin loading strategies in Confluence
- Removing Malfunctioning Plugins
- Enabling and Configuring Macros
    - Adding, Editing and Removing User Macros
    - Configuring a URL Whitelist
    - Configuring the userlister Macro
    - Enabling HTML macros
        - Enabling the html-include Macro
    - Troubleshooting the Gallery Macro

- Configuring the Office Connector

Installing and Configuring Plugins Manually

On this page:

- Installing a Plugin Manually
- Enabling and Disabling Plugins

This document is for **administrators** who wish to install new plugins or manage the plugins installed in their Confluence server. For an overview of how plugins work in Confluence, read the Confluence Plugin Guide.

You need to have System Administrator permissions in order to install and configure plugins.

Looking for existing plugins?
See the existing plugins and extensions written by the community in the Confluence Extensions space.

The Plugin Repository
The Plugin Repository provides an alternative way to install plugins directly from the the plugin libraries.

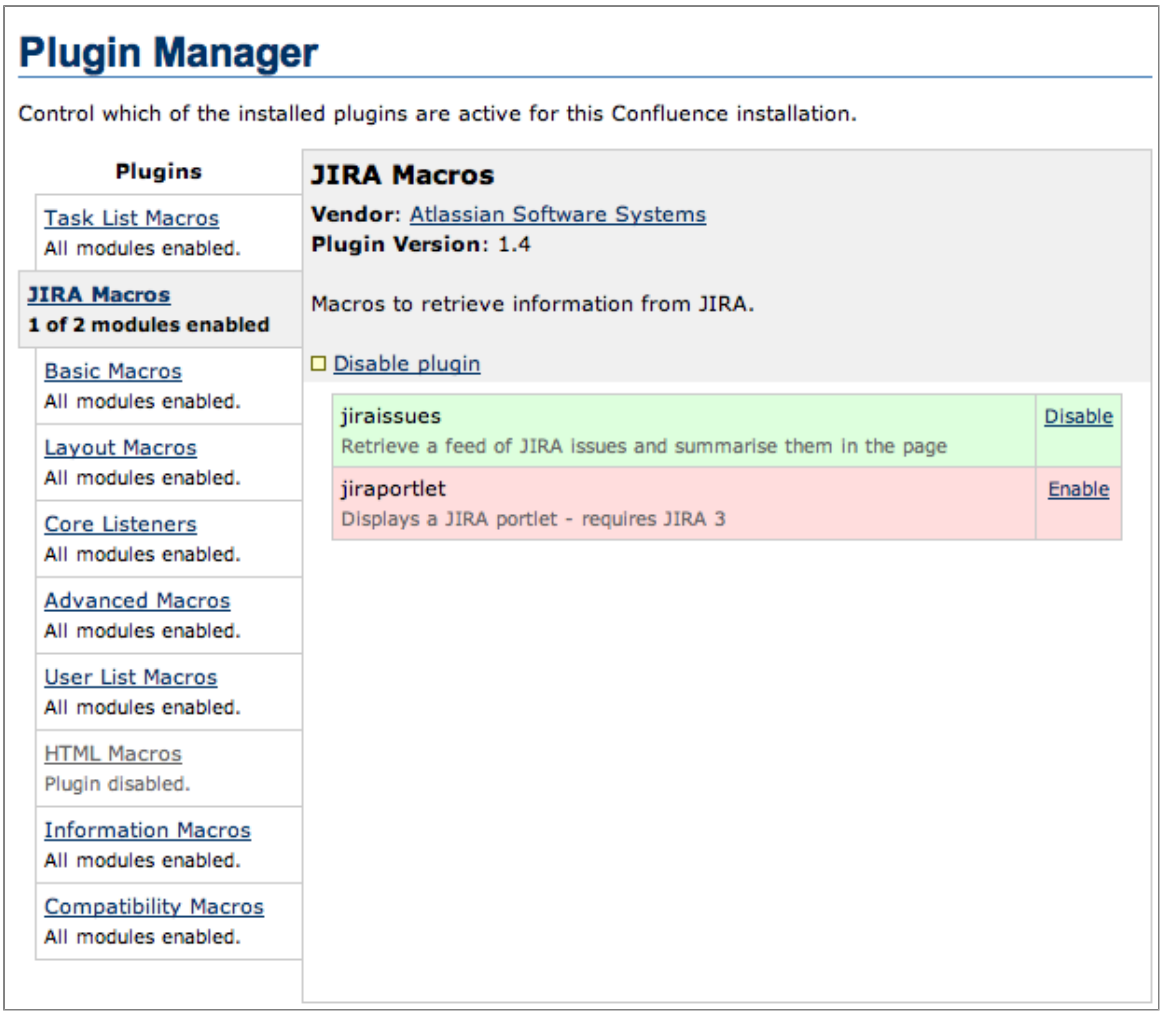
Plugin Safety
Plugins are very powerful: they can change the behaviour of almost any part of the Confluence server. This makes it **very important** that you trust a plugin before you install it. Always be aware of where (and who) a plugin comes from.

Installing a Plugin Manually

Plugins are distributed as a jar file. To install a plugin:

1. In the '**Administration**' section of Confluence, click the '**Plugins**' link.
2. Use the '**Choose file**' button to find the plugin jar you wish to install from your hard drive or network location, and select it.
3. Click '**Upload**'.
4. The plugin will be uploaded to Confluence and will be automatically installed.
5. Check the '**Plugin Administration**' screen to ensure if the plugin is available.
6. Enable the plugin if necessary. (Some plugins will be enabled by default when they are installed. Others will have to be manually enabled from the Plugin Administration screen.)

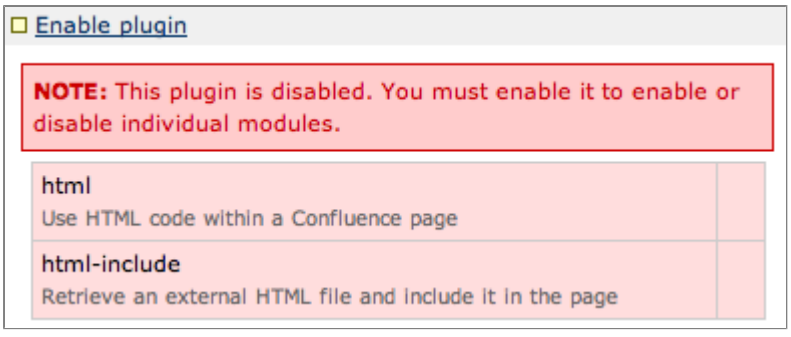
Enabling and Disabling Plugins



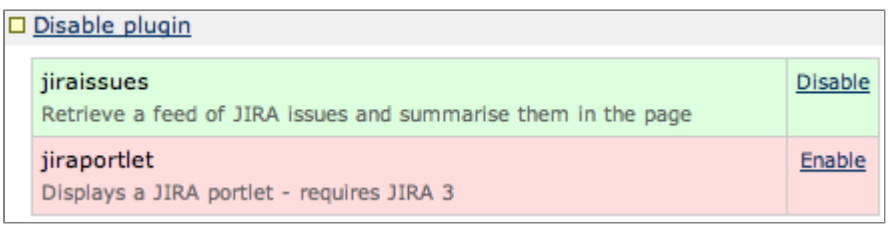
As administrator, you can enable and disable plugins, and the plugin modules which form part of each plugin. You can do this from the 'Plugins' section of the administration screen. All plugins installed in the Confluence server are listed on the left-hand side. To enable or disable a plugin or its modules, click the plugin name.

On the right-hand side, a description of the plugin is shown, including its component plugin modules.

You can enable or disable the whole plugin:



Or each module individually:





Disabling a plugin module *may* cause other modules in the same plugin to cease to function correctly. When in doubt, make sure you disable or enable the entire plugin.

RELATED TOPICS

Installing and Configuring Plugins using the Plugin Repository Client
Installing Plugins and Macros

Installing and Configuring Plugins using the Plugin Repository Client

Plugin Safety
Plugins are very powerful: they can change the behaviour of almost any part of the Confluence server. This makes it very important that you trust a plugin before you install it. To help ensure the authenticity of a plugin, always be aware of its origins and/or its creator.

The Plugin Repository provides an easy way to install and configure plugins.

- If you are using Confluence version 2.3 or later, the Plugin Repository is included as part of Confluence.
- If you are using Confluence 2.0 - 2.2.10, you can use the Plugin Repository after you install the Confluence Repository Client.
- Confluence versions prior to 2.0 cannot use the Plugin Repository, so you will need to install and configure plugins manually instead.

You need to have System Administrator permissions in order to install and configure plugins.

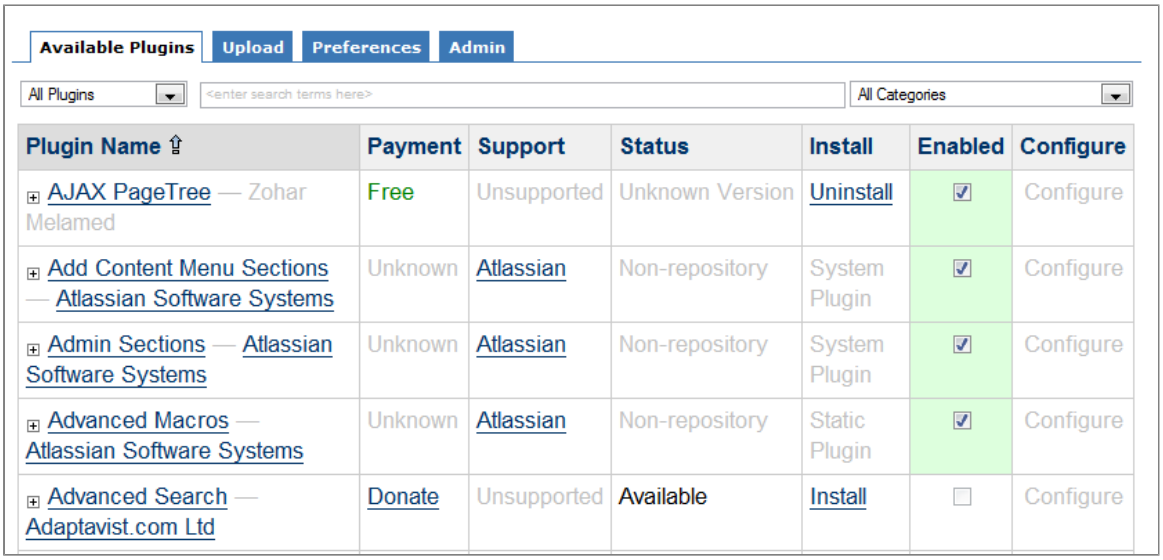
On this page:

- Using the Plugin Repository
  - Filtering the List of Plugins
  - Using the List of Plugins
- Viewing Plugin Modules and Versions
  - Administering the Plugin Repository
  - Uploading a Non-Repository Plugin
  - Obtaining More Information About Plugins in the Plugin Repository
  - Configuring a Web Proxy
- Troubleshooting

Using the Plugin Repository

Go to the '**Administration Console**' and click '**Plugin Repository**' in the left-hand panel. The following will be displayed:

Screenshot: Plugin Repository



Filtering the List of Plugins

Along the top of the page, you'll see three items which allow you to choose the plugins you want displayed:

- **Status filter** — Defaults to 'All Plugins'. Choose one of the following to limit the list of plugins displayed:
  - 'Installed Plugins' – Plugins which have been installed on your Confluence instance.
  - 'Outdated Plugins' — Plugins for which updates are available
  - 'Available Plugins' — Plugins which are available, but have not yet been installed.
  - 'Supported Plugins' — All plugins which are supported by Atlassian or a third-party.
- **Search** — Enter text in the middle textbox to quickly find what you're looking for:
  - Quickly locate plugins by searching on their title, description, vendor and other details. Just type what you are looking for and let Confluence do the rest

    - The search results are filtered by the status filter (as discussed above), so if you want to see all installed plugins from Adaptavist.com, for example, set the filter to 'Installed Plugins' and type 'Adaptavist.com' into the search box.
- **Categories list** — Filter by category:
    - Simply choose the desired category from the list and only plugins relating to that category will be displayed.
    - Note that plugins can exist in more than one category.

Using the List of Plugins

Under the filter options, the list of plugins matching the current filter settings is shown in a table.

Each column in the table shows information about a particular plugin, and allows you further configuration options:

- **Plugin Name** — Displays the name of the plugin (linked to the detailed information page)and the plugin vendor (linked to their website if applicable)
    - Click the '**+**' icon to expand the information display showing plugin description.
    - Click the '**-**' icon to hide the description again.
- **Payment** — Can be one of the following:
    - Free (self explanatory)
    - Donate (it's free, but you should consider donating to keep it that way)
    - Buy (it's commercial - click the link to show a price list and purchase online)
- **Status** — Shows the current status of this plugin in respect to your Confluence installation:
    - Installed - installed and up-to-date
    - Outdated - installed, but there are new versions available
    - Available - not installed yet
    - Non Repository - a version is installed which is not in the repository
- **Support** — Tells you who supports the plugin:
    - 'Atlassian' — The plugin is supported by Atlassian. If you have any problems, please raise a ticket at the Atlassian Support System.
    - 'Unsupported' — The plugin has been developed by a third party, not by Atlassian, and is not currently supported by Atlassian. In addition, the third party has not yet given detailed information about support arrangements. This does **not** necessarily mean that the plugin is not supported. Please refer to the plugin's home page in the Confluence Extension space or the Atlassian Plugins Exchange site.
    - The 'Support' column can also contain a link to the third-party plugin support site.
    - Read more information about supported plugins.
- **Install** — Install, upgrade or uninstall a plugin:
    - When installing or upgrading, everything is automatic (i.e. it downloads and installs for you, etc). Although the client (since 1.0.2) warns you of dependencies and (since 1.0.3) will do its best to check what has been downloaded is what you asked for - Confluence may break as if you had uploaded the plugins to the Plugin Manger yourself. Where it can, the client will error constructively allowing you to choose the best course of action for yourself. In general, **things usually work - and if they don't its a bug with the client or the metadata**.
    - If the Confluence Repository Client encounters a password request when downloading the plugin (usually case with Commercial plugins), you will be prompted for a username and password.
    - If the plugin is installed into WEB-INF (or otherwise uninstallable) it will display 'Manually Installed. where the actions would be.
- **Enabled** — If the box is ticked, the plugin is enabled, otherwise it's either disabled or partially disabled. You can enable or disable individual modules within the plugin from the plugin details screen (see later).
- **Configure** — If the plugin offers further configuration options, you can click the 'Configure' link. A new screen will open, showing the specific options offered by the plugin.

You can click the table headings to sort the table. Click a second time to reverse the sort.

Viewing Plugin Modules and Versions

When you click the name of a plugin in the plugin list, you'll be taken to the detailed view for that plugin.

General information and plugin module details are shown at the top of the display and from here you can disable or enable individual modules.

Note: Disabling a plugin module *may* cause other modules in the same plugin to cease to function correctly. When in doubt, make sure you disable or enable the entire plugin.

Near the bottom of the display a table outlines all plugin versions and shows which you have installed. Just like the plugins list, you can click the '**+**' to expand the details shown for a specific version.
This screen also allows you to quickly install, upgrade, downgrade and uninstall any version of the plugin.

Administering the Plugin Repository

There are various settings on the '**Admin**' tab.

The most important of these is the **'Data Source'** — without this, you'll see no plugins in the list and will get a fair number of errors.

The **'Earliest Plugin State'** allows you to filter the plugin list to versions at or above a specific state: Alpha, Beta, Release Candidates, Stable. If you are running in a production environment, you will usually want to set this setting to 'Stable + Release Candidates' or 'Stable Only'.

The **'Plugin Compatibility'** setting allows you to restrict the list to only show plugin versions that are specifically known to work with your version of Confluence.

In many cases, plugins will work with your version of Confluence, but they might be marked as 'unknown' compatibility if the plugin author hasn't been able to test with that specific Confluence version. As such, it's extremely useful if you find an 'unknown' version of a plugin to work (or not work) with your Confluence that you let us know (by any means possible) so that we can update the repository to reflect this.

The **'Category Visibility'** setting allows you to trim down the categories list by hiding categories that don't contain a plugin yet.

Uploading a Non-Repository Plugin

There are several plugins that are not currently listed by the Plugin Repository which need to be manually uploaded either as a file or from a remote server URL.

If you want to install such a plugin:

1. Click the '**Upload**' tab.
2. Enter either the file name or the URL of the remote server URL.
3. Click the 'Install' button.

Obtaining More Information About Plugins in the Plugin Repository

Above we have described the **'Plugin Repository'** screen in the **Administration Console**. For more information about these plugins, please refer to its page on the Atlassian Plugins Exchange site.

To find a specific plugin quickly on this site, ensure that you are in the **Plugins** section of this site and use the search tool. You should only need to enter a few keywords of the plugin's name to find it.

Configuring a Web Proxy

If the plugin repository connection times out, it's likely that there is a Web Proxy intermediary blocking the outbound connection. See Configure Web Proxy Support for Confluence for instructions how to configure.

Troubleshooting

Problem Loading the Plugin Repository with IE7

There is a problem using the Plugin Repository with Internet Explorer 7. You may see an error message saying 'Object error'. In this case, the browser will just hang until you close the browser window or tab.

This problem is caused by a bug in IE7, which Microsoft have fixed in Windows Service Pack 3. Details are on our JIRA site at CONF-10837.

There are some workarounds for those who cannot install the Microsoft Service Pack. You can use a different browser instead of IE7, such as Firefox. Or you can upload the plugin jar into Confluence manually, via the 'Plugins' option in the Administration Console. See Installing and Configuring Plugins Manually.

RELATED TOPICS

Confluence Plugin Guide

Plugin loading strategies in Confluence

The categories

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

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

Category	Description	Example
Static	**cannot be installed or upgraded without a Confluence restart**	
Core	Included with Confluence and cannot be uninstalled. The classes and plugin.xml are not bundled into plugin jars, but mixed in with Confluence source on the main classpath. Additionally, the plugin.xml definitions are not called "atlassian-plugin.xml" as they are everywhere else, but are named for the plugin e.g., "basic-macros.xml". We would like to separate some of them out and turn them into *Bundled* plugins.	Admin Sections
WEB-INF/lib	Confluence also places some plugin jars inside **WEB-INF/lib**. They are inserted during the build process by Maven. These plugins, likewise, cannot be uninstalled. In ancient times, this was the only way to install plugins, so users are also free to install plugins here. We try to discourage them from doing so, however. As of version 3.0, most of the JAR files in this directory are library dependencies, not plugins.	
Dynamic	**the opposite of static, these can be installed/upgraded while Confluence is running**	

Bundled	Bundled plugins can be administered from the Plugins console from Administration >> Plugins. You can upload or disable them there. *Bundled* plugins are included in a zip of jars called **atlassian-bundled-plugins.zip** which is on the main Confluence classpath, in a resources directory - `<confluence-install>/confluence/WEB-INF/classes/com/atlassian/confluence/setup`. At Confluence startup, they are extracted and copied into the **$CONFLUENCE_HOME/bundled-plugins** directory, from whence they are loaded. To remove a bundled plugin (you shouldn't normally have to do this), remove the plugin from the atlassian-bundled-plugins.zip file and the bundled-plugins directory, otherwise Confluence will just put it back in place on the next startup. In versions later than 2.6, you'll have to recreate the .jar file (if the jar file is from the lib folder) or recreate the zip folder(if its in the classes folder). Bundled plugins can be upgraded or disabled.	Office Connector
Uploaded	Installed by the user via the plugin repository or the Plugin Manager page. These plugins are stored in the database and then copied to the **$CONFLUENCE_HOME/plugins-cache** folder on each Confluence node.	could be anything

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

Use of the categories in Confluence

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

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

All of the *Core* plugins once were labeled as "System", but it seems the practice has faded over time. If a plugin is designated as "System", then it will not show up in the Plugin Manager page in Confluence and thus cannot be enabled/disabled. However, it *will* show up in the Plugin Repository Client, where it can be disabled; allowing disabling there is probably incorrect behavior.

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

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

Upgrading plugins

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

RELATED TOPICS

Removing Malfunctioning Plugins

Removing Malfunctioning Plugins

Confluence goes to some lengths to prevent itself being unusable due to a problematic plugin. However, sometimes a plugin will manage to do this anyway. This page describes what to do if a plugin cannot be disabled or deleted from the Administration console (from `Administration >> Plugins`).

Plugin Loading Strategies

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

Check these locations when troubleshooting plugin loading issues.



Check the How to display classpath utility for tips on what's loading, and the Knowledge Base Article on plugin malfunctioning.

Deleting a plugin from the Database

To remove a plugin from Confluence when Confluence is not running:

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

```
1.select plugindataid, pluginkey,filename,lastmoddate from plugindata;
```

3. After you have found the plugindataid for the offending plugin, please run the following:

```
1.delete from plugindata where plugindataid='XXXXXX';
```

   where XXXXXX is the plugindataid value.

4. Restart Confluence.

Deleting a Bundled Plugin

Bundled plugins can be administered from the Plugins console from Administration >> Plugins. You can upload or disable them there.

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

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

Usually, the problem is that an old plugin is getting loaded along with the properly bundled one, but if you need to remove a bundled plugin, check Plugin loading strategies in Confluence.

Enabling and Configuring Macros

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

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

RELATED TOPICS:

- Adding, Editing and Removing User Macros
- Configuring a URL Whitelist
- Configuring the userlister Macro
- Enabling HTML macros
    - Enabling the html-include Macro
- Troubleshooting the Gallery Macro

Adding, Editing and Removing User Macros

User macros allow you use the Confluence web interface to create simple formatting macros that people can use on their Confluence pages.

You need to have System Administrator permissions in order to perform this function.

To add a user macro,

1. Go to the Confluence '**Administration Console**'. To do this:
    - Open the '**Browse**' menu and select '**Confluence Admin**'. The 'Administration Console' view will open.
2. Click '**User Macros**' in the left-hand panel.
3. Click '**Create a User Macro**' at the top of the list of macros.
4. Enter the macro details as explained in the Confluence development guide.
5. Click the '**Add**' button.

To edit a user macro,

1. Go to the Confluence '**Administration Console**'. To do this:
    - Open the '**Browse**' menu and select '**Confluence Admin**'. The 'Administration Console' view will open.
2. Select '**User Macros**' in the left-hand panel. This will list the currently configured user macros.
3. Click '**Edit**' next to the relevant macro.
4. Update the macro details as explained in the Confluence development guide,
5. Click the '**Save**' button.

To remove a user macro,

1. Go to the Confluence '**Administration Console**'. To do this:
    - Open the '**Browse**' menu and select '**Confluence Admin**'. The 'Administration Console' view will open.
2. Select '**User Macros**' in the left-hand panel. This will list the currently configured user macros.
3. Click '**Remove**' next to the relevant macro.

RELATED TOPICS

Include Page Macro

Enabling HTML macros

Enabling the html-include Macro

Adding, Editing and Removing User Macros

Administrators Guide Home Confluence Documentation Home

Configuring a URL Whitelist

The RSS and HTML-include macros are used to include content dynamically from other websites onto a Confluence page. The included content may possibly be malicious or harmful to your Confluence instance.

Confluence administrators can set up a list of trusted URLs, thus limiting the locations from which the RSS macro and the HTML-include macro can draw their content.

The form below allows you to define specific URLs and/or URL patterns which are trusted, or to allow inclusion from all URLs without restriction.

To configure the URL whitelist,

1. Go to the Confluence '**Administration Console**'. To do this:
    - Open the '**Browse**' menu and select '**Confluence Admin**'. The 'Administration Console' view will open.
2. Select '**Configure Whitelist**' in the left-hand panel.
3. The 'Configure Whitelist' screen will appear, as shown in the screenshot below.
4. Select one of the radio buttons as follows:
    - **Allow all domains** — There will be no restrictions to the content which can be included onto your Confluence pages.
    - **Restrict to listed domains** — Confluence will allow content from trusted URLs only. When you select this option, a textbox will open allowing you to enter specific URLs and/or URL patterns. Enter one or more URLs, each on its own line. You can enter the full URL, or use the pattern matching rules described below.
5. Click '**Save**'.

Screenshot: Configuring a URL whitelist



URL Pattern-Matching Rules

Enter one URL or URL pattern per line. You can enter a full URL or use pattern-matching as described below:

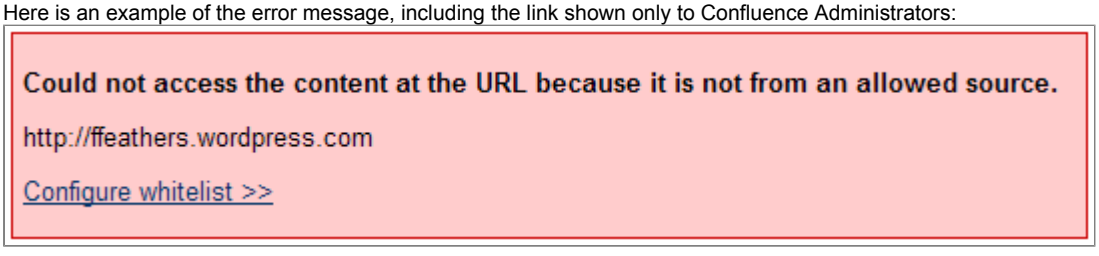
- If the rule starts with an equals sign (=), only the exact URL following the '=' will be allowed.
- If the rule starts with a slash (/) then the whole rule will be treated as a regular expression.
- Otherwise, any asterisk (*) will be treated as a wildcard to match one or more characters.

What Happens to a Page Containing a Disallowed URL?

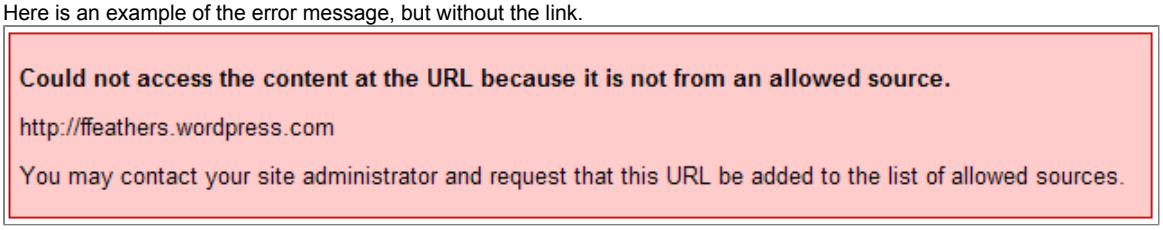
A user can add the RSS macro or the HTML-include macro to a Confluence page. The macro code includes a URL from which the content is drawn. When the page is displayed, Confluence will check the URL against the whitelist. If the URL is not allowed, Confluence will display an error message on the page.

The error message says that Confluence "could not access the content at the URL because it is not from an allowed source" and displays the offending URL. If the person viewing the page is a Confluence Administrator, they will also see a link to the Administration page where they can configure the URL whitelist.

Here is an example of the error message, including the link shown only to Confluence Administrators:



Here is an example of the error message, but without the link.



Notes

Some things to be aware of:

- By default, the RSS and HTML-include macros are disabled in Confluence. A System Administrator can enable them on the '**Plugins** ' screen of the Confluence Administration Console.

- A user who has the 'Confluence Administrator' permission, but not necessarily the 'System Administrator' permission, can configure the URL whitelist (for the HTML-include and RSS macros).

RELATED TOPICS

Enabling HTML macros
RSS Feed Macro
HTML Include Macro

Administrators Guide Home Confluence Documentation Home

Configuring the userlister Macro

The userlister macro has an optional 'online' parameter. If the '**User Listener**' plugin is configured to allow this feature, then the page author can specify 'online=true' to show a list of all online users.

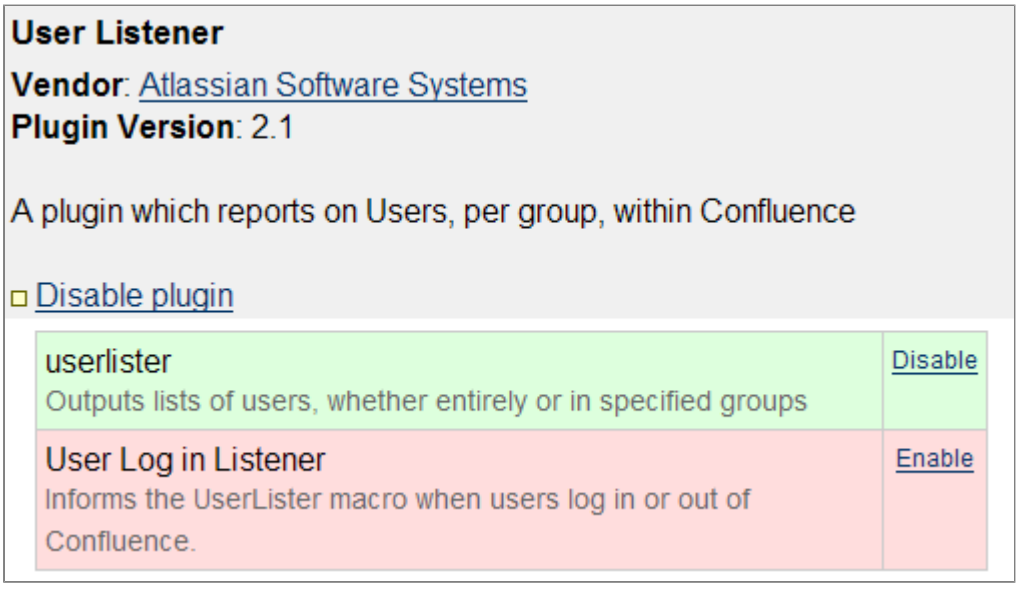
You need to have System Administrator permissions in order to perform this function.

To enable the 'online' filter in the `userlister` macro,

1. Go to the Confluence '**Administration Console**'. To do this:
    - Open the '**Browse**' menu and select '**Confluence Admin**'. The 'Administration Console' view will open.
2. Select '**Plugins**' in the left-hand panel. This will list the currently installed plugins.
3. Scroll down and click the '**User Listener**' link. The User Listener plugin panel will appear at the top of the screen.
4. Enable the '**User Log In Listener**' module by clicking the '**Enable**' link on its right.
5. Restart Confluence.

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

Screenshot: Enabling the User Log In Listener



RELATED TOPICS

Userlister Macro
Enabling and Configuring Macros

Enabling HTML macros

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

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

CAUTION: Including unknown HTML inside a webpage is dangerous. Because HTML can contain active scripting components, it would be possible for a malicious attacker to present a user of your site with script that their web browser would believe came from you. Such code could be used, for example, to steal a user's authentication cookie and give the attacker their Confluence login password.
By default, the HTML macros are disabled. You should only turn on these macros if you trust all your users not to attempt to exploit them.

 You need to have System Administrator permissions in order to perform this function.

To enable the HTML macros,

1. Go to the Confluence '**Administration Console**'. To do this:
    - Open the '**Browse**' menu and select '**Confluence Admin**'. The 'Administration Console' view will open.
2. Select '**Plugins**' in the left-hand panel. This will display the installed plugins active for this Confluence installation.
3. Click' '**HTML macros**', then click '**Enable Plugin**'

RELATED TOPICS

 Include Page Macro

 Enabling HTML macros

 Enabling the html-include Macro

 Adding, Editing and Removing User Macros

 Administrators Guide Home  Confluence Documentation Home

Enabling the html-include Macro

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



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

Enabling the HTML Macros

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

 You need to have System Administrator permissions in order to perform this function.

To enable the HTML macros,

1. Go to the Confluence '**Administration Console**'. To do this:
    - Open the '**Browse**' menu and select '**Confluence Admin**'. The 'Administration Console' view will open.
2. Select '**Plugins**' in the left-hand panel. This will display the installed plugins active for this Confluence installation.
3. Click' '**HTML macros**', then click '**Enable Plugin**'.

RELATED TOPICS

HTML Include Macro

 Include Page Macro

 Enabling HTML macros

 Enabling the html-include Macro

 Adding, Editing and Removing User Macros

 Administrators Guide Home  Confluence Documentation Home

Troubleshooting the Gallery Macro

Gallery Macro

The full list of parameters is shown in the following table.

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 i.e. include all the pictures on the page.	The gallery will ignore any pictures specified by `exclude=picture file name` i.e. they will not be included in the gallery. You can specify more than one picture, separated by commas. Example: `exclude=my picture.png,my picture2.gif`
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. Format is `include=picture file name`. You can specify more than one picture, separated by commas. Example: `include=my picture.png,my picture2.gif`
Use Images in these Pages (`page`)	If no page is specified, the gallery displays the images attached to the page containing the macro.	Specify the title of the page which contains the images you want displayed. If the page is in the same space as the page containing the macro, use the format `page=My Page Name`. To specify a page in a different space, use `page=SPACEKEY:My Page Name`, such as `page=DOC:Gallery Macro`
Reverse Sort (`reverseSort`)	Nothing, i.e. sort order is ascending	Used in conjunction with 'sort' parameter above. Use 'reverseSort' to reverse the sort order, from ascending to descending.
Sort Images By (`sort`)	None i.e. the sort order is unspecified and therefore unpredictable.	Specify an attribute to sort the images by that attribute. Sort order is ascending, unless you specify the 'reverseSort' 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.



If the actual name of an attachment file or page contains a comma, you can refer to it in the **exclude**, **include**, or **page** parameters above by enclosing it in single or double quotes, for example `"this,that.jpg", theother.png`.

For more information, refer to Gallery Macro.

Troubleshooting

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

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

Also see CONF-1737



Please note that gallery-ext.jar is available at CONF-6620

Configuring the Office Connector

The Office Connector is a Confluence plugin that allows Confluence users to interact with Microsoft Office and Open Office in various ways. You can display content from Office documents on a wiki page, import content from an Office document into Confluence, and edit a Confluence page in Microsoft Word. Please refer to the User Guide for details of these interactions.

A System Administrator can enable or disable parts of the Office Connector and can configure options as described below.

On this page:

- Enabling and Disabling the Office Connector and its Modules
- Configuring the Office Connector Options

Enabling and Disabling the Office Connector and its Modules

A System Administrator can install, enable or disable plugins and plugin modules. You can read a general overview in Installing Plugins and Macros.

The Office Connector is bundled with Confluence 2.10 and later, so you should not need to install it. But you may wish to enable or disable some of its modules.

To enable or disable the Office Connector and its modules,

1. Select '**Plugins**' in the left-hand panel of the Confluence Administration Console.
2. Search the page for 'Office Connector plugin' and select the link.
3. The '**Office Connector plugin**' panel will appear near the top centre of the page, as shown in the screenshot below.
4. Now you can do one of the following:
    - **Configure plugin** — This will take you to the separate plugin configuration screen described below.
    - **Disable plugin** — Click this link if you want to disable all modules of the plugin, but leave the plugin installed on your Confluence site.
    - **Uninstall plugin** — Click this link if you want to remove the Office Connector permanently from your Confluence site. To restore it at a later date, you will need to re-install it from the Confluence Plugin Repository.
    - You can also enable or disable one or more of the Office Connector modules, as described in the table below.

Screenshot: Enabling the Office Connector plugin and its modules



The following modules are available for the Office Connector plugin:

Module	Description
OC Settings Manager	Component to read and write persistent settings for the Office Connector.
Slide Cache Manager	Component to cache slide-based conversions when displaying PowerPoint and PDF documents.
Html Cache Manager	Component to cache HTML-based conversions when displaying Word and Excel documents.
File Cache Cleanup Job	This module is a recurring task that cleans up the Office Connector file cache.
File Cache Cleanup	This module is the trigger for the File Cache Cleanup Job.
Office Connector administration link	This module supplies the 'Office Connector Configuration' link in the left-hand panel of the Confluence Administration Console. The link gives access to the plugin configuration screen described below.

Link for previewing a search result	This modules supplies the 'View' link which appears next to attachments displayed in search results, where the attachment is an Office document.
Link for previewing an attachment	This modules supplies the 'View' link which appears next to attachments displayed on the 'Attachments' view of a page, where the attachment is an Office document.
Edit in Word UI on page actions	This module supplies an 'Edit in Word' icon which appears in older versions of Confluence, at the top right of the page with other icons such as the 'Browse Space' and 'Add Page' icons. Not relevant to Confluence 2.10 or later, except for custom themes. You can configure the location of this option on the Office Connector configuration screen, as described below.
Edit in Word UI on drop down menu	This modules supplies the 'Edit in Word' link which appears in the Confluence 'Tools' dropdown menu.
editinwordlink	This modules supplies the {editinwordlink} macro. See Edit in Word Link Macro.
viewfile	This module supplies the {viewfile} macro. See View File Macro.
viewdoc	This module supplies the Word document component of the {viewfile} macro.
viewxls	This module supplies the Excel document component of the {viewfile} macro.
viewppt	This module supplies the PowerPoint document component of the {viewfile} macro.
viewpdf	This module supplies the PDF document component of the {viewfile} macro.
editgrid	This module is used to migrate editgrid users to the Office Connector.
Edit in Word UI on page tabs	This module supplies an 'Edit in Word' tab which appears in older versions of Confluence, next to the 'View', 'Edit', 'Attachments' and 'Info' tabs. Not relevant to Confluence 2.10 or later, except for custom themes. You can configure the location of this option on the Office Connector configuration screen, as described below.
Import Word UI on page tabs	This module supplies a 'Doc Import' tab which appears in older versions of Confluence, next to the 'View', 'Edit', 'Attachments' and 'Info' tabs. Not relevant to Confluence 2.10 or later, except for custom themes.
Import Word UI on drop down menu	This modules supplies the 'Doc Import' link which appears in the Confluence 'Tools' dropdown menu.
Edit in Office javascript resource	This module contains the javascript resources for launching the desktop applications for editing Office documents.
Office Connector Servlet	This module allows Confluence users to edit their Confluence pages in Microsoft Word. It performs the conversion to and from Word.
Office Authenticator Filter	This module authenticates HTTP requests from Office applications.
PPT slide web service	This module allows Confluence users to view a PowerPoint presentation on a wiki page. It provides the slide images to the Flash control which displays the slides on the wiki page.
DOC and XLS image cache web service	This module is required if Confluence users want to view a Word document or an Excel spreadsheet on a wiki page. It allows images to be stored in a cache on the server, so that they can be retrieved when the browser renders the HTML page.
Office Connector Actions	This module must be enabled if the Office Connector is used.

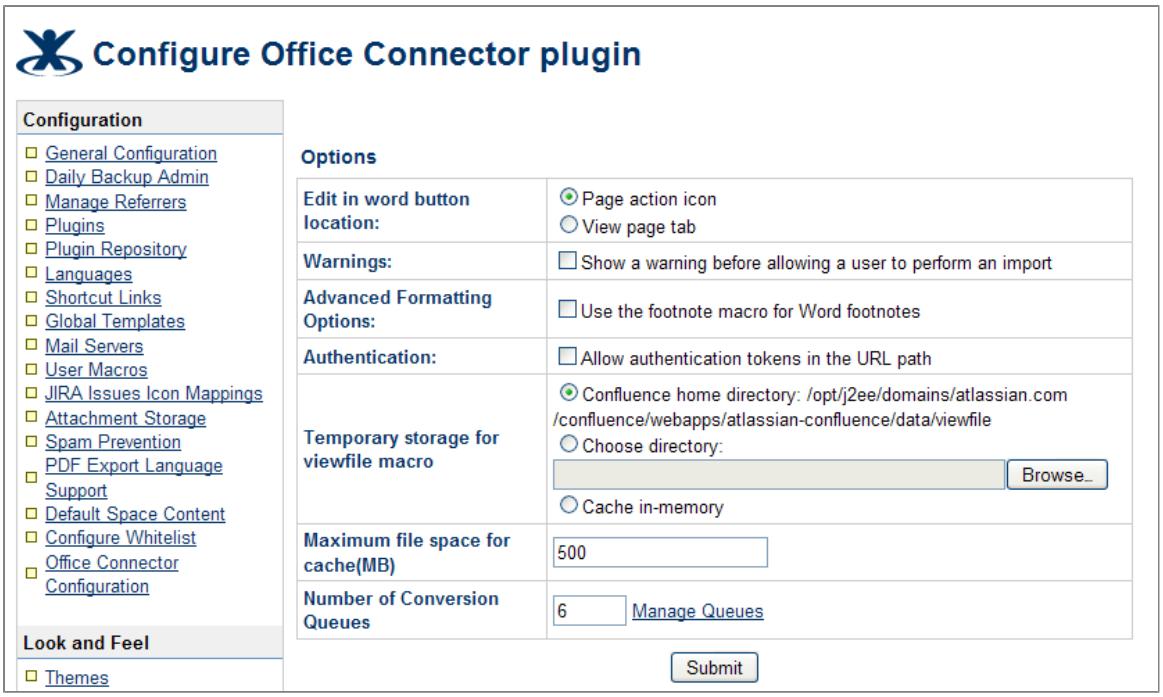
Configuring the Office Connector Options

A Confluence administrator can set the options described below, to determine the behaviour of the Office Connector on your Confluence site.

To set the configuration options for the Office Connector,

1. Select '**Office Connector Configuration**' in the left-hand panel of the Confluence Administration Console.
2. The '**Configure Office Connector plugin**' screen will appear. Set the configuration options as described in the table below.

Screenshot: Configuring the Office Connector options



The configuration options are described in the table below:

Option	Default Value	Description
Edit in word button location	'Tools' menu.	This setting determines the location of the **'Edit in Word'** option on the Confluence menus or screens. You may want to change the location if you are using a theme which does not support the default location. Note that you must ensure that the relevant plugin modules are enabled, as described above . Available settings are: • Page action icon — An icon at the top right of the page with other icons such as the 'Browse Space' and 'Add Page' icons. • View page tab — A page tab, next to the 'View', 'Edit', 'Attachments' and 'Info' tabs.
Warnings: Show a warning before allowing a user to perform an import	Disabled	If this option is enabled, the user will receive a warning when importing a Word document. The warning will tell the user when they are about to overwrite existing content.
Advanced Formatting Options: Use the footnote macro for Word footnotes	Disabled	If this option is enabled, a Confluence page created from an imported Word document will use the {footnote} macro from Adaptavist to render any footnotes contained in the document. Note that you will need to install the Footnotes Plugin onto your Confluence site. For more information about this plugin and macro, please refer to the Footnotes Plugin.
Authentication: Allow authentication tokens in the URL path	Disabled	If this option is enabled, the Office Connector will use authentication tokens in the URL.
Temporary storage for viewfile macro	The Confluence Home directory.	The {viewfile} macro will cache data temporarily. This option allows you to set the location of the cache. If you are running in a clustered environment, we recommend that you use the Cache in-memory setting.
Number of Conversion Queues	6	This is the maximum number of threads used to convert PowerPoint 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 the **Manage Queues** link to view attachments that are still pending conversion.

RELATED TOPICS

Office Connector Prerequisites
Office Connector Limitations and Known Issues
Working with the Office Connector
Installing Plugins and Macros

Operating Large or Mission-Critical Confluence Installations

This page gives guidelines for operational management teams who are responsible for a large Confluence installation, or for a Confluence installation which is crucial to the business of their organisation.

On this page:

- Introduction to this Page
  - Motivation for Presenting these Guidelines
  - Who should Read these Guidelines?
- Requirements of Large or Mission-Critical Confluence Installations
  - Dedicated Hardware for Confluence
  - Dedicated Qualified Staff
    - Operations Team with General Administrators
    - Network Staff
    - Database Staff
    - Developers
  - Constant Monitoring of Production Systems
  - Adherence to Strict Upgrade Procedures
  - Testing of Upgrades before Production Implementation
  - Load-Testing Environments
  - Tuning
    - Optimising your System
    - Limiting Third-Party Plugins
    - Selecting and Tuning your JVM
    - Customising Confluence to Optimise Performance

Introduction to this Page

Motivation for Presenting these Guidelines

Most Confluence installations start off small. Ten people in an early-adoption department use it for a couple of weeks. Everything works well and the good news starts spreading. Adoption increases throughout the organisation. More and more people use the wiki, and more and more rely on Confluence being up and running. After a while even the CEO starts blogging. And then a system outage occurs.

Now what?

Wikis like Confluence often grow into mission-critical applications within just a few months. Often adoption is so fast that IT departments haven't had the time to scale up their support.

We have assembled some requirements to help you make sure that your installation of Confluence can be mission critical. There are no surprises to be found here — all of the requirements would apply to any other piece of software that is mission critical within your organisation.

Who should Read these Guidelines?

The guidelines **do not apply to you** if you are using Confluence with just a few dozen users, and no one really minds if Confluence is down for a couple of hours because your database has crashed.

But if any one of the following applies to you, then these guidelines are **a must read for you**!

- The wiki has become your organisation's documentation base.
- Your users can't work properly when Confluence is down.
- Your boss or customer threatens to terminate your contract if you don't meet a strict service level agreement (SLA), such as 99.9% availability.

Requirements of Large or Mission-Critical Confluence Installations

Dedicated Hardware for Confluence

In a small work group with a few dozen or even hundreds of users, your Confluence installation can happily share the CPUs, memory and disks with other low-profile applications and a database.

But with thousands or even tens of thousands of users, you need dedicated hardware that runs Confluence and nothing else, and it needs to be fast hardware with plenty of RAM. While you *can* run Confluence in a virtualised environment such as VMware, we suggest you don't do it for mission-critical or high-load installations unless you are a real expert in virtualisation. Otherwise your other VMs might have performance problems which propagate to Confluence.

If you experience database-related problems, you should consider moving the Confluence database to a dedicated machine. Confluence itself can run queries that impact the performance of other applications, and other application problems or scheduled tasks can have an adverse affect on the usability of Confluence.

Dedicated Qualified Staff

If your Confluence installation is mission critical and your service level agreements require 24/7 up time, you need to be able to pinpoint problems quickly. You need qualified staff, dedicated to looking after Confluence, who are available during business hours and possibly beyond.

If you require assistance from the Atlassian Support team, you may need to answer some pretty technical questions to help us diagnose what is going on in your systems. Also keep in mind that Atlassian support assists you in finding problems in Confluence, but we can't help you administer your systems.

In particular, we recommend that you have dedicated staff in the roles listed below.

Operations Team with General Administrators

If your organisation relies on Confluence being up and running around the clock with very little downtime, you need people who can set up, maintain, tune and improve your Confluence installation. This requires at least one person, but ideally you will have a team of operational engineers.

If your wiki is mission critical, chances are that other IT systems within your organisation have already made it necessary to have such an operations team. So you will probably not need to hire someone specifically to administrate Confluence. But it is vital that supporting and maintaining Confluence is added to the list of responsibilities of that operations teams, and that you can get them to troubleshoot and analyse Confluence at short notice.

If problems arise and you need to contact Atlassian Support, these engineers will be our first point of contact. We may ask them to provide details of log files, application-server settings, monitoring systems, and so on.

Network Staff

If Confluence is mission critical for large numbers of users, it is vital that you have dedicated network staff available to track down problems when they arise.

A mission-critical installation will usually be used by hundreds or even thousands of users, and you don't want to keep them waiting because a network card breaks, or because someone has made an undocumented change to the network and you don't have an expert around who can figure it out.

Again, this only applies to mission-critical systems. If you use Confluence for less critical collaboration and knowledge sharing, and a broken network cable causing a day's downtime is no major catastrophe, then you will not need dedicated networking staff.

Database Staff

If Confluence is mission critical for a large number of users, you need an experienced database administrator (DBA) available to troubleshoot database performance issues and other potential problems. It is dangerous not to have an experienced full-time DBA at hand at short notice when running a mission critical application. While small installations of Confluence basically work 'out of the box', any system that involves high load or high-availability requirements needs continual monitoring, optimising and fine tuning of the Confluence database. Database monitoring is no trivial task — it's not something that anyone can learn quickly.

Developers

You may have decided to customise Confluence by changing its source-code, or by writing your own plugins. If your server is mission-critical, you must nominate staff who will be responsible for that code, and they must be up for the task. Otherwise you might end up in a situation in which your server experiences downtimes because of custom code is broken, or does not work with a newer version of Confluence anymore, but you can't fix the problem because no one knows how the customized code works, and you can't uninstall it either because it has become critical for your Confluence usage pattern. Keep good track of changes, and have someone available to jump into action if there is a problem Don't let the summer intern write mission-critical plugins, unless you have more senior staff to maintain that code as long as it is in use.

Constant Monitoring of Production Systems

You will need to monitor your production systems constantly.

When the wiki is the lifeblood of your organisation, you need know exactly what is going on inside, so that you can plan for future needs and analyse potential bottlenecks.

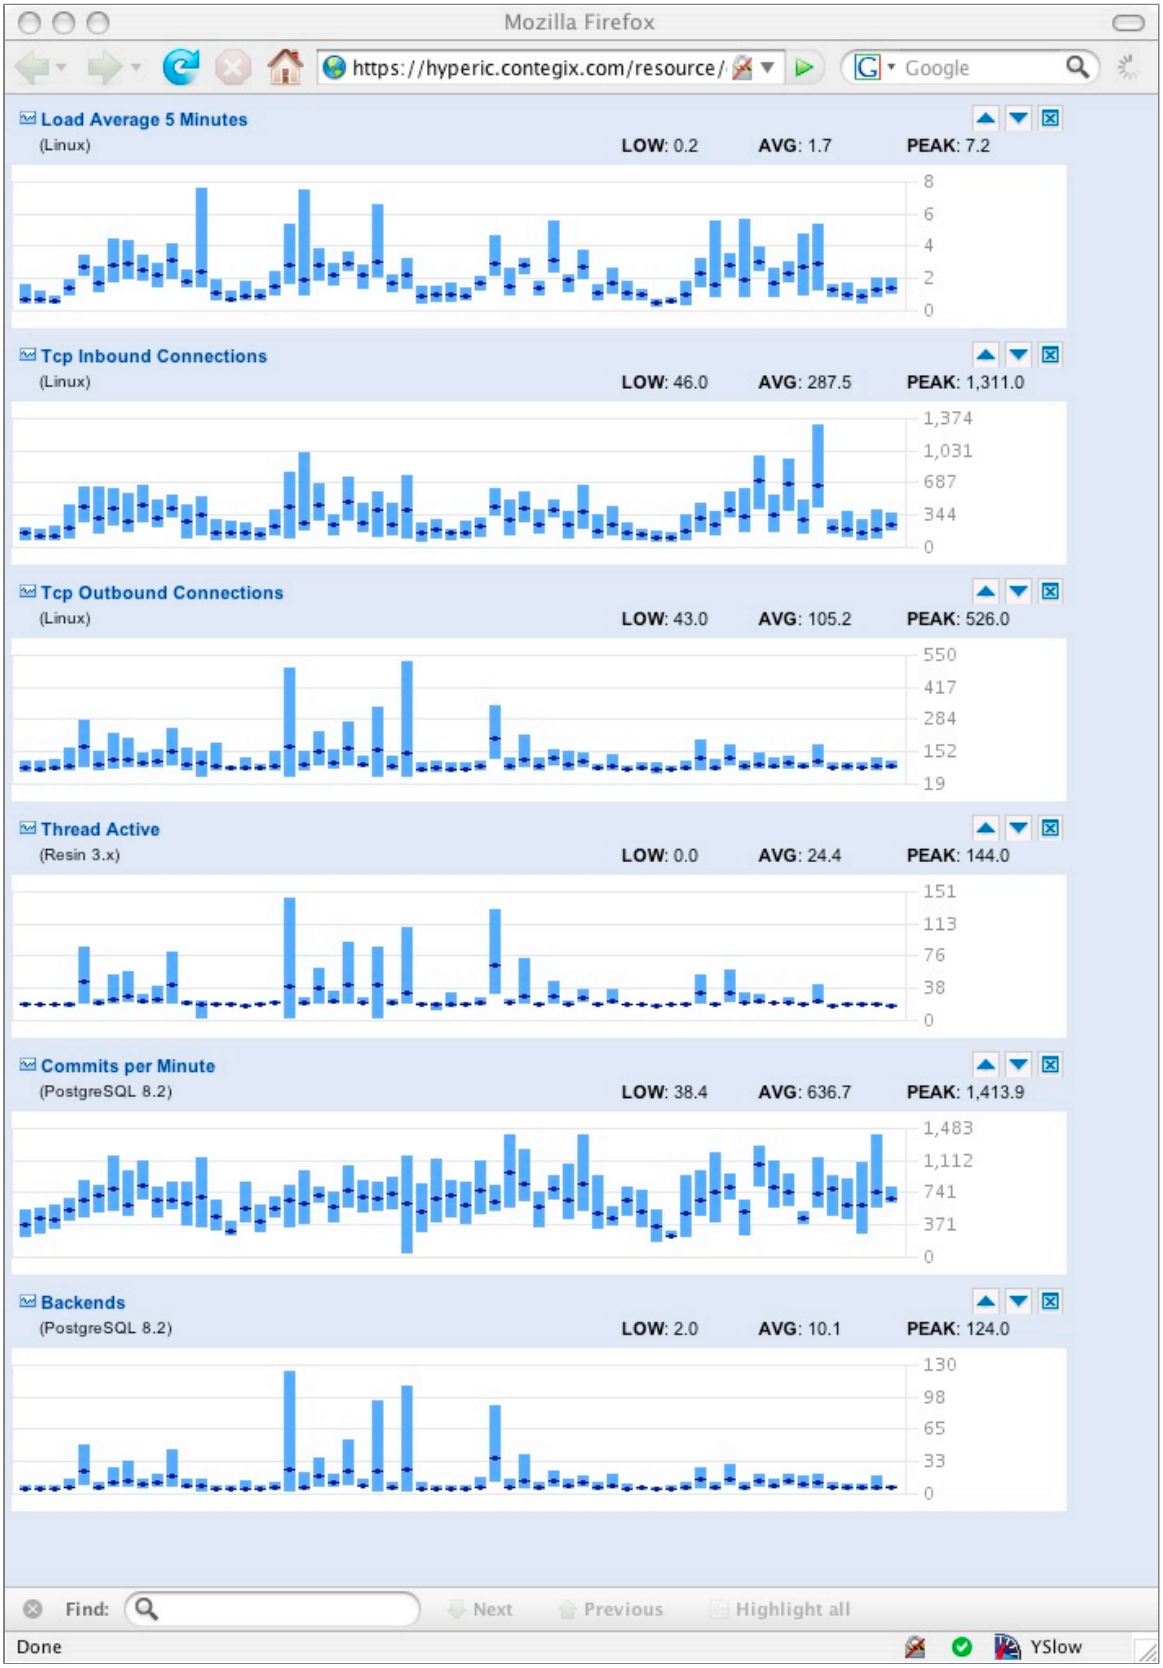
Monitoring involves a number of essential tasks, including those listed below:

- Monitoring log files.
- Checking for HTTP-availability and performance (e.g. by getting the same page every five minutes and displaying the time on a graph).
- Looking at many different parameters such as load, connections, IO, database-trends, and so on.
- Charting long-term trends.
- Keeping an access log of requests to the web server. This is vital, especially when requesting performance-related support from Atlassian.

Monitoring a web application like Confluence implies also monitoring the subsystems it uses. Many outages and downtimes are caused by broken mail servers, databases running out of space, file systems filling up and so on. It is often possible to detect these trends way before the actual web application breaks down. Keep an eye on the file system, and if you see it is getting closer to 90% utilisation, you can mend the situation without Confluence breaking down. Or even if the worst case happens (e.g. the database breaks down and Confluence is affected straight away) then having the proper monitoring for the database server makes troubleshooting a lot easier.

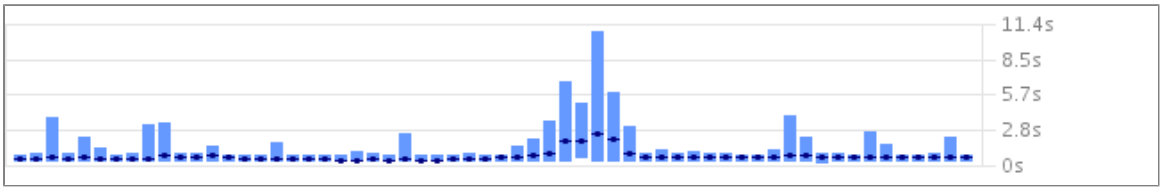
Tools for Monitoring Confluence
At Atlassian we use Hyperic. But the list of monitoring systems is long and we can't recommend a specific product over the other. If your organisation has a monitoring system already, make sure you hook up Confluence to it. If you don't have a monitoring system yet, you need to install one as soon as you feel Confluence is mission critical.

As an example of what our monitoring UI looks like, have a look at this screenshot:



The following screenshot shows one of our sensors looking at the HTTP response times of our documentation wiki over the last 8 days. You

can clearly see an incident four days ago. Having the graph (and regularly looking at it) allowed us to pinpoint the problem. We analysed the access logs and found that webpage-profiling had been enabled but not disabled again, which caused performance problems.



This page would get too long if we described all our monitoring sensors - but just to give you an impression, this is what we monitor on the JVM level alone.

JVM basics

- Current Loaded Classes
- Daemon Thread Count
- Heap Memory Committed
- Heap Memory Max
- Heap Memory Used
- Loaded Classes
- Loaded Classes per Minute
- Object Pending Finalization Count
- Peak Thread Count
- Thread Count
- Unloaded Classes
- Unloaded Classes per Minute

JVM garbage collection

- Collection Count
- Collection Count per Minute
- Collection Time
- Collection Time per Minute

JVM memory: (Metrics for Eden space, Old Gen, Survivor space, Perm Gen)

- Commited Memory
- Used Memory

We get the same level of detail for our database, for the file system, for the CPU, for the network, and so on. Not all of this is needed all the time. But if your company depends on an application, then the more information you have at your fingertips the better. Fortunately these metrics can be extracted quite easily once you have a monitoring system in place.

Adherence to Strict Upgrade Procedures

Your organisation will have its own upgrading procedure. Here are a few recommendations that you should add to your list:

- Our main recommendation: Never change more than one component at a time. Sometimes it may be tempting to upgrade the server hardware when you upgrade Confluence, but we recommend you don't do that. It makes pinpointing errors much more difficult. So, for example, don't upgrade hard disks in conjunction with a Confluence version upgrade, don't change the Confluence configuration at the same time as you upgrade your Apache software, and don't upgrade a major third-party plugin the day you move your database system to a new machine. The list is endless, these were just a few examples to get you thinking.
- After each upgrade step, run Confluence for a couple of days to check that everything is still fine.
- Keep track diligently of what you change, and when. It will be nearly impossible for us to help you if you can't tell us what exactly you changed at what time.
- Keep a copy of all log files produced during the upgrade, together with notes about what changed between successive restarts.

Always take careful note of the upgrade notes published with the Release Notes of each Confluence version, as well as the Confluence Upgrade Guide.

Example

Here you can see an extract of our change log for `http://confluence.atlassian.com` — the server that hosts this very page.

Sydney time	Server time	Event	Reason/Purpose (including JIRA issues)
	2008-03-25 22:18	Started upgrade to 2.8-m9-r3 (build #1314)	
	2008-03-25 22:25	App server brought down due to failed database upgrade	
	2008-03-26 00:51	Server brought back up after database restored from backup. Running 2.8-m9-r3.	

	2008-03-28 04:18	GC algorithm changed from concurrent to parallel collector. Max heap increased from 1.4 GB to 2.0 GB	
	2008-04-24	Hyperic agent started with connection to Resin.	
	2008-05-08 20:30 - 22:30	Manual updates to menu.css, comments.js and comments.css in webapp	Temporary fix for , which was impacting performance
	2008-05-12	Updated cache sizes for five caches, bounced server.	Cache efficiency was low on these caches.
2008-05-13 18:00-18:20	2008-05-13 03:00-03:20	Upgrade from Resin 3.0 to Tomcat 5.5	
2008-05-14 16:30-17:00		Upgrade from Confluence 2.8.1-rc2 to 2.8.1-rc3	
	2008-05-14 20:30	Install new cronjob as j2ee for automating access log analysis	

Testing of Upgrades before Production Implementation

You should test upgrades in a staging environment.

Before rolling out a new version of Confluence (or of the software or hardware that it uses, e.g. database systems, application servers, data storage), make sure that you test the upgrade with real data (e.g. a database dump) on a completely independent machine.

Here's an example of what such a test would pick up: The new release of Confluence may not be compatible with a custom third party plugin you have previously installed, thus breaking the plugin's functionality. You may not even know that anyone installed that plugin — but maybe many people are already using it. You'll want to find out about this before you actually roll out the new version of Confluence.

Here is an outline for a simple upgrade test:

1. Create a clone of your production environment, using a database dump to obtain a copy of the Confluence data. We'll call this your 'staging environment'.
2. Upgrade the staging environment to the new version of Confluence.
3. Ask a few selected users from different departments to check the pages they commonly access, but have them do it in the staging environment.

Hint: In addition to finding weirdnesses with plugins, this may also show whether training for new functionality is needed in some of the departments. The IT department staff may be able to handle the upgrade to a new version of Confluence without training, but perhaps the sales representatives who use the wiki less often will need some training.

Getting a license for your staging environment

Only a technical contact for your commercial/academic license is able to create a Developer license

Atlassian supplies 'developer' licenses which can be used by existing commercial license holders who wish to deploy non-production installations of our software to use in QA/staging environments. Developer licenses are free of charge to commercial license holders and, like our commercial offerings, they include 12 months of updates starting from the date of purchase of the commercial license.

If you hold a commercial license, you can obtain a free developer license by performing the following:

1. Log in to your Atlassian account.
2. Under the "Licenses" heading, all of your licenses will be displayed. Click the plus sign next to a license to view its details.
3. Click the **'View Developer License'** link in the bottom right corner of the license detail panel, below your commercial license key.

Load-Testing Environments

Many customers ask us,

> *So, how many users and spaces can I put into Confluence, and what is the best hardware do to so?*

The answer is, 'It depends'.

It depends a lot on your use case. Confluence is so successful because it can cover a huge range of use cases. If most of your users only access Confluence infrequently, it is no problem to have 70 000 to 100 000 users. But if each user is a power-user who uses the system the whole day, there's a substantial decrease in number Confluence can take without tuning. If your pages are short, simple, and don't contain a lot of macros, then the situation will be vastly different from a system that relies heavily on macros, background-tasks, or other features.

If your system is large (for example serving more than 10 000 users or storing more than 1000 spaces) or mission-critical (which it could be with as few as 1000 users who use it all the time) you need one or more more load-testing environments.

Even if your system is working nicely for 20 000 users right now, it might take just another 2000 users to push it over the edge.

We recommend the following basic procedure:

- Set up an environment that closely resembles your production environment.
- Gather statistics from your production system.
- Regularly apply a similar kind of load (and slightly higher) to the load-testing environment.
- Analyse how well Confluence scales for your usage patterns.

The Confluence development team has load-testing scripts available which you can use to simulate load. You can also contact Atlassian Support for more details.

Tuning

You may need to be able to tune your installation in the ways mentioned below.

Optimising your System

If you have large numbers of users, then downloading all the static content (CSS, default images, JavaScript-files) may result in a high additional load on the application server that can be offloaded to a caching web server.

Please refer to the following additional information:

- Our general Performance Tuning page.
- Information on configuring a large Confluence installation.

Limiting Third-Party Plugins

You may have to restrict the number of third-party plugins installed on your Confluence instance.

Most third-party plugins are not specifically written for high-load environments. What works fine in low-load environments could have unexpected and adverse effects when thousands of users are competing for your application server's CPU time or for database IO.

A common source of problems is access to database connections. If you have fewer users than database connections, it does not matter if an operation holds on to a database connection for two seconds while it downloads some data from the internet. With hundreds of concurrent users, this could quickly become a bottleneck.

Confluence itself is tested and optimised to handle high loads and avoids these kinds of problems. But if you install a number of plugins that have not been tested against high load, your system may become unstable.

We recommend that you load test the common use cases of each unofficial third-party plugin if your Confluence installation is mission critical. Only activate plugins that are vital to your business, and never allow experimental plugins onto your production system until they have been tested in a staging environment.

Selecting and Tuning your JVM

You should select your JVM carefully and you may need to be able to tune it.

The selection of the JVM for your large Confluence instance can have a huge impact on the performance perceived by the users. Between versions 1.4 and 6 of the Sun Java JVM there have been some impressive improvements in performance, especially under high concurrent load.

Here are some essential guidelines:

- Always run the most recent point release of your selected JVM.
- Where ever possible run the most recent major release from your selected JVM manufacturer. The Sun JVM version 6 is much faster than 1.4, especially under high loads.
- Tune your garbage collection algorithms. Experiment with different algorithms and settings to get the response times you desire in your environment. Here are some specific guidelines for Sun JVM in the Sun documentation:
    - Java 6
    - Java 5
    - Java 1.4

Customising Confluence to Optimise Performance

You may need to customise Confluence for performance reasons. Depending on your usage scenario, there may be ways to enhance Confluence performance that become necessary when you reach a certain level of usage.

Here are some things you might decide to do:

- Remove the display of the space list on the Dashboard. See Customising the Dashboard.
- Configure any search appliances or other crawlers which are configured to index the Confluence site:
    - These should be suitably rate limited.
    - Configure them to crawl only pages in the `/display/` URL path, and only current versions of pages.

Please refer to our general Performance Tuning page for more details.

RELATED TOPICS

Performance Tuning
Configuring a Large Confluence Installation
Confluence Clustering Overview
Requesting Performance Support
Confluence Administrator's Guide
Confluence Configuration Guide
Server Hardware Requirements Guide
Managing Application Server Memory Settings

Performance Tuning

- Description
- Use the latest version of your tools
- Avoid swapping due to not enough RAM
- Careful about those other systems using the same infrastructure
- Choice of Database
- Database Connection Pool
- Database in general
- Database indexes
- Database Statistics and Query Analysers
- Cache Tuning
- Antivirus Software
- Enabling HTTP Compression
- Virtual Operating Systems
- Performance Testing
- Access logs
- Built-in Profiler
- Adjust Application Server Memory Settings
- Use A Web Server
- Parallel GC
- Troubleshoot possible memory leaks
- Some 3rd-party plugins were not written to scale to large enterprises' needs
    - RELATED TOPICS



This document describes tuning your application for improved performance. It is not a guide for troubleshooting Confluence outages. Check Troubleshooting Confluence Hanging or Crashing for help if Confluence is crashing. NEW: Garbage Collector Performance Issues

Description

Like any server application, Confluence may require some tuning as it is put under heavier use. We do our best to make sure Confluence performs well under a wide variety of circumstances, but there's no single configuration that is best for everyone's environment and usage patterns.

If you are having problems with the performance of Confluence and need our help resolving them, you should read Requesting Performance Support.

Use the latest version of your tools

Use the latest versions of your application servers and Java runtime environments. Newer versions are usually better optimized for performance. As an example, our internal performance tests show a **20% speed-up** (when viewing pages under load) between Tomcat 6 on Java 6 vs Tomcat 5.5 on Java 5 **out of the box**.

Avoid swapping due to not enough RAM

Always watch the swapping activity of your server. If there is not enough RAM available, your server may start swapping out some of Confluence's heap data to your hard disk. This will slow down the JVM's garbage collection considerably and affect Confluence's performance. In clustered installations, swapping can lead to a Cluster Panic due to Performance Problems. This is because swapping causes the JVM to pause during Garbage Collection, which in turn can break the inter-node communication required to keep the clustered nodes in sync.

Careful about those other systems using the same infrastructure

It may sound tempting: Just have one powerful server hosting your database and/or application server, and run **all** your crucial programs on that server. **If** the system is set up perfectly, then you might be fine. Chances are however that you are missing something, and then one application's bug might start affecting other applications. So if Confluence is slow every day around noon, then maybe this is because

another application is using the shared database to generate complicated reports at that time? Either make sure applications can't harm each other despite sharing the same infrastructure, or get these systems untangled, for example by moving them to separate instances that can be controlled better.

Choice of Database

The **embedded database** that is provided with Confluence is meant only to be used for evaluation, not for production Confluence sites. After the evaluation finishes, you will certainly need to switch to an external relational database management system. Beyond this, we do not recommend any particular RDBMS over another. We recommend using what you are familiar with, because your ability to maintain the database will probably make far more difference to what you get out of it than the choice of database itself.

Database Connection Pool

If load on Confluence is high, you may need more simultaneous connections to the database.

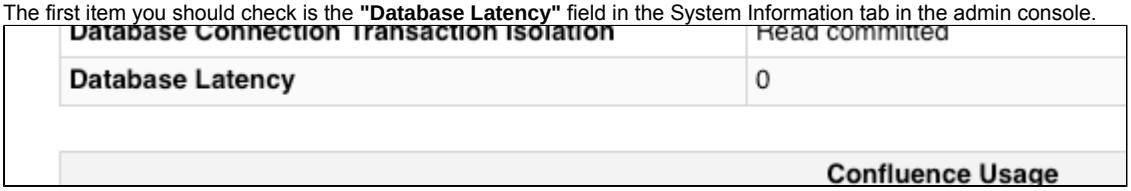
- If you are using JNDI data-sources, you will do this in your application server's configuration files.
- If you have configured Confluence to access the database directly, you will need to manually edit the hibernate.c3p0.max_size property in the confluence.cfg.xml file in your confluence.home directory. After you have changed the URL in this file, restart Confluence.

To assess whether you need to tune your database connection pool, take thread dumps during different times (including peak usage). Inspect how many threads have concurrent database connections.

Database in general

If Confluence is running slowly, one of the most likely cause is that there is some kind of bottleneck in (or around) the database.

The first item you should check is the **"Database Latency"** field in the System Information tab in the admin console.



The latency is calculated by sending a trivial request to the database, querying a a table which is known to have only one column and one row. ("select * from CLUSTERSAFETY"). Obviously this query should be blazing fast, and return within 1 or 2 milliseconds. If the value displayed is between 3 and 5 milliseconds, you might already have an issue. If the value is above 10ms, then you **definitely** need to investigate and improve something! A few milliseconds may not sound so bad, but consider that Confluence sends quite a few database queries per page request, and those queries are a lot more complex too! High latency might stem from all sorts of problems (slow network, slow database, connection-pool contention, etc), so it's up to you to investigate. Don't stop improving until latency is below 2ms on average.

Obviously, latency is just the very first thing to look at. You may get zero latency and still have massive database problems, e.g. if your tables are poorly indexed. **So don't let a low latency fool you either.**

Database indexes

Especially if you have more than a few thousand active users, and all most obvious measures have been tried out but the database still seems to be under high load, you should consider enaging a database administrator (DBA) to tune the database specifically to the demands that your particular Confluence installation is placing on it. If you do not have a full-time DBA and can't even get one for temporary consulting, you may want to consult the database indexing advice that we have been gathering from customer reports and our own experience running and developing Confluence. The instructions on that page are for Oracle, but most of the indexes can be applied to (and will help with) any database.

(These database indexes are now created automatically when Confluence is installed, but existing installations upgrading to a more recent version may still need to add them manually)

Database Statistics and Query Analysers

Modern databases have query optimisers based on collecting statistics on the current data. Using the SQL EXPLAIN statement will provide you information on how well the query optimiser is performing. If the cost estimate is wildly inaccurate then you will need to run statistics collection on the database. The exact command will depend on your database and version. In most cases you can run statistics collection while Confluence is running, but due to the increased load on the database it's best to do this after normal hours or on a week-end.

Cache Tuning

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). Future versions of Confluence will allow you to tune the size of this cache from within the web application. Vote for tuning the cache from the UI and getting cache recommendations to encourage Atlassian to build this feature into a Confluence release.

Please have a look at our documentation on Cache Performance Tuning for information on how to tune Confluence caches.

Antivirus Software

Antivirus software greatly decreases the performance of Confluence. Antivirus software that intercepts access to the hard disk is particularly detrimental, and may even cause errors with Confluence. You should configure your antivirus software to ignore the Confluence home directory, its index directory and any database-related directories.

Enabling HTTP Compression

If bandwidth is responsible for bottlenecking in your Confluence installation, you should consider enabling HTTP compression. This may also be useful when running an external facing instance to reduce your bandwidth costs.

 Take note of the known issues with HTTP compression in versions of Confluence prior to 2.8, which may result in high memory consumption.

Virtual Operating Systems

Virtual Environments such as VMWare can cause Confluence CPU to spike. Run Confluence on a native OS. Refer to the list of supported operating systems for Confluence in the Supported Platforms topic.

Performance Testing

You should try out all configuration changes on a demo system. Ideally, you should run and customize loadtests that simulate user behaviour. Learn about how to test performance issues using the Performance Testing Scripts.

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.

Adjust Application Server Memory Settings

See Managing Application Server Memory Settings.

Use A Web Server

For high-load environments, performance can be improved by using a web server such as Apache in front of the application server. There is a configuration guide to Running Confluence behind Apache.

When configuring your new web server, make sure you configure sufficient threads/processes to handle the load. This applies to both the web server and the application server connector, which are typically configured separately. If possible, you should enable connection pooling in your web server connections to the application server.

Parallel GC

If you have multiple CPU's on your server, you can add -XX:+UseParallelOldGC to your JAVA_OPTS options. This will allow garbage collection of the Tenured Space to happen in parallel with the application and can boost performance and can reduce slow performance spikes. For more information, please refer to our detailed page on Garbage Collector Performance Issues, and Sun's summary of collectors.

Troubleshoot 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.

Some 3rd-party plugins were not written to scale to large enterprises' needs

Confluence has been optimized to work under high load and with many pages. Some 3rd party plugins however have been written with small size companies in mind, and can't cope with large numbers of concurrent users, or large numbers of pages and permissions, or large numbers of spaces. It is impossible to tell which ones will fail under which conditions, but it will always help to turn off 3rd-party plugins that are not strictly mission-critical while investigating performance issues.

RELATED TOPICS

Garbage Collector Performance Issues
Cache Performance Tuning
Cache Performance Tuning for Specific Problems
Performance Testing Scripts
Working with Confluence Logs
Operating Large or Mission-Critical Confluence Installations
Confluence Clustering Overview

Requesting Performance Support
Confluence Administrator's Guide
Confluence Configuration Guide

Cache Performance Tuning

Confluence performance can be significantly affected by the performance of its caches. It is essential for the administrator of a large production installation of Confluence to tune the caches to suit its environment. There are several configurable parameters for each of the cache regions, most notably cache size, cache expiry delay and eviction policy. In the majority of the cases, cache size is the parameter you would want to change. Fortunately, from Confluence 3.0, it is very easy to adjust cache sizes through the Administration Console. However, if you need to modify parameters other than a cache size, you would need to modify the relevant configuration files manually.



If you only need to modify Confluence's maximum cache sizes, you can do this through the Cache Statistics feature of the Administration Console.

The cache performance information for your Confluence installation is available under **Administration > Cache Statistics**. More information about the numbers displayed here is available on Cache Statistics.

On this page:

- Cache tuning example
- Finding the configuration file
- Cache Key Mappings
- Standard Editions of Confluence
  - Understanding the Ehcache Configuration File
  - Converting your Coherence configuration to Ehcache
- Clustered Editions of Confluence
  - Understanding the Coherence configuration file
  - Defining Caching Scheme Mappings in Coherence Cache config file
- Important Caches
- Cache Tuning Follow-Up

Cache tuning example

As an example of how to tune Confluence's caches, let's have a look at the following table:

Caches	% Used	% Effectiveness	Objects/Size	Hit/Miss/Expiry
Attachments	87%	29%	874/1000	78226/189715/187530
Content Attachments	29%	9%	292/1000	4289/41012/20569
Content Bodies	98%	81%	987/1000	28717/6671/5522
Content Label Mappings	29%	20%	294/1000	4693/18185/9150
Database Queries	96%	54%	968/1000	105949/86889/83334
Object Properties	27%	18%	279/1000	5746/25386/8102
Page Comments	26%	11%	261/1000	2304/17178/8606
Users	98%	5%	982/1000	6561/115330/114279

The caches above are of size 1000 (meaning that it can contain up to 1000 objects), which is the default size for caches in the default cache scheme. Refer to Confluence Cache Schemes for more explanation.

You can tell when a cache size needs to be increased because the cache has both:

- a high usage percentage (above 75%)
- a low effectiveness percentage.

Check the 'effectiveness' versus the 'percent used'. A cache with a low percent used need not have its size lowered; it does not use more memory until the cache is filled.

Based on this, the sizes of the "Attachments", "Database Queries", and "Users" caches should be increased to improve their effectiveness.

As the stored information gets older or unused it will expire and be eliminated from the cache. Cache expiry may be based on time or on frequency of use.



There is not much that you can do with a cache that has both a low percentage of usage and effectiveness. Over time, as the cache is populated with more objects and repeat requests for them are made, the cache's effectiveness will increase.

Finding the configuration file

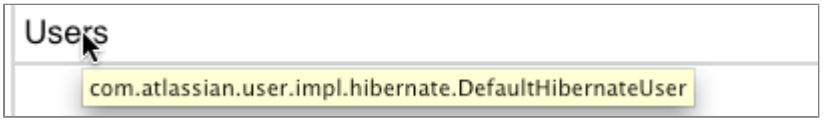
The caches are configured in `ehcache.xml` (for standard editions) or `confluence-coherence-cache-config-clustered.xml` (for clustered editions) which is stored in `<confluence-home>/config/`.

Oracle Coherence Licensing Change:

- Due to a license agreement change between Atlassian and Oracle over the Coherence technology, from September 2009, Confluence will be made available in two editions:
  - **Standard Edition** — Confluence with Ehcache's caching technology (available to customers with non-clustered Confluence licenses).
    If you are currently running a clustered installation of Confluence, please do not upgrade it with a standard edition of Confluence.
  - **Clustered Edition** — Confluence with Oracle's Coherence clustering and distributed caching technology (available to customers with Confluence clustered licenses only).
- For more information about these changes, please refer to the Coherence License Changes document.
- If you have a Confluence clustered license, are running a clustered installation of Confluence and wish to upgrade to Confluence version 2.6 or later after late September 2009, please ensure that you download only a clustered edition of Confluence and please refer to the Confluence 3.0.1 Upgrade Notes for additional upgrade information.

Cache Key Mappings

The cache configuration file configures caches by their keys. When you move your mouse over the the cache names displayed on the cache statistics page, a tooltip will indicate the actual cache key for that cache name.



Using our example from the table above, if we were to modify parameters for the `Users` cache we would need to change the cache with the key `com.atlassian.user.impl.hibernate.DefaultHibernateUser`. Do not get confused with `Users (External Mappings)` and `Users (External Groups)` which are in themselves, two separate caches. "Users" is the friendly name for `com.atlassian.user.impl.hibernate.DefaultHibernateUser`.

Standard Editions of Confluence

In standard editions of Confluence, the caching layer is Ehcache.

Understanding the Ehcache Configuration File

For more information about the Ehcache configuration file and a full reference on Ehcache configuration, please refer to the Ehcache configuration documentation.

Converting your Coherence configuration to Ehcache

This section only applies to customers who:

- Have an installation of Confluence that was downloaded before the 4th of September 2009.
- Intend to (or have already) upgraded to Confluence 3.0.1 or later (or to Confluence versions 2.6.3, 2.7.4, 2.8.3, 2.9.3 and 2.10.4).
- Will use a non-clustered Confluence license for the Confluence upgrade.
- Have implemented customisations to their Confluence installation's cache configuration file (`confluence-coherence-cache-config.xml`).

To maintain your existing cache configuration file settings, you will need to transfer any cache customisations you have implemented in the Coherence cache configuration file (`confluence-coherence-cache-config.xml`) to the relevant entries in the Ehcache cache configuration file (`ehcache.xml`).

Each cache has a `cache-mapping` element in the Coherence file (of which there is an equivalent `cache` element in the `ehcache.xml` file). Unfortunately, copying across your customisations is not quite a straightforward process because the Coherence file defines several 'caching schemes' to store the actual cache values, which in turn are referenced by the `cache-mapping` elements. In contrast, the `ehcache.xml` file does not support caching schemes and a cache's values are expressed explicitly in separate parameters of a `cache` element.

To convert your Coherence cache configuration file customisations across to the equivalent Ehcache file:

1. Open both the `confluence-coherence-cache-config.xml` and `ehcache.xml` files in a text editor. These files are located in the `<confluence-home>/config` directory.

If you implemented your customisations in a version of Confluence prior to 3.0, you will most likely find the `confluence-coherence-cache-config.xml` file in the `<confluence-install>/confluence/WEB-INF/classes` directory.

2. In the customised `confluence-coherence-cache-config.xml` file:
    a. Identify the caching schemes that were customised in this file and make a note of the values of all its child elements.
    Typically, each caching scheme is located inside a `local-scheme` element and all of these are enclosed within the `cache-schemes` element, which appears towards the end of this file.
    b. Note each customised caching scheme by the content of its `scheme-name` element.
    c. For each `cache-mapping` element (which typically appears towards the top of this file), identify if it has a `scheme-name` element whose content matches one noted in the previous step and if so, make a note of its associated `cache-name` element.
3. In the `ehcache.xml` file:
    a. Identify each `cache` element whose 'name' parameter matches the `cache-name` elements noted in step '2c'.
    b. Using the mappings table below, apply the values noted in step '2a' to the appropriate parameters of the `cache` elements identified in the previous step ('3a').

Mappings table showing how elements of the Coherence cache configuration file map to parameters of the equivalent Ehcache file.

Coherence Element	Ehcache Attribute
`high-units`	`maxElementsInMemory`
`expiry-delay` > 0s	`timeToIdleSeconds` - Use this attribute for expiry delays greater than 0s along with the `eternal` attribute set to 'false'
`expiry-delay` = 0s	`eternal` - For expiry delays of 0s, set this attribute to 'true'.

Clustered Editions of Confluence

Understanding the Coherence configuration file

The Coherence configuration file is a mapping of *cache keys* to *cache schemes*. Each cache scheme controls the expiry, eviction policy and size of the caches linked to it. A cache scheme can extend another scheme.

For a full reference, see the Oracle's Coherence cache configuration documentation.

Defining Caching Scheme Mappings in Coherence Cache config file

If a cache key does not have an explicit definition in the caching scheme mappings (defined in `confluence-coherence-cache-config.xml`) then it will use the "default" `cache-mapping`.

In our example, `com.atlassian.user.impl.hibernate.DefaultHibernateUser` is not explicitly defined in the caching scheme mappings. Hence to increase the expiry-delay to 2 hours, we will need to define the mapping ourselves and add the following within the `<caching-scheme-mapping>...</caching-scheme-mapping>` tags:

```
1.<cache-mapping>
2.<cache-name>com.atlassian.user.impl.hibernate.DefaultHibernateUser</cache-name>
3.<scheme-name>cache:com.atlassian.user.impl.hibernate.DefaultHibernateUser</scheme-name>
4.</cache-mapping>
```

Then we will need to define a cache schema with name `cache:com.atlassian.user.impl.hibernate.DefaultHibernateUser` within `<caching-schemes>...</caching-schemes>` tags.

```
1.<local-scheme>
2.<scheme-name>cache:com.atlassian.user.impl.hibernate.DefaultHibernateUser</scheme-name>
3.<scheme-ref>default</scheme-ref>
4.<high-units>10000</high-units>
5.<expiry-delay>7200</expiry-delay>
6.</local-scheme>
```

It's possible to define a local-scheme mapping for a cache key without defining certain parameters (e.g. `<high-units>`). In such a cases, their parameters will be inherited from `scheme-ref` scheme, which is the `default` scheme in our case.

Important Caches



The following suggestions are general guidelines. In cases of large databases, 20-30% of the size of the table may be unnecessarily large. Check the effectiveness and Percent Used categories in the cache for more specific assessments.

- **`com.atlassian.confluence.core.ContentEntityObject`** (known as **`Content Objects`** cache)
  should be set to at least 20-30% of the number of content entity objects (pages, comments, emails, news items) in your system. To find the number of content entity objects, use the query `select count(*) from CONTENT where prevver is null`.
- **`com.atlassian.confluence.core.ContentEntityObject.bodyContents`** (known as **`Content Body Mappings`** cache)
  should be set to at least 20% of the number of content entity objects (pages, comments, emails, news items) in your system. To find the number of content entity objects, use the query `select count(*) from CONTENT where prevver is null`.
- **`com.atlassian.confluence.security.PermissionCheckDispatcher.isPermitted()`** (known as **`User Authorized URLs`** cache)
  should be set to at least the number of concurrent users you expect to access Confluence at the same time
- **`com.atlassian.user.impl.hibernate.DefaultHibernateUser`** (known as **`Users`** cache)
  should be set to the number of users you have: `select count (*) from users`. Note that by default, this will also control the LDAP user's cache, including expiration.
- **`com.atlassian.confluence.security.SpacePermission`** (known as **`Permissions`** cache)
  should be set to the number of space permissions in your deployment (a good rule of thumb is 20 times the number of spaces). You can find the number of space permissions using the query `select count(*) from SPACEPERMISSIONS`.

Cache Tuning Follow-Up

After you have made changes to your cache config, doing a follow up on the changes in the next week or after the expected performance spike would be important.

Make sure that you take a screenshot of the cache statistics before and after the change. Then compare them with the cache statistics in the later period where performance improvement is expected.



You can monitor what's in the cache by using a JSP included in the Confluence distribution. Browse to <base-URL>/admin/cachecontents.jsp to monitor the cache contents.

RELATED TOPICS

Cache Performance Tuning for Specific Problems
Confluence Cache Schemes
Performance Testing Scripts
Working with Confluence Logs
Operating Large or Mission-Critical Confluence Installations
Confluence Clustering Overview
Requesting Performance Support
Confluence Administrator's Guide
Confluence Configuration Guide

Cache Performance Tuning for Specific Problems

The following are more specific performance problems that can be resolved from tuning the cache.

LDAP cache sizes and expiry does not appear to be picked up.

This is a known problem, please refer to CONF-11858 for the solution.

"Edit Page" screen takes a long time to load

If your installation of Confluence is suffering from this problem, it may be due to a insufficient SpacePermissions cache size. To address this problem, first determine the number of space permission objects in your Confluence instance. You can do this by running this query against your database:

```
> select count(*) from SPACEPERMISSIONS
```

Now locate the cache entry for SpacePermissions in your `confluence-coherence-cache-config.xml`:

```
1.<local-scheme>
2.
<scheme-name>cache:com.atlassian.confluence.security.CachingSpacePermissionManager.permissions</schem
3.                    <scheme-ref>default</scheme-ref>
4.                    <high-units>10000</high-units>
5.                    <expiry-delay>0s</expiry-delay>
6.</local-scheme>
```

Adjust the **maxElementsInMemory** or **high-units** property to the number of space permissions you have (in the example above, I've used 10000). Also, just as important, you need to adjust the **timeToLiveSeconds** or **expiry-delay** property to `0`.

Note: 10K of space permissions consumes approximately 8MB of memory. Please ensure there is enough memory allocated to your instance to cater for this.

How to set specific cache settings

1. Find the cache name from the cache name mappings:
   - For **Confluence 2.5.x and earlier**, the cache name mappings are in file `confluence/WEB-INF/classes/com/atlassian/confluence/admin/actions/cache-name-mappings.proper` .
   - For **Confluence 2.6.0 and later**, you will find the cache name mappings in the file `com/atlassian/confluence/core/ConfluenceActionSupport.properties` which is packed into the `confluence-2.x.*.jar file`.
2. Find the appropriate `<cache-mapping>` tag in `confluence-coherence-cache-config.xml` or `confluence-coherence-cache-config-clustered.xml`. If the tag doesn't exist, you can create it within the `<caching-scheme-mapping>` tag.

   Attached to this page are corrected copies of confluence-coherence-cache-config.xml and confluence-coherence-cache-config-clustered.xml. These are updated from a bug CONF-11857.

3. The `<scheme-name>` will correspond to a `<local-scheme>` tag below. It refers to a scheme reference. Either change the high-units tag in the scheme reference, or add a high-units tag to override the scheme reference. For example, the following tag would change the Content Bodies cache from the default 1000 units to 2000 units:

   ```
   <local-scheme>
   <scheme-name>cache:com.atlassian.confluence.core.ContentEntityObject.bodyContents</scheme-name
   ```

   Another popular cache to change is the LDAP related User cache:

   ```
   <local-scheme>
   <scheme-name>user</scheme-name>
   <scheme-ref>default</scheme-ref>
   <high-units>5000</high-units>
   <expiry-delay>300s</expiry-delay>
   </local-scheme>
   ```

4. After updating the appropriate file, you do not need to repack it into the jar to use it. You can simply place the file in your `confluence/WEB-INF/classes/` directory. The file in this directory will override the settings in your jar file. If you want to back out the changes, you only need to remove the file from your `confluence/WEB-INF/classes/` directory — then the default values in the `confluence-coherence-cache-config.xml` located in your jar file will apply.

You can find more information about configuring the Coherence cache in the Coherence cache documentation.

RELATED TOPICS

Cache Performance Tuning
Performance Testing Scripts
Confluence Cache Schemes
Working with Confluence Logs
Operating Large or Mission-Critical Confluence Installations
Confluence Clustering Overview
Requesting Performance Support
Confluence Administrator's Guide
Confluence Configuration Guide

Confluence Cache Schemes

Default Scheme

If a cache has not been defined, then it will use the default cache size and expiry. As the start of your `confluence/WEB-INF/classes/confluence-coherence-cache-config.xml` file you will notice the following:

```
1.<cache-mapping>
2.      <cache-name>*</cache-name>
3.      <scheme-name>default</scheme-name>
4.</cache-mapping>
```

So basically all caches will default to using the default scheme, which is defined as below:

```
<!-- Default scheme -->
<local-scheme>
    <scheme-name>default</scheme-name>
    <class-name>com.atlassian.confluence.cache.tangosol.ExpiryCountingLocalCache</class-name>
    <high-units>1000</high-units>
    <expiry-delay>3600</expiry-delay>
</local-scheme>
```

I.e. with a size of 1000 Objects and an expiry of 3600 seconds. Other schemes use the above as their default and either override the size of the cache, or the length of the expiry.

Common Schemes

In addition to the default scheme, there are also common schemes used in Confluence caches:

```
<!-- Common schemes -->
<local-scheme>
    <scheme-name>large</scheme-name>
    <scheme-ref>default</scheme-ref>
    <high-units>10000</high-units>
</local-scheme>
<local-scheme>
    <scheme-name>medium</scheme-name>
    <scheme-ref>default</scheme-ref>
    <high-units>5000</high-units>
</local-scheme>
<local-scheme>
    <scheme-name>small</scheme-name>
    <scheme-ref>default</scheme-ref>
    <high-units>100</high-units>
</local-scheme>
<local-scheme>
    <scheme-name>large-transient</scheme-name>
    <scheme-ref>default</scheme-ref>
    <high-units>10000</high-units>
    <expiry-delay>300s</expiry-delay>
</local-scheme>
<local-scheme>
    <scheme-name>user</scheme-name>
    <scheme-ref>default</scheme-ref>
    <high-units>5000</high-units>
    <expiry-delay>300s</expiry-delay>
</local-scheme>
```

RELATED TOPICS

Cache Performance Tuning
Confluence Cache Schemes
Cache Performance Tuning for Specific Problems
Requesting Performance Support
Confluence Administrator's Guide
Confluence Configuration Guide

Configuring a Large Confluence Installation

Deploying *any* application to several thousand users requires care and planning, expecially if those users are going to be relying on the application to get their work done.

General Advice

Staged Rollout

Do not try to deploy Confluence immediately to your whole organisation. Instead, roll it out department by department, or project by project.

How Confluence will scale given a particular software and hardware configuration depends very much on how Confluence is likely to be used in your organisation. Launching Confluence to everybody at once may seem like a neat idea, but it also means that any problems you might experience scaling the system up to your entire organisation will hit you *all at once*, annoy everyone and possibly hurt adoption.

Rolling Confluence out gradually will give you the chance to tune it as you go, resulting in a much more painless experience. There will also be organisational advantages: you can identify those teams or projects who are most likely to be successful 'early adopters', and those teams can experiment with how best a wiki might suit your organisation, and pass on their 'best wiki practices' as usage of Confluence expands.

Plugin Governance

Confluence plugins can add tremendous value. Before adding one, visit the plugin's page and explore its issues (available from the issue management link). Try the plugin in a test environment, and make sure to note any adverse effects after adding it to a production environment. Test plugins independently when upgrading.

Backup strategy

Disable the XML backup and use the Production Backup Strategy.

New Spaces Governance

For both performance and good practice, put some modest governance in place around the creation of new spaces, such as a simple request that includes a check for duplicates and some strategy around how to best use a space. Duplicates and unused spaces should be purged by a wiki gardener. Try to keep it to one space per group.

Choose User Management and Single Signon

It is possible to integrate with an LDAP repository or add a Single Signon solution later (especially with the addition of Crowd), but if possible it's best to configure this up front. You can configure access for only a specific group or set of groups, thereby keeping the gradual rollout.

Configuring your Application Server, Web Server and Database

Because Confluence can be deployed in so many server combinations, we do not currently have guides on the best tuning parameters for each individual server. We will be happy to provide support, however. If you have any tuning parameters that you find particularly useful for Confluence instances, feel free to share them with other Confluence users in the Confluence Community space.

Best Practices

Troubleshoot possible memory leaks

The Troubleshooting Confluence Hanging or Crashing guide is a good place to start. Some of the known causes listed there could result in performance issues short of a crash or hang. Many of the issues reported there are exacerbated with a large installation.

Memory Usage

The Java virtual machine is configured with a "maximum heap size" that limits the amount of memory it will consume. If Confluence fills up this maximum heap size it will run out of memory, and start behaving unpredictably. You can keep track of Confluence's memory usage from the System Information screen of the administration console:

Java VM Memory Statistics	
Total Memory	313 MB
Free Memory	140 MB
Used Memory	173 MB
Memory Graph	 **45 % Free**

This example shows that, at the time of writing, confluence.atlassian.com is using 173MB of an allocated 313MB of heap. (The JVM was configured with a maximum heap size of 450MB, but this information is not available in the graph. The 313MB figure shows that the full 450MB of heap has not yet been needed)

Database Connection Pool

Confluence will need a database connection for each simultaneous user connection to the server. It is also a good idea to have 5-10 connections spare for Confluence internal processes such as backups, re-indexing or daily notification jobs.

Running out of pooled connections will cause the server to slow down as more users are waiting for a connection to be freed before starting their own request, and will eventually cause visible system errors as Confluence times out waiting for a database connection.

If you are using Confluence's internal connection pool, you can increase the number of available connections by modifying the `hibernate.c3p0.max_size` property in {`confluence_home`}`/confluence-cfg.xml`, and restarting Confluence. **Make sure** you have also configured your database to be able to support that many simultaneous connections.

Cache Sizes

The Performance Tuning page includes some useful rules of thumb for configuring the sizes of Confluence's internal caches.

RELATED TOPICS

Operating Large or Mission-Critical Confluence Installations
Performance Tuning
Confluence Clustering Overview
Requesting Performance Support
Confluence Administrator's Guide
Confluence Configuration Guide

How Adaptavist Runs Confluence

Key information about performance tuning and how Adapatavist runs many instances of Confluence. This was posted to the mailing list but not captured in the forums so I wanted to copy it here:

Confluence Performance Recommendations from Adaptavist

Hi,
> 1) Opterons – for better or worse, we're solely an Intel shop at this point (i.e. I'll get some resistance going down the AMD path). It seems like the latest Woodcrest Xeons have caught up with the Opterons performance-wise (although AMD was smacking around Intel for a while before that). Have you the Woodcrest Xeons at all? (I'm guessing at this point you're probably just interested in pursuing an Opteron path but figured I'd ask.)
We found the Opterons are perfect for running multithreaded apps - especially when running lots of Confluence / JIRA installs on a server in their own Resin containers. In addition, the Opterons use a less electrical power than their Intel counterparts which reduces our energy bills - for that reason we've not tried the Woodcrest chip yet. The chassis used for Opteron chips is also more established than the Woodcrest counterpart and is thus cheaper at the moment and has more options. Any dual-core will likely give better results, but you'll obviously need your OS and other software on the server to be set-up to take advantage of them.
> 2) Resin vs. Tomcat – would you mind guessing at performance numbers from what you've seen? In a ton of Googling, it seems like people are saying that Tomcat was slow back with 4.x but got much faster with 5.x and even more with 5.5.x. See the comments in the first link.

Resin, properly configured, is still faster  In addition, it uses a lot less RAM (important when you have 50 web apps in their own containers). We spent a lot of time super-fine-tuning Resin (about 4 months if memory serves - huge thanks to Caucho and the chaps at BeJUG) to run Confluence very nicely indeed. I can't vouch for it's speed with regards to other apps. We've found Resin to be highly stable (when correctly configured) and it deals very well with that elusive Confluence memory leak (something you notice on a site like JavaPolis[1] with over 17,600 registered users). It's garbage collection, again when properly tuned, was better than Tomcat and we found many tasks easier to automate with Resin as compared to Tomcat.

Admittedly, a lot of the reasons that we chose Resin for are geared to an environment where we're running up to 50 Confluences on a single server, each in their own web app. That's quite a different scenario to what you are doing where you maybe have one or two Confluences on a server. Although, having said that, we use the same set-up for our dedicated hosting (we're hosting some real BIG Confluence installs as you'll know if you followed the discussion about the import routine we've been working on) and it works great in that environment too.

We use the commercial version of Resin[2] - it's much better than the OS version as it has fewer bugs, runs more smoothly and has some real nice features (read: absolutely critically essential for the sanity of our staff thus reducing our monthly bills for padded cells and therapy) for the type of environment we use it in. We also really liked Caucho's licensing of resin[3]: $500 per physical server with 2 cores (additional cores @ $500/core which is very reasonable) regardless of the number of Resin containers on that server.

It should be noted that some of the stats you provided links to were done on Windows running Cygwin - hardly an ideal server environment  The second link (with all the graphs that people like me understand) was far more representative. We run on SuSE Linux 10 EMT64 (or something like that - whatever the latest version of their 64-bit OS is) so there's no Windows bloat getting in the way of the web apps, etc.
> 3) Memory. I think I'll go for DDR667 and see if I can bump Confluence up to 2 GB. Is there ever a point where you can allocate too much RAM? (i.e. java and/or Confluence just don't handle tracking that many cached objects well)
I think we're up to 16GB in most of our servers now. Confluence does enjoy more RAM (although containers such as Resin bring the overall RAM consumption down a fair bit - very noticeable on servers with 50 containers/apps). More RAM means more space to cache and longer gaps between forced GC. RAM allocation is also vital when it comes to the nightly backup (or "the dreadded backup hour" as we refer to it) - you can imagine the CPU and RAM spikes caused by 50 large Confluence installations all deciding to backup at the same time (roll on Confluence 2.3!)...

Should you have too much memory, you can always run a Quake server on there :o)

FYI: We also separate our database out on to a separate server.

Best Regards,

Guy

[1] http://www.javapolis.com - at last year's conference the Belgian's were somewhat annoyed at the term "SOA" which is an obscenity over there. They were also less than happy about the spoons in sexual positions plastered all over Antwerp (and several thousand Javapolians wearing the conference t-shirts). So this year Stephan and the crew have decided to push the boundaries to hitherto unimaginable levels - anyone who's seen the promo video will know exactly what I mean (and no, not the white painted bloke next to the urinal - the video goes waaaay beyond that - how they got James Gosling to... well, you'll have to wait and see)
[2] http://www.caucho.com/resin-3.0/features/overview.xtp
[3] http://www.caucho.com/sales/sales.xtp

Dan will probably be along in the morning to correct any mistakes I've made 
-

Memory usage and requirements

Managing Confluence's performance and memory usage really depends on what resources are available - Confluence will run faster if you give it lots of memory for its caches, but it should still be able to run quite well in low-memory environments, with the right tuning. Below are some tips on getting the most out of your Confluence site:

Increasing the amount of memory available to Confluence

See Increasing JIRA Memory for details on how to increase the memory available to web application servers typically used to run

Confluence.

Embedded Database

The embedded HSQL database that comes with Confluence essentially holds all your data in memory while the Confluence server is running. If you are running out of memory, you should consider migrating Confluence to some external RDBMS.

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. The default size of these caches will be reduced significantly in Confluence 1.5 (although this means that administrators of larger Confluence sites may need to configure the size of their caches to improve performance).

To customise confluences cache to meet your needs, see cache tuning.
To increase the amount of memory available to confluence, see 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. Since Confluence 1.4.1 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 / 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. This issue was resolved in Confluence post 1.4.x. To increase the amount of memory available to Confluence, see 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, file 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 logging, start Confluence with the option `-XX:+PrintGCDetails -XX:+PrintGCTimeStamps -verbose:gc -Xloggc:gc.log`. Replace `gc.log` with an absolute path to a `gc.log` file.

For example, with a Windows service, run:

```
tomcat5 //US//Confluence ++JvmOptions="-XX:+PrintGCDetails -XX:+PrintGCTimeStamps -verbose:gc
-Xloggc:c:\confluence\logs\gc.log"
```

or in `bin/setenv.sh`, set:

```
export CATALINA_OPTS="$CATALINA_OPTS -XX:+PrintGCDetails -XX:+PrintGCTimeStamps -verbose:gc
-Xloggc:${CATALINA_BASE}/logs/gc.log"
```

If you modify `bin/setenv.sh`, you will need to restart Confluence for the changes to take effect.

What can you do to minimise the time taken to handle the garbage collection? See http://java.sun.com/docs/hotspot/gc1.4.2/ for details on tuning the JVM to minimise the impact that garbage collection has on the running application.

Requesting Performance Support

Basic Performance Troubleshooting Steps

Begin with the following procedures:

1. Go through the Troubleshooting Confluence Hanging or Crashing page to identify the major known performance problems
2. Proceed with the Performance Tuning tips to help optimize performance

Requesting Basic Performance Support

If those tips don't help or you're not sure where to start, open a support ticket starting with at least the basic information:

1. The atlassian-confluence.log
2. The catalina.out log (or your application server log), with a series of three thread dumps separated by 10 seconds
3. A description with as much detail as possible regarding:
    a. What changes have been made to the system?
    b. When did performance problems begin?
    c. When in the day do performance issues occur?
    d. What pages or operations experience performance issues?
    e. Is there a pattern?

Continue with as much of the Advanced Performance Troubleshooting information as you can.

Advanced Performance Troubleshooting

Please gather **all** of the information listed below and include it in your support request, even if you think you have a good idea what's causing the problem. That way we don't have to ask for it later.

System Information

Confluence Server

- Take a screenshot of Confluence's `Administration  System Information` (or save the page as HTML)
- Take a screenshot of Confluence's `Administration  Cache Statistics` (or save the page as HTML)
- Find out the exact hardware Confluence is running on
    - How many CPUs? What make and model? What MHz?
    - How much memory is installed on the machine?
    - How much memory is assigned to Confluence's JVM? (i.e. what are the -Xmx and -Xms settings for the JVM?)
    - What other applications are being hosted on the same box?

Confluence Content

- How many *users* are registered in Confluence?
- On average, to how many *groups* does each user belong?
- How many *spaces* (global and personal) are there in your Confluence server?
- How many of those spaces would be viewable by the average user?
- Approximately how many *pages*? (Connect to your database and perform '`select count(*) from content where prevver is null and contenttype = 'PAGE'`')
- How much data is being stored in Bandana (where plugins usually store data)? (Connect to your database and perform '`select count(*), sum(length(bandanavalue)) from bandana`')

The Database

- What is the exact version number of Confluence's database server?
- What is the exact version number of the JDBC drivers being used to access it? (For some databases, the full filename of the driver JAR file will suffice)
- Is the database being hosted on the same server as Confluence?
- If it is on a different server, what is the network latency between Confluence and the database?
- What are the database connection details? How big is the connection pool? If you are using the standard configuration this information will be in your confluence_cfg.xml file. Collect this file. If you are using a Data source this information will be stored in your application server's configuration file, collect this data.

User Management

- Are you using external user management or authentication? (i.e. JIRA or LDAP user delegation, or single sign-on)
- If you are using external JIRA user management, what is the latency between Confluence and JIRA's database server?
- If you are using LDAP user management:
    - What version of which LDAP server are you using?

- What is the latency between Confluence and the LDAP server?

Diagnostics

Observed Problems

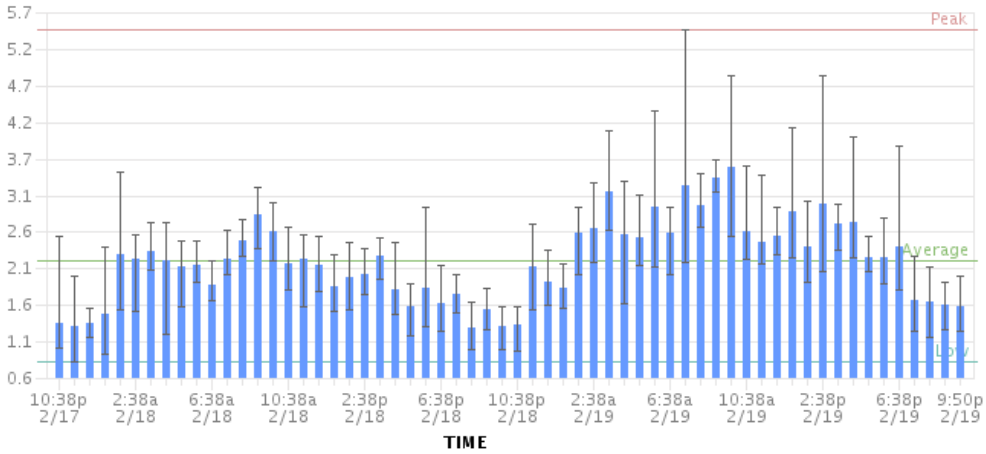
- Which pages are slow to load?
  - If it is a specific wiki page, attach the wiki source-code for that page
- Are they always slow to load, or is the slowness intermittent?

Monitoring data

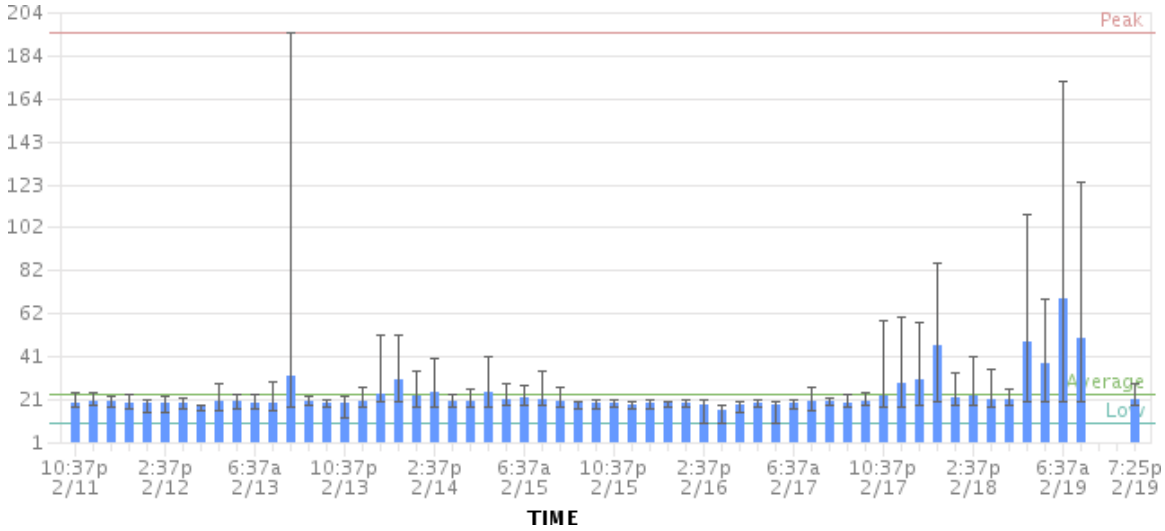
Before drilling down into individual problems, helps a lot to understand the nature of the performance problem. Do we deal with sudden spikes of load, or is it a slowly growing load, or maybe a load that follows a certain pattern (daily, weekly, maybe even monthly) that only on certain occasions exceeds critical thresholds? It helps a lot to have access to continuous monitoring data available to get a rough overview.

Here are sample graphs from the confluence.atlassian.com system, showing

Load
This graph shows the load for two consecutive days. The obvious pattern is that the machine is under decent load, which corresponds to the user activity, and there is no major problem.

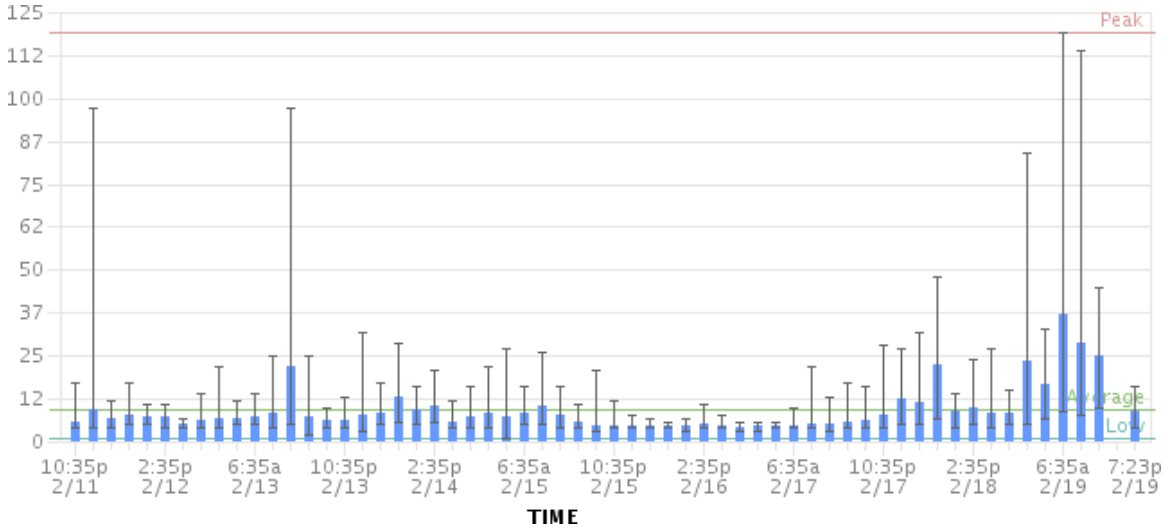


Resin Threads and Database Connections



Active number of Java Threads

These two charts show the active threads in the application server (first chart) and the size database connection pool (second chart). As you can see, there was a sudden spike of server threads and a corresponding spike of db-connections.



The database connection pool size

The database connection pool size peaked over 112, which happened to be more than the maximum number of connections the database was configured for (100). So it was no surprise that some requests to Confluence failed and many users thought it had crashed, since many requests could not obtain the crucial database connections.

We were able to identify this configuration problem quite easily just by looking at those charts. The next spikes were uncritical because more database connections were enabled.

The bottom line being: it helps a lot to monitor your Confluence systems continuously (we use Hyperic, for example), and it helps even more if you are able to send us graphs when you encounter problems.

Access logs

- How to audit Confluence - enabling user access logging, including redirecting the logs to a separate file
  - You can run this file through a log file analyser such as AWStats, or manually look through for pages which are slow to load.

Profiling and Logs

- Enable Confluence's built-in profiling for long enough to demonstrate the performance problem using Troubleshooting Slow Performance Using Page Request Profiling.
  - If a single page is reliably slow, you should make several requests to that page
  - If the performance problem is intermittent, or is just a general slowness, leave profiling enabled for thirty minutes to an hour to get a good sample of profiling times
- Find Confluence's standard output logs (which will include the profiling data above). Take a zip of the entire logs directory.
- Take a thread dump during times of poor performance

CPU Load

- If you are experiencing high CPU load, please install the YourKit profile and attach two profiler dumps taken during a CPU spike. If the CPU spikes are long enough, please take the profiles 30-60 seconds apart. The most common cause for CPU spikes is a virtual machine operating system.
- If the CPU is spiking to 100%, try Live Monitoring Using the JMX Interface, in particular with the Top threads plugin.

Instance Metrics and Scripts

- It is essential to understand the user access and usage of your instance. Please use the access log scripts and sql scripts to generate Usage statistics for your instance.

Next Step

Open a ticket on https://support.atlassian.com and attach all the data you have collected. This should give us the information we need to track down the source of your performance problems and suggest a solution. Please follow the progress of your enquiry on the support ticket you have created.

If your site is non-responsive, please use our Live Support during business hours once you have created the ticket to escalate your problem.

Access Log Scripts

The access log scripts are attached to this page. To use the scripts:

1. Unzip the 7z file.
2. Copy all the daily access logs to a folder called `logs`.
3. Run `Atlassian-processDailyLog.rb`. This will generate a `csv` file called `summary.csv` and several directories which contain the access logs of each defined user action.
4. Run the appropriate script `Atlassian-processDailyLog-hourly.rb <admin/comment/create/edit/search/rss>`.

Each script will generate a different `csv` file. For example, `Atlassian-processDailyLog-hourly.rb admin` will process the admin logs extracted in step 3.

5. Import the `csv` files to `www-log-Analysis.xls` (`summary.csv` to 'raw stats - daily' sheet and `admin.csv` to 'admin -hours' sheet, etc) to generate the load profiles and graphs. You may need to modify the number of rows in each sheet depending on the number of logs.



Note
All scripts are written in Ruby and assume the log file name contains the string 'confluence.atlassian.com-access.log'. Scripts need to be changed if another name is used. Modify the line: `filenameRegexp = Regexp.new('confluence.atlassian.com-access.log')`

Obtaining Confluence Instance Metrics

This page can be used as a guide to obtain detailed performance information of your instance.

Please read the Confluence Reporting HOWTO for information about the reporting capabilities of Confluence, including the {`sql`} macro, charting and security.

Users and usage

Users

What is the typical number of concurrent active users i.e. number of concurrent requests being processed?

- users with currently active requests
- users currently using Confluence: eg including reading a page, editing a page, viewing search results.
- users with sessions held in application server memory.
- users logged in *active users* (Note that Confluence uses "Remember Me" Session cookies and in my experience of Confluence, users never explicitly log out).
- define **user types** (viewer, editor, etc)

– light viewer
– rss reader
– searcher
– infrequent editor
– frequent editor
– administrator
– commenter

Usage

What is the average number of pages created per day, and similar usage stats (AWStat reports are a good starting place when User Access Logging is enabled)
To help interpret the raw access data, consider these important URL patterns:

Searches: `http://<host>/dosearchsite.action`
Rss requests: `http://<host>/createrssfeed.action`
Dashboard: `dashboard.action`
Creation: `createpage.action`
Editing: `http://<host>/pages/editpage.action`
Administrators: `http://<host>/admin/*`

Database usage statistics



The following SQL statements are for PostgreSQL database. You may need to adjust the queries to suit your database.

Note: specify the date range

Table sizes

```
1.SELECT relname, reltuples, relpages FROM pg_class ORDER BY relpages DESC ;
```

- Example result:

relname	reltuples	relpages
pg_toast_1404472	115842	30116
trackbacklinks	451197	23832
bodycontent	170462	18197

The column `reltuples` is the number of rows in the table, `relpages` is the number of 8 KB pages used by the table. Indexes are included in this list as well.

In this example, the bodycontent table includes 170462 rows and is approximately 142 MB (18197 * 8 KB) in size.

Content created per day

```
1.select contenttype, min(number_of_changes), max(number_of_changes), avg(number_of_changes)
2.from (
3.           select contenttype, date_trunc('day', creationdate) , count(*) as number_of_changes
4.           from content
5.           where content.creationdate > date '2007-01-01' and version = 1
6.           group by contenttype, date_trunc('day', creationdate)
7.              ) as dates
8.group by contenttype
```

- Example result:

contenttype	min	max	avg
DRAFT	4	6	5.0000000000000000
MAIL	1	1	1.00000000000000000000
COMMENT	1	54	20.5040000000000000
USERINFO	1	45	15.8112449799196787
SPACEDESCRIPTION	1	3	1.1403508771929825
PAGE	1	119	21.4593495934959350
BLOGPOST	1	64	5.5925925925925926

Content edited per day

```
1.select contenttype, min(number_of_changes), max(number_of_changes), avg(number_of_changes)
2.from (
3.           select contenttype, date_trunc('day', lastmoddate) as changedate, count(*) as
number_of_changes
4.           from content
5.           where content.creationdate > date '2007-01-01'
6.           group by contenttype, date_trunc('day', lastmoddate)
7.              ) as dates
8.group by contenttype
```

- Example result:

contenttype	min	max	avg
BLOGPOST	1	718	14.4705882352941176
COMMENT	1	73	23.5120000000000000
DRAFT	4	6	5.0000000000000000
MAIL	1	1	1.00000000000000000000
PAGE	1	4658	130.2650602409638554
SPACEDESCRIPTION	1	4	1.2033898305084746
USERINFO	1	48	16.7991967871485944

Number of existing pages

```
1.select contenttype, count(*) from content group by content.contenttype
```

- Example result:

contenttype	count
MAIL	7914
COMMENT	12983
SPACEDESCRIPTION	232
DRAFT	10
PAGE	81465

USERINFO	13782
BLOGPOST	3308

Number of links per page

```
1.select http, max(linkcount), min(linkcount), avg(linkcount), stddev_pop(linkcount),
stddev_samp(linkcount), var_pop(linkcount), var_samp(linkcount)
2.from
3.(
4.          select contentid, (links.destspacekey = 'http') as http, count(*) as linkcount
5.          from links group by contentid, (links.destspacekey = 'http')
6.) as links_per_page
7.group by http
```

- Example result:

http	max	min	avg	stddev_pop	stddev_samp	var_pop	var_sa
false	1994	1	5.8366957470010905	32.7082672608353032	32.7104967872521825	1069.8307472062305489	1069.9
true	189	1	2.9633190883190883	6.3609167066017375	6.3614831031752836	40.4612613483250948	40.468

http	max	min	avg	stddev_pop	stddev_samp	var_pop	var_samp
f	600	1	5.8769371011850501	21.0221241817454213	21.0245200019346061	441.9297051127255987	442.03044131174
t	695	1	3.9395946999220577	12.8270107456162932	12.8282606370439241	164.5322046681558531	164.56427097193

Number of characters per content body

```
1.select max(blength), min(blength), avg(blength), stddev(blength), variance(blength)
2.from (select length(body) as blength from bodycontent) as bodylengths
3.where blength > 0
```

- Example result:

max	min	avg	stddev	variance
488707	1	2826.5649320388349515	8858.740996699238	78477292.046599816739

(Note this query takes a long time to execute.)

Number of characters per *page* body

```
1.select max(blength), min(blength), avg(blength), stddev(blength), variance(blength)
2.from (select length(bodycontent.body) as blength
3.                from bodycontent, content
4.                where bodycontent.contentid = content.contentid and contenttype='PAGE'
5.             ) as bodylengths
6.where blength > 0
```

- Example result:

max	min	avg	stddev	variance
488707	1	3333.0885906386048069	9884.337162920180	97700121.150284961908

Attachments

```
1.select count(*), max(filesize), min(filesize), avg(filesize), stddev(filesize), sum(filesize)
from attachments;
```

- Example result:

count	max	min	avg	stddev	sum
16082	107431588	0	207641.309725158562	2447904.9322	3339287543

Attachments per page

List the stats for attachments per page, only for those pages that actually have attachments.

```
1.select count(*) as pages_with_attachments, avg(attachments_per_page), max(attachments_per_page),
min(attachments_per_page), stddev(attachments_per_page) from
2.( select count(*) as attachments_per_page from attachments group by attachments.pageid ) as app
```

- Example result:

pages_with_attachments	avg	max	min	stddev
4197	3.8317846080533715	231	1	10.7013051235493489

Configuration / plugin data stored in Bandana

- Just the global context

```
1.select count(*), sum(length(bandanavalue)) from bandana where bandanacontext = '_GLOBAL'
```

- Example result:

count	sum
84	47729

- All of the information

```
1.select count(*), sum(length(bandanavalue)) from bandana
```

- Example result:

count	sum
665	153094

Content

It is essential to obtain the typical configuration of database (#pages, #spaces, #registered users, etc), based on Global Stats Plugin

Home directory usage statistics

On Unix-based environments like Linux and Mac OS X, you can use the following commands to gather information about the home directory usage.

Size of home directory components

```
du -sh /path/to/home/directory/*
```

- Example output:

```
5.9G attachments
4.0K backups
13M bundled-plugins
8.0K config
4.0K confluence.cfg.xml
216K fonts
3.4M framework-bundles
2.9G index
12M plugin-cache
114M plugins-temp
412K resources
4.0K restore
201M temp
84M thumbnails
222M viewfile
```

Number of attachments, including all versions

```
find /path/to/home/directory/attachments -type f | wc -l
```

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.

In this page:

- Enabling Page-Request Profiling
- Profiling an Activity
- Example of a Profile
- Start Confluence with Profiling Enabled

Enabling Page-Request Profiling

To see just the slow performing macros, see Identifying Slow Performing Macros.

From **Confluence 2.7**, you can use the '**Logging and Profiling**' option to enable or disable profiling.

You need to have System Administrator permissions in order to perform this function.

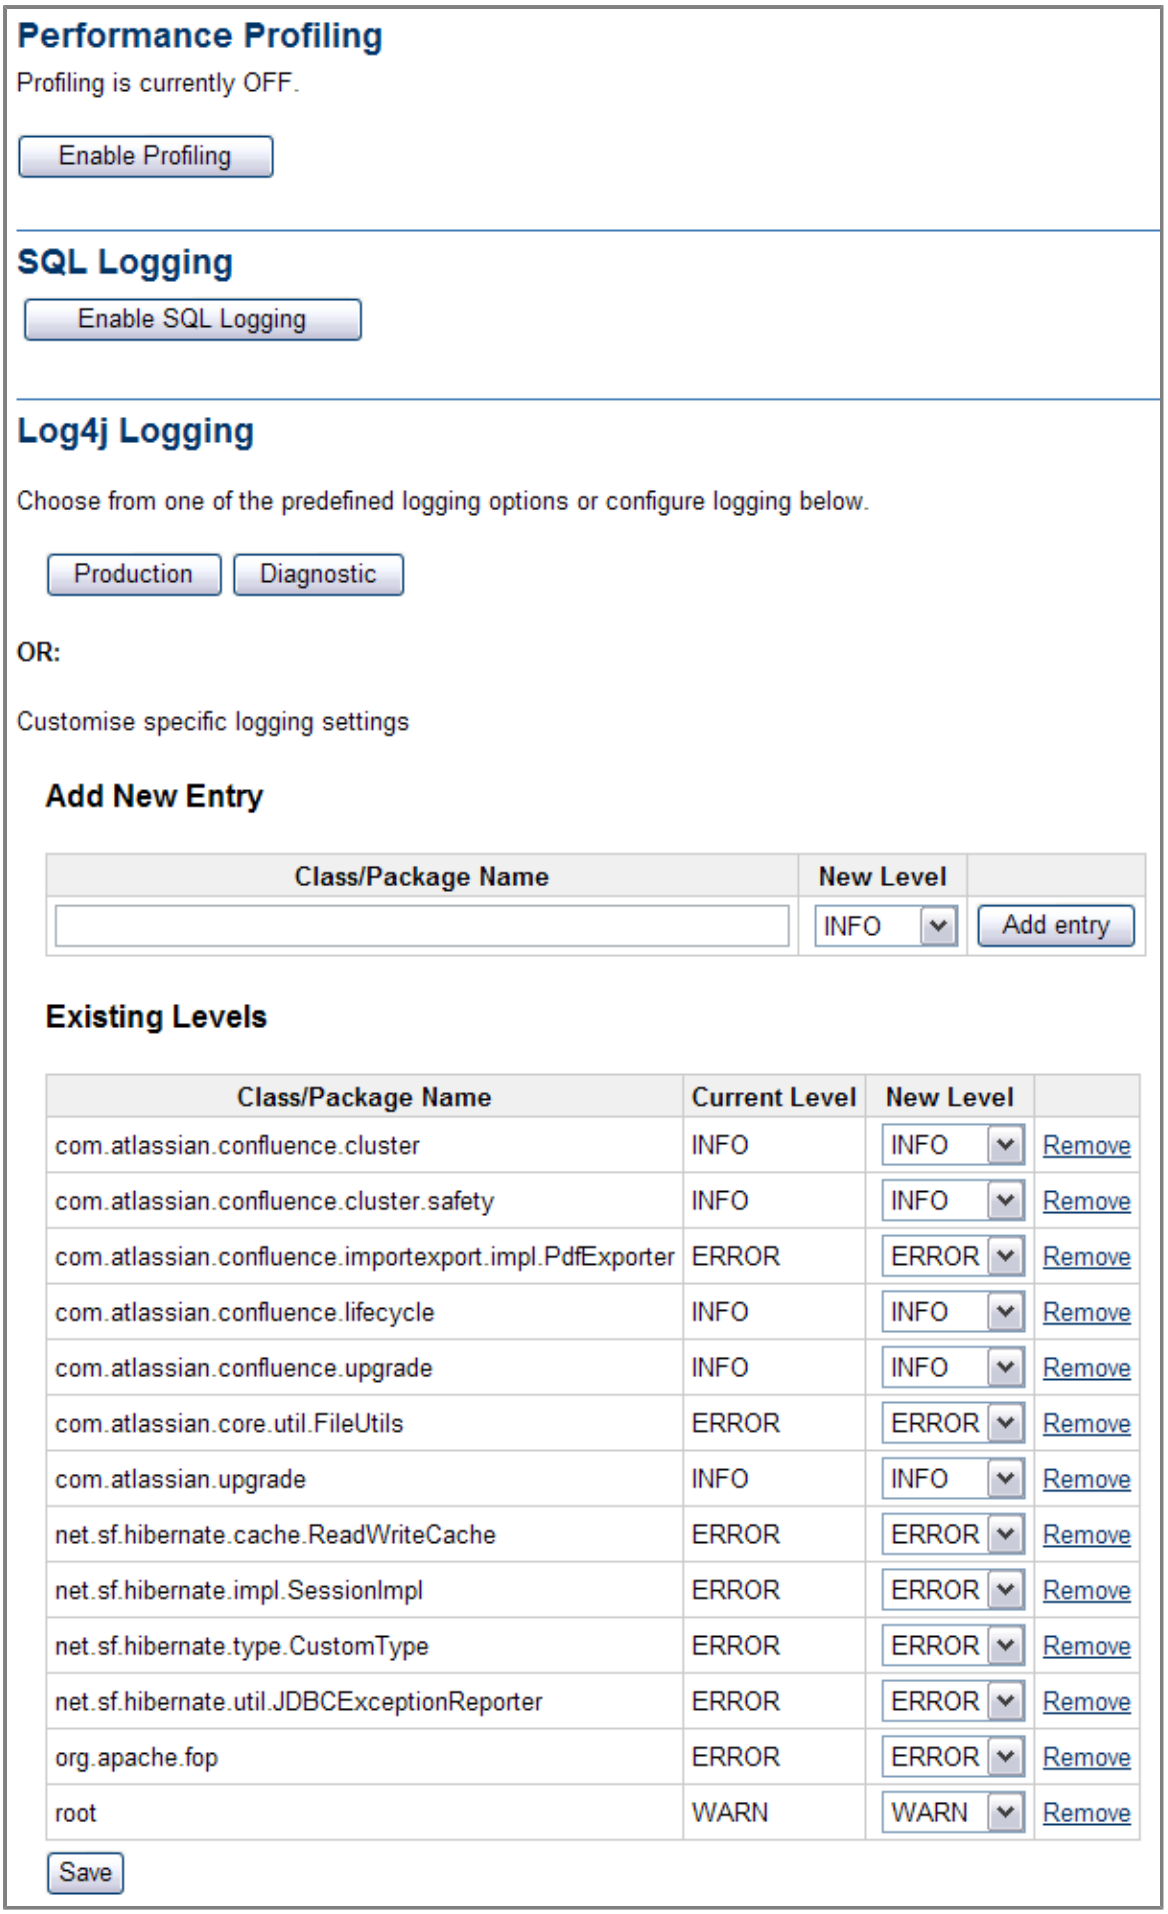
To enable page profiling,

1. Go to the '**Administration Console**' and click '**Logging and Profiling**' in the '**Administration**' section of the left-hand panel.
2. The '**Logging and Profiling**' screen appears. Click the '**Enable Profiling**' button.
   If profiling is already enabled, the button will be labelled 'Disable Profiling' instead.

To disable page profiling,

1. Go to the '**Administration Console**' and click '**Logging and Profiling**' in the '**Administration**' section of the left-hand panel.
2. The '**Logging and Profiling**' screen appears. Click the '**Disable Profiling**' button.
   If profiling is already disabled, the button will be labelled 'Enable Profiling' instead.

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())
           ...
```

Start Confluence with Profiling Enabled

There may be some situations where you may wish to have Confluence profiling enabled during startup. This may be useful if you restart often and may forget to enable profiling for Support/Trouble-shooting purposes.

Edit the file *CONFLUENCE_HOME\confluence\WEB-INF\web.xml*. You should see a stanza similar to the one below. Set the parameter value for **autostart** to **true**:

```
<filter>
        <filter-name>profiling</filter-name>
        <filter-class>com.atlassian.core.filters.ProfilingAndErrorFilter</filter-class>
        <init-param>
            <!-- specify the which HTTP parameter to use to turn the filter on or off -->
            <!-- if not specified - defaults to "profile.filter" -->
            <param-name>activate.param</param-name>
            <param-value>profile</param-value>
        </init-param>
        <init-param>
            <!-- specify the whether to start the filter automatically -->
            <!-- if not specified - defaults to "true" -->
            <param-name>autostart</param-name>
            <param-value>true</param-value>
        </init-param>
    </filter>
```

Remember to turn it back to **false** or your logs will grow very large.

RELATED TOPICS

Requesting Performance Support
Working with Confluence Logs

Compressing an HTTP Response within Confluence

Confluence supports HTTP GZip transfer encoding. This means that if a user's web browser supports it, Confluence will compress the data it sends to the user. This will speed up Confluence over slow or congested Internet links, and reduce the amount of bandwidth consumed by a

Confluence server.

Gzipping the HTTP Response is available in Confluence 1.4 and later.

You should turn on Confluence's GZip encoding if:

- Users are accessing Confluence over the Internet, or a WAN connection with limited bandwidth.
- You wish to reduce the amount of data transfer between the Confluence server and client.

If you are accessing Confluence over a Local Area Network or over a particularly fast WAN, you may wish to leave GZip encoding disabled. If the network is fast enough that transferring data from Confluence to the user isn't a limiting factor, the additional CPU load caused by having to compress each HTTP response may in fact slow Confluence down.

Known issues in Confluence 2.7 and earlier
There are known issues with the GZip filter and memory consumption evident in versions 2.7 of Confluence and earlier (CONF-9930). If you are running a large instance of Confluence 2.7 or earlier and frequently experiencing 'out of memory' errors, we recommend that you do not enable HTTP compression. These issues have been resolved in Confluence 2.8.

Enabling HTTP Compression

1. Go to the Confluence '**Administration Console**'. To do this:
    - Open the '**Browse**' menu and select '**Confluence Admin**'. The 'Administration Console' view will open.
2. Select '**General Configuration**' in the left-hand panel.
3. Enable '**Compress HTTP Responses**'.

In Confluence 2.8 and later, you can configure which types of content are compressed within Confluence. By default, the following mime types will be compressed:

- text/htmltext
- javascript
- text/css
- text/plain
- application/x-javascript
- application/javascript

If you wish to change the types of content to be compressed, add a replacement `urlrewrite-gzip-default.xml` file within the `WEB-INF/classes/com/atlassian/gzipfilter/` directory in your Confluence Installation Directory. A sample file is provided as an attachment. Generally speaking, it is unlikely that you will need to alter this file.

RELATED TOPICS

Performance Tuning
Confluence Administrator's Guide

Performance Testing Scripts

Load Testing Confluence

This page contains scripts and hints on load-testing your Confluence installations.

Contents

Introduction

Before making a new Confluence instance available to your users it is useful to get a feel for how it will perform under your anticipated load and where you may need to consider improving your configuration to remove bottlenecks. Likewise, before making changes to your Confluence instance it would again be useful to assess the impact of these changes before making them live in a production context.

This kind of testing is not an exact science but the tools and process described here are intended to be a straightforward, configurable and extensible way of allowing you to begin this kind of load testing.

It will rarely be the case that these scripts will perform representative testing for you 'out of the box'. But either through configuration or by extending the scripts it should be possible to build an appropriate load test.

Load testing scripts are not designed for a production environment
The load testing scripts will update the data within the targeted Confluence instance and are not designed to be run against a production server. If you want to load test your production environment you will need to perform these tests on a backup of your data and restore your real data after the tests.

Setup

You will need the following -

- A Confluence server, set up and running with an admin user. The scripts assume a default username and password for this user: 'admin'/'admin'.
- Ensure the Confluence Remote API is enabled in the administration options. See Enabling Remote APIs for details on how to configure this.
- Apache JMeter (currently version 2.3.4).
- The load testing scripts and resources which are available in our public Maven repository - version 3.2.1



The Maven Repository has both ZIP and TAR archives. These archives contain the same files - if in doubt, download the ZIP file archive.



Users have reported problems when using the Windows built-in UNZIP utility. Please use a 3rd party unzip program such as WinZIP to extract these Performance Tests

The test scripts have been updated to work with Confluence 3.2 in version 3.2.1. Using an older version of the tests will result in errors when running the test.

Quick, Just Tell Me How To Run It.

If you don't want to read the rest of this document, here are the main points:

1. Create the test data:

```
<jmeter location>/bin/jmeter -n -t jmeter-test-setup.jmx -Jscript.base=<scripts location>
-Jspace.zip=<path to a demo space ZIP file> \
-Jadmin.user=<username> -Jadmin.pass=<password>
```

2. Run the test:

```
<jmeter location>/bin/jmeter -n -t jmeter-test-fixedload.jmx -Jscript.base=<scripts
location>
```

The remainder of this document is just an elaboration of those two steps.



For information on how to use JMeter please refer to the manual

Creating the Test Data

A known data set is required to run the testing against. By default this is the Confluence demo space (space key = DS) although this can be changed (more on this later). If you decide to use the Confluence demo space, ensure that the group "confluence-users" is able to update content in this space.

The script `jmeter-test-setup.jmx` is used to:

- create a set of users to be used in the test
- import the Confluence demo space for running tests against.

You should first ensure that you don't already have the demo space (key = DS) on your test instance. Delete it if you do.

Run the script from the `performance-testing` directory as follows:

```
<jmeter location>/bin/jmeter -n -t jmeter-test-setup.jmx -Jscript.base=<scripts location>
-Jspace.zip=<path to a space export.zip> \
-Jadmin.user=<username> -Jadmin.pass=<password>
```

Where:

- `<scripts location>` is the absolute path to where you expanded the scripts e.g. `/Users/YourName/Download/performanceTest`. This defaults to the current directory. This is needed for the script to find its external resources and must be specified absolutely since JMeter occasionally does unexpected things with the working directory when it is running.
- `<path to a space export.zip>` is the absolute path to the space export zip you want to be used in your testing. For example, the path to `demo-site.zip` as found in your Confluence distribution or source: `<confluence`

```
install>/confluence/WEB-INF/classes/com/atlassian/confluence/setup/demo-site.zip
```

- <username>> and <password> are the username and password for an admin user that is able to create Confluence users and to import spaces.

By default the setup process will create 250 users — 50 each of the following formats: tstreader<n>, tstcommentor<n>, tsteditor<n>, tstcreator<n> and tstsearcher<n>. The password for each matches the username.

A typical run of the setup script will only take a few seconds.

Removing the Test Data

You can reverse the effects of the setup script by setting the `remove.data` parameter to `true`, e.g.

```
<jmeter location>/bin/jmeter -n -t jmeter-test-setup.jmx -Jscript.base=<scripts base>
-Jremove.data=true -Jadmin.user=<username> -Jadmin.pass=<password>
```

Setup Script Parameters

You can modify the behaviour of the setup script via JMeter parameters. These are supplied on the command line in the form `-J<parameter name>=<parameter value>`.

Parameter	Default	Explanation
script.base	.	The absolute path to the script. Default to the current directory.
space.zip	N/A	The absolute path to space export zip file to be imported as test data.
remove.data	false	Run the script in reverse — remove all test data.
admin.user	admin	The admin user name used to import data and create users.
admin.pass	admin	The password for the admin user.
confluence.context	confluence	The confluence webapp context.
confluence.host	localhost	The address or host name of the test instance.
confluence.port	8080	The port of the test instance.
space.key	ds	The space key for the space import that will be tested against.
space.setup	true	Control whether the test space will be created (or removed).
commentor.max	250	The number of users to be created for making comments.
creator.max	250	The number of users to be created for adding pages.
editor.max	250	The number of users to be created for editing existing pages.
reader.max	250	The number of users to be created for viewing existing pages.
searcher.max	250	The number of users to be created for performing searches.
resource.max	250	The number of users to be created for downloading site resources.
attachments.max	250	The number of users to be created for downloading attachments.

Setup Script Output

On the console you will see no obvious indication of success or otherwise. JMeter will output something similar to this:

```
Created the tree successfully
Starting the test @ Mon Apr 14 17:35:08 EST 2008 (1208158508222)
Tidying up ... @ Mon Apr 14 17:35:08 EST 2008 (1208158508928)
... end of run
```

The `scripts location/results` directory will contain the file `jmeter-result-setuptest.jtl`. There were failures or errors if there are any assertions in this file that have the value `true` for the failure or error element, e.g.

```
1.<assertionResult>
2.<name>Manage Users</name>
3.<failure>true</failure>
4.<error>false</error>
5.<failureMessage>Test failed: URL expected to contain /browseusers.action/</failureMessage>
6.</assertionResult>
```

Running the Test

The test script itself will put Confluence under a fixed load. Each thread group will attempt to do a certain amount of work for a prescribed period of time (30 minutes by default). This is by design so that load during test runs can accurately be compared against each other.

Execute the test as follows:

```
<jmeter location>/bin/jmeter -n -t jmeter-test-fixedload.jmx -Jscript.base=<scripts location>
```

Where:
`<scripts location>` is the absolute path to where you extracted the scripts e.g. `/Users/YourName/Download/performanceTest`. This is needed for the script to find its external resources.

Test Behaviour

The test has a number of parameters to tweak its behaviour but generally speaking it has the rough format of:

- 5 groups of users - readers, commentors, searchers, editors and creators.
  - readers simply view a set of individual pages or browse space functionality.
  - commentors add comments to a set of pages.
  - searchers perform searches on a fixed set of keywords.
  - editors make small additions to the end of a set of pages.
  - creators add new pages to a particular space.
- Each individual user in each group will repeat for a fixed amount of time with a small pause between each request.

Note that there is **no execution of JavaScript** by the client. Keep this in mind if you use this test to gauge Confluence performance in a production environment.

There is also very little use of permissions in these tests. All data involved is accessible to all of the test users.

Test Script Parameters

You can modify the behaviour of the test script via JMeter parameters. These are supplied on the command line in the form `-J<parameter name>=<parameter value>`.

Parameter	Default	Explanation
script.base	.	The absolute path to the script. Defaults to the current working directory.
confluence.context	confluence	The confluence webapp context.
confluence.host	localhost	The address or host name of the test instance.
confluence.port	8080	The port of the test instance.
create.page.prefix	Nihilist	The title prefix for any created page e.g. Nihilist00001.
script.runtime	1800	The amount of time the script will run for in seconds.

Test Thread Parameters

Parameter	Default	Explanation
threads.reader	15	Number of readers.
pause.reader	2000	The approximate (within 500ms) millisecond pause between reader repeats.
threads.searcher	8	Number of searchers.
pause.searcher	2000	The approximate (within 500ms) millisecond pause between searcher repeats.
threads.creator	3	Number of page creators.
pause.creator	2000	The approximate (within 500ms) millisecond pause between creator repeats.
threads.editor	3	Number of page editors.
pause.editor	2000	The approximate (within 500ms) millisecond pause between editor repeats.

threads.commentor	4	Number of page commentors.
pause.commentor	2000	The approximate (within 500ms) millisecond pause between commentor repeats.



In version 3.0 of the tests, it's now possible to control the percentage executions of certain actions. These percentages are defined in the "Thread Details" configuration screen.

So with the default parameters, you are emulating a load on Confluence of 33 concurrent users who will each be hitting the server approximately every 2 seconds (16 users per second).

23 of these users are read only (searchers or readers) and 10 of them are read/write — 11 read only users per second and 5 read/write users per second.

Test Script Output

During the run of the test script Jmeter will output progress to the console of the form:

```
Created the tree successfully
Starting the test @ Fri Apr 18 00:07:39 EST 2008 (1208441259523)
Display Summary Results During Run + 462 in 77.6s = 5.9/s Avg: 1564 Min: 18 Max: 33738 Err: 1
(0.22%)
Display Summary Results During Run + 1338 in 189.9s = 7.0/s Avg: 3596 Min: 24 Max: 34545 Err: 0
(0.00%)
Display Summary Results During Run = 1800 in 257.6s = 7.0/s Avg: 3074 Min: 18 Max: 34545 Err: 1
(0.06%)
Display Summary Results During Run + 1046 in 200.9s = 5.2/s Avg: 4529 Min: 40 Max: 50461 Err: 0
(0.00%)
Display Summary Results During Run = 2846 in 438.2s = 6.5/s Avg: 3609 Min: 18 Max: 50461 Err: 1
(0.04%)
Display Summary Results During Run + 677 in 201.2s = 3.4/s Avg: 6638 Min: 46 Max: 27636 Err: 0
(0.00%)
Display Summary Results During Run = 3523 in 618.1s = 5.7/s Avg: 4191 Min: 18 Max: 50461 Err: 1
(0.03%)
Display Summary Results During Run + 561 in 197.5s = 2.8/s Avg: 8326 Min: 171 Max: 39494 Err: 0
(0.00%)
Display Summary Results During Run = 4084 in 798.3s = 5.1/s Avg: 4759 Min: 18 Max: 50461 Err: 1
(0.02%)
Display Summary Results During Run + 555 in 199.2s = 2.8/s Avg: 8247 Min: 160 Max: 45270 Err: 0
(0.00%)
Display Summary Results During Run = 4639 in 978.0s = 4.7/s Avg: 5177 Min: 18 Max: 504
```

Garbage Collector Performance Issues



This document relates broadly to memory management with Sun's Hotspot JVM. These are recommendations based on Support's successful experiences with customers and their large Confluence instances.

Summary

- Set the Eden space up to 30-50% of the overall heap: -XX:NewSize=<up to half of your Xmx value>
- Use a parallel collector: `-XX:+UseParallelOldGC` (make sure this is Old GC)
- Disable explicit garbage collection triggered from distributed remote clients (`-XX:+DisableExplicitGC`)
- set the minimum and maximum Xmx and Xms values as the same (eg. `-Xms1024m -Xmx1024m`)
- Turn on GC logging (add the flags `-verbose:gc -Xloggc:<full-path-to-log> -XX:+PrintGCTimeStamps -XX:+PrintGCDetails`) and submit the logs in a support ticket
- Use Java 1.6
- Read below if heap > 2G

See Configuring System Properties for how to add these properties to your environment.

Background

Performance problems in Confluence generally manifest themselves in either:

- frequent or infrequent periods of viciously sluggish responsiveness, which requires a manual restart, or, the application eventually and almost inexplicably recovers
- some event or action triggering a non-recoverable memory debt, which in turn envelops into an application-fatal death spiral (Eg. overhead GC collection limit reached, or Out-Of-Memory).
- generally consistent poor overall performance across all Confluence actions

There are a wealth of simple tips and tricks that can be applied to Confluence, that can have a significantly tangible benefit to the long-term stability, performance and responsiveness of the application.

On this page:

- Summary
- Background
- Why Bad Things Happen
- Appreciate how Confluence and the JAVA JVM use memory
- Memory is contiguous
- Figure out which (default) collector implementation your vendor is using
- Use the Parallel Garbage Collector
- Disable remote (distributed) garbage collection by JAVA clients
- Virtual Machines are Evil
- Use Java 1.6
- Use -server flag
- If using 64bit JRE for larger heaps, use CompressedOops
- Use NUMA if on SPARC, Opteron or recent Intel (Nehalem or Tukwila onwards)
- Use 32bit JRE if Heap < 2G
- JVM coredumps can be instigated by memory pressures
- Artificial Windows memory limit
- Instigate useful Monitoring techniques
- Tuning the frequency of full collections
- Performance Tuning works

Why Bad Things Happen

Confluence is basically a gel. Multiple applications, data-types, social networks and business requirements can be efficiently amalgamated together, leading to more effective collaboration. The real beauty of Confluence, however, is it's agility to mold itself into your organizations' DNA - your existing business and cultural processes, rather than the other way around - your organization having to adapt to how the software product works.

The flip side of this flexibility is having many competing demands placed on Confluence by it's users. Historically, this is an extraordinarily broad and deep set of functions, that really, practically can't be predicted for individual use cases.

The best mechanism to protect the installation is to place Confluence on a foundation where it is fundamentally more *resilient* and able to react and cope with competing user requirements.

Appreciate how Confluence and the JAVA JVM use memory

The java memory model is naive. Compared to a unix process, which has four *intensive* decades of development built into time-slicing, inter-process communication and intelligent deadlock avoidance, the JAVA thread model really only has 10 years at best under it's belt. As it is also an interpreted language, particular idiosyncrasies of the chosen platform Confluence is running can also influence how the JRE reacts. As a result it is sometimes necessary to *tune* the jvm parameters to give it a "hint" about how it should behave.

There are circumstances whereby the JAVA jvm will take the 'mediocre' option in respect to resource contention and allocation, and, 'struggle' along with ofttimes highly impractical goals. For example, The JRE will be quite happy to perform at 5 or 10% of optimum capacity if it means overall application stability and integrity can be ensured. This often translates into periods of extreme sluggishness, which effectively means that the application isn't stable, and isn't integral (as it cannot be accessed).

This is mainly because JAVA shouldn't make assumptions on what kind of runtime behavior an application needs, but, it's plain to see that the charter is to assume 'business-as-usual' for a wide range of scenarios and really only react in the case of 'dire' circumstances.

Memory is contiguous

The Java memory model *requires* that memory be allocated in a *contiguous* block. This is because the heap has a number of side data structures which are indexed by a scaled offset (ie n*512 bytes) from the start of the heap. For example, updates to references on objects within the heap are tracked in these "side" data structures.

Consider the differences between:

1. Xms (the *allocated* portion of memory)
2. Xmx (the *reserved* portion of memory)

Allocated memory is fully backed, memory mapped physical *allocation* to the application. That application now owns that segment of memory.

Reserved memory (the difference between Xms and Xmx) is memory which is *reserved* for use, but not physically mapped (or backed) by memory. This means that, for example, in the 4G address space of a 32bit system, the *reserved* memory segment can be used by other applications, but, because JAVA requires *contiguous* memory, if the *reserved* memory requested is occupied, the OS must swap that memory out of the reserved space either to another non-used segment, or, more painfully, it must swap to disk.

Permanent Generation memory is also contiguous. The net effect is even if the system has vast quantities of *cumulative* free memory, Confluence demands *contiguous* blocks, and consequently undesirable swapping may occur if segments of requested size do not exist. See Causes of OutOfMemoryErrors for more details.

Please be sure to position Confluence within a server environment that can successfully complete competing requirements (operating system, contiguous memory, other applications, swap, and Confluence itself).

Figure out which (default) collector implementation your vendor is using

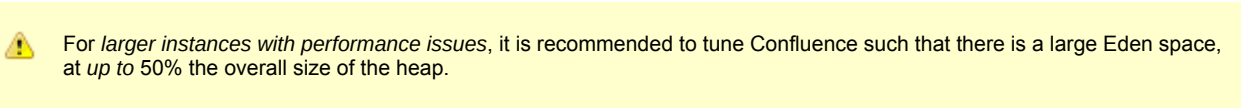
Default JVM Vendor implementations are subtly different, but in production can differ enormously.

SUN *by default* splits the heap into three *spaces*

1. Eden (Nursery, or Scavenger)
2. Tenured (Old)
3. Permanent Generation (classes & library dependencies)

Objects are central to the operation of Confluence. When a request is received, the Java runtime will create new objects to fulfill the request in the Eden Space. If, after some time, those objects are still required, they may be moved to the Tenured (Old) space. But, typically, the *overwhelming majority* of objects created die young, within the Eden Space. These are objects like method local references within a while or for loop, or Iterators for scanning through Collections or Sets.

But in IBM J9 the default policy is for a single, contiguous space - one large heap. The net effect is that for large Websphere environments, garbage collection can be terribly inefficient - and culpable to suffer outages during peak periods.



`-XX:NewSize=XXXm` where XXX is the size in megabytes, is the command line parameter. `-XmnXXXm` can also be used interchangeably. Ie. -XX:NewSize=700m, -Xmn700m

By setting a larger NewSize, the net effect is that the JRE will spend less time garbage collecting, clearing dead memory references, compacting and copying memory between spaces, and more time *doing actual work*.

Use the Parallel Garbage Collector

Confluence out of the box, and Sun Java as default, uses the *serial* garbage collector on the Full Tenured heap. The Eden space is collected in parallel, but the Tenured is not. This means that at a time of load, if a full collection event occurs, since the event is a 'stop-the-world' serial event, *all application threads* other than the garbage collector thread are taken off the CPU. This can have severe consequences if requests continue to accrue during these 'outage' periods. As a rough guide, for every gigabyte of memory allocated, allow a full second (exclusive) to collect.

If we parallelize the collector on a multi-core/multi-cpu architecture instance, we not only *reduce* the total time of collection (down from whole seconds to fractions of a second) but we also improve the *resiliency* of the JRE in being able to recover from high-demand occasions.

Additionally, Sun provide a CMS, Concurrent Mark-Sweep Collector (-XX:+UseConcMarkSweepGC), which is more optimized for higher-throughput, server-grade instances. As a general rule the Parallel Collector (-XX:+UseParallelOldGC) is suitable for most installations, if you are keen to extract the best performance available, then the CMS collector is an option.

Disable remote (distributed) garbage collection by JAVA clients

Many clients integrate 3rd party, or their own custom applications to interrogate, or add content to Confluence via it's RPC interface. The Distributed Remote Garbage Collector in the client uses RMI to trigger a remote GC event in the server, Confluence. Unfortunately, as of this writing, a System.gc() call via this mechanism triggers a full, serial collection of the entire Confluence heap (as it needs to remove references to remote client objects in it's own deterministic object graph). This is a deficiency in the configuration and/or implementation of the JVM, and, has the potential to impact severely if the remote client is poorly written, or operating within a constricted JVM.

This can be disabled by using the flag `-XX:+DisableExplicitGC` at startup.

Virtual Machines are Evil

Vmware Virtual Machines, whilst being extremely convenient & fantastic, also obviously being a very strong growth segment of the industry, cause particular problems for JAVA applications because it's very easy for host operating system resource constraints (ie, temporarily insolvent memory availability, or I/O swapping), to cascade into the JAVA VM and manifest itself in extremely unusual, frustrating and seemingly illogical problems. We already document some disk I/O metrics with VMware images, and, although we now *officially* support the use of virtual instances, we absolutely do not recommend them unless maintained correctly.

This is not to say that vmware instances cannot be used, but, they must be used with due care, proper maintenance and configuration. Besides, if you are reading this document because of poor performance, the first port of call should be to remove any virtualization - emulation will never beat the real thing, and at the least introduces black box variability into the system.

Use Java 1.6

Java 1.6 is generally regarded via public discussion to have an approximate 20% performance improvement over 1.5. Indeed, our own internal testing revealed this statistic to be credible. 1.6 is compatible for all supported versions of Confluence, and we **strongly recommend** that installations not using 1.6 migrate.

Use -server flag

The hotspot server JVM has specific code-path optimizations which yield an approximate 10% gain over the client version. Most installations *should* already have this selected by default, but it is still wise to force it with -server, especially on some Windows machines.

If using 64bit JRE for larger heaps, use CompressedOops

For every JDK release Sun also build a "Performance" branch in which specifically optimized performance features can be enabled; It is available on the Sun Java SE page after a brief survey. These builds are certified production grade.

Some blogs have suggested a 25% performance gain AND reduction in Heap size when using this parameter. The use and function of the `-XX:+UseCompressedOops` parameter is more deeply discussed on Sun's Official Wiki (Co-incidentally, using Confluence!)

Use NUMA if on SPARC, Opteron or recent Intel (Nehalem or Tukwila onwards)

-XX:+UseNUMA flag enables the java heap to take advantage of Non-Uniform-Memory-Architectures. JAVA will place data structures relevant to the thread which it owns / operates on, in memory locations closest to that particular processor. Depending on the environment, gains can be substantial. Intel market NUMA as Quick Path Interconnect™.

Use 32bit JRE if Heap < 2G

Using a 64bit JRE when the heap is under 2G will hurt substantially in heap and performance, as nearly every object, reference, primitive, class and variable will cost twice as much to address.

A 64bit JRE/JDK is only recommended if heaps greater than 2G are required. If so, use CompressedOops.

JVM coredumps can be instigated by memory pressures

If your instance of Confluence is throwing Java coredumps, it's known that memory pressure and space / generation sizings can influence the frequency and/or occurrence of this phenomena.

If your Tomcat process completely disappears and the logs record similar to:

```
#
# An unexpected error has been detected by HotSpot Virtual Machine:
#
#  SIGSEGV (0xb) at pc=0xfe9bb960, pid=20929, tid=17
#
# Java VM: Java HotSpot(TM) Server VM (1.5.0_01-b08 mixed mode)
# Problematic frame:
# V  [libjvm.so+0x1bb960]
#

---------------  T H R E A D  ---------------

Current thread (0x01a770e0):  JavaThread "JiraQuartzScheduler_Worker-1" [_thread_in_vm, id=17]

siginfo:si_signo=11, si_errno=0, si_code=1, si_addr=0x00000000

Registers:
 O0=0xf5999882 O1=0xf5999882 O2=0x00000000 O3=0x00000000
 O4=0x00000000 O5=0x00000001 O6=0xc24ff0b0 O7=0x00008000
 G1=0xfe9bb80c G2=0xf5999a48 G3=0x0a67677d G4=0xf5999882
 G5=0xc24ff380 G6=0x00000000 G7=0xfdbc3800 Y=0x00000000
 PC=0xfe9bb960 nPC=0xfe9bb964
```

You should upgrade the JVM. See SIGSEGV JVM Crash.

Artificial Windows memory limit

On Windows, the maximum heap allocatable to the Tomcat 32bit wrapper process is around 1400m. If the instance is allocated too close to this limit, **chronic and decimating garbage collection is likely to result**, often producing JAVA core dumps similar to:

```
#
# A fatal error has been detected by the Java Runtime Environment:
#
# java.lang.OutOfMemoryError: requested 8388608 bytes for GrET in
C:\BUILD_AREA\jdk6_18\hotspot\src\share\vm\utilities\growableArray.cpp. Out of swap space?
#
#  Internal Error (allocation.inline.hpp:39), pid=11572, tid=12284
#  Error: GrET in C:\BUILD_AREA\jdk6_18\hotspot\src\share\vm\utilities\growableArray.cpp
#
# JRE version: 6.0_18-b07
# Java VM: Java HotSpot(TM) Server VM (16.0-b13 mixed mode windows-x86 )
# If you would like to submit a bug report, please visit:
#   http://java.sun.com/webapps/bugreport/crash.jsp
#

---------------  T H R E A D  ---------------

Current thread (0x002af800):  GCTaskThread [stack: 0x00000000,0x00000000] [id=12284]
```

or,

```
#
# A fatal error has been detected by the Java Runtime Environment:
#
# java.lang.OutOfMemoryError: requested 123384 bytes for Chunk::new. Out of swap space?
#
#  Internal Error (allocation.cpp:215), pid=10076, tid=4584
#  Error: Chunk::new
#
# JRE version: 6.0_18-b07
# Java VM: Java HotSpot(TM) Server VM (16.0-b13 mixed mode windows-x86 )
# If you would like to submit a bug report, please visit:
#   http://java.sun.com/webapps/bugreport/crash.jsp
#

---------------  T H R E A D  ---------------

Current thread (0x6ca4d000):  JavaThread "CompilerThread1" daemon [_thread_in_native, id=4584,
stack(0x6cd10000,0x6cd60000)]
```

Workarounds include:

- upgrading to linux
- switching to the 64 bit Tomcat wrapper (this is not supported).
- reducing memory allocation to the Tomcat process (try backing off 100M a piece and observe the results)

Instigate useful Monitoring techniques

At all times, the best performance tuning recommendations are based on current, detailed metrics. This data is easily available and configurable, and helps us at Atlassian, **tremendously**, when diagnosing reported performance regressions.

1. enable JMX monitoring
2. enable Confluence Access logging
3. enable Garbage Collection Logging
4. Take Thread dumps at the time of regression. If you can't get into Confluence, you can take one externally.
5. Jmap can take a memory dump in real time without impacting the application. Syntax: `jmap -heap:format=b <process_id>`

Great tools available include:

- The very excellent VisualVM, Documentation.
- Thread Dump Analyzer - great all round thread debugging tool, great for identifying deadlocks.
- Samurai, an excellent alternative thread analysis tool, great for iterative dumps over a period of time.
- GC Viewer - getting a bit long in the tooth, but is a good mainstay for GC analysis.
- GChisto - A new GC analysis tool written by members of the Sun Garbage Collection team.

Documentation:

- Sun's White Paper on Garbage Collection in JAVA 6.
- Sun's state-of-the-art JavaOne 2009 session on garbage collection (registration required).
- IBM stack: Java 5 GC basics for WebSphere Application Server.
- An Excellent IBM document covering native memory, thread stacks, and how these influence memory constricted systems. Highly recommended for additional reading.
- The complete list of JRE 6 options



Atlassian recommends at the very least to get VisualVM up and running (You need JMX), and to add Access and Garbage Collection logging.

Tuning the frequency of full collections

The JVM will generally only collect on the full heap when it has no other alternative, because of the relative size of the Tenured heap (it is typically larger than Eden), and the natural probability of objects within tenured not being eligible for collection. (ie. they are still alive).

Some installations can trundle along, only ever collecting in Eden space. As time goes on, some object will survive the initial Eden object slaughterhouse and be promoted to Tenured. At some point, it will be dereferenced and no longer reachable by the deterministic, directed object graph. However, the occupied memory will still be held in limbo as "dead" memory until a collection occurs in the Tenured space to clear and compact the space.

It is not uncommon for moderately sized Confluence installations to reclaim as much as 50% of the current heap size on a full collection; This is because full collections occur so infrequently. By reducing the occupancy fraction heap trigger, this means that more memory will be available at any time, meaning less swapping / object collections will occur *when needed*, ie, during the busy hour.

Atlassian would classify frequency tuning on collections as an **advanced** topic for further experimentation, and is provided for informational purposes only. Unfortunately, it's impractical for Atlassian to support this degree of 'ergonomics'.

Performance Tuning works

Atlassian has a number of high profile and some *extremely* high demanding, mission-critical clients who have successfully, through some trial and error, applied these recommendations to production instances and have very positively improved their Instances. For more information, please lodge a support case at support.atlassian.com.

Scheduled Jobs

This page provides a quick overview of the jobs that are scheduled to run regularly in your Confluence instance.

Job Name	Description
backupJob	performs a site backup
mailQueueFlushJob	sends notifications that have been queued up
referralQueueFlushJob	referrals to Confluence pages are queued up. This job writes this referrals to the database
taskQueueFlushJob	flushes the task queue
cleanTempDirectoryJob	this cleans up temp files created in Confluence home temp directory (created by exports etc.)
dailyReportJob	sends out an email summary of all changes in Confluence to all subscribers
clearOldMailErrorsJob	notifications that fail to send due to errors are added to the mail error queue. This job resets this clear periodically.
indexQueueFlushJob	each content update to Confluence needs to be updated in index so search results are accurate. This job flushes changes to the index.
indexOptimizerJob	index optimization is performed to compact the index and maintain searching performance. This task is expensive and does not need to be performed too regularly. If you see Confluence performance deteriorate around 3pm, you can try scheduling this job for 3am only and check that search performance remains reasonable.
indexQueueCleanJob	this job is responsible for periodically triggering an Index Queue clean to ensure that size of the index queue does NOT grow indefinitely.
mailPollJob	polls POP accounts on all spaces that have them configured.
clusterSafetyJob	ensures that only one cluster is ever writing to the database at one time. For non-clustered instances, this job is still useful for alerting customers that have accidentally deployed two instances of Confluence against the same database.

Search

 Setup External Search Tool To Index Confluence

 Setup Confluence To Index External Sites

Setup Confluence To Index External Sites

Confluence Indexing External Sites

Confluence cannot easily index external sites due to technical reasons, but there are two alternatives:

1. Embed External Pages Into Confluence
2. Replace Confluence Search

Technical Reasons

Confluence indexes pages using a customised Lucene search engine that returns matching pages, mail and blog posts for which the searcher has view permission. It would require significant source code modifications to enable Confluence to process search results from external pages, as the indexing process has been customised to utilise internal Confluence metadata. Note that users can still index content from new attachment filetypes.

Embed External Pages Into Confluence

If you only have a small number of external sites to index, you may prefer to enable the HTML-include Macro and use it embed the external content inside normal Confluence pages.

Replace Confluence Search

Use your own programmer resources to replace Confluence's internal search with a crawler that indexes both Confluence and external sites. This advanced option is easier than modifying the internal search engine. It requires removing Confluence internal search from all pages and replacing the internal results page with your own crawler front-end.

1. Setup a replacement federated search engine to index the Confluence site, as well as your other sites, and provide the results that way. You would need to host a web crawler, such as these open-source crawlers. Note that you can perform a search in Confluence via the remote API
2. Replace references to the internal search by modifying the site layout so that it links to your search front-end
3. Host another site containing the search front-end. You may wish to insert it into a suitable context path in your application server so that it appears to be from a path under Confluence. Tomcat sets Confluence's paths from the Confluence install\confluence\WEBINF\web.xml file.

RELATED TOPICS

Setup External Search Tool To Index Confluence

Setup External Search Tool To Index Confluence

Any web crawler can be configured to index Confluence content, for example the Google Search Appliance or similar. If a login is required to view content that will be indexed, you should create a Confluence user specifically for the search crawler to use. Grant this user view rights to all content you wish to index, but deny that user all delete and administration rights. This ensures that an aggressive crawler will not be able to perform actions that could modify the site. There is also a forum thread on Google Mini integration.

External applications can also use the search function in the Confluence Remote API.

Related Information

- Setup Confluence To Index External Sites
- Setup Confluence To Index External Sites
- Setup Confluence To Index External Sites
- Integrate Confluence Search to Jira Search
- Setup External Search Tool To Index Confluence
- Setup External Search Tool To Index Confluence
- Setup Confluence To Index External Sites
- Setup External Search Tool To Index Confluence
- Setup Confluence To Index External Sites
- Setup Confluence To Index External Sites

Setup External Search Tool To Index Confluence
Setup Confluence To Index External Sites
Setup External Search Tool To Index Confluence
Setup External Search Tool To Index Confluence
Setup External Search Tool To Index Confluence
Setup Confluence To Index External Sites
Setup External Search Tool To Index Confluence

Security

- Adding SSL for Secure Logins and Page Security
- Anonymous Access to Remote API
- Confluence Cookies
- Enabling or Disabling Public Signup
- Hiding External Links From Search Engines
- Hiding the People Directory
- Managing External Referrers
    - Excluding external referrers
    - Hiding external referrers
    - Ignoring External Referrers
- User Email Visibility
- Spam Prevention via Captcha

Adding SSL for Secure Logins and Page Security

This document describes how to configure Confluence Stand-Alone or WAR file using Tomcat to use a HTTPS encrypted secure socket layer for user logins and page data.

> For other Application Servers, or if using Apache HTTPd Web Server, see Using Apache with mod_proxy for instructions on how to terminate an SSL connection at Apache Web Server.

Unencrypted confidential data within Confluence may be intercepted by an attacker. To secure user logins, you can enable access via HTTPS (HTTP over SSL), and require its use for pages where passwords are sent. In some cases where issue data is sensitive, all pages can be set to be accessed over HTTPS.

Enabling SSL access is different for each application server, but specifying which pages to require protection for is generic. This document is specific to Tomcat, the default application server shipped with Confluence.

On this page:

- Adding Secure User Logins
    - Creating A New SSL Certificate
    - Verify the Certificate is in the Correct Location
    - Specifying URL Patterns to be Redirected
- Troubleshooting

Adding Secure User Logins

Adding HTTPS requires a valid SSL certificate. If you have a Certificate prepared, skip to the 'Modify the `<INSTALL>/conf/server.xml` File' section.

Creating A New SSL Certificate

Creating a self-signed certificate

> The following commands are in reference to JDK 1.5. For commands/syntax relevant to JDK 1.6, please refer to this document.

On Windows, perform the following at the command prompt:

```
1."%JAVA_HOME%\bin\keytool" -genkey -alias tomcat -keyalg RSA
```

Or on other platforms, perform the following at the command prompt:

```
1.$JAVA_HOME/bin/keytool -genkey -alias tomcat -keyalg RSA
```

Some questions will be asked, including a password for the certificate (the default is 'changeit'). Please note down what you choose, as it will be used in the next step.



"IE7 on Vista Issue"
If your clients will access Confluence from Internet Explorer 7 on Vista, please ensure that you specify the -keyalg RSA flag. By default the SHA1 algorithm is used, which results in 'Internet Explorer cannot display the webpage' errors on IE7 on Vista. Apparently on JDK 1.6 you also need to specify the -sigalg MD5withRSA flag since -keyalg RSA will still result in SHA1 being used (see this blogpost for more information).

Modify the `<INSTALL>/conf/server.xml` File

In the confluence directory, open the *conf/server.xml* file and insert one of the following just after the closing `</Engine>` tag:

1. For users of Confluence 2.10 or later:
   Open `conf/server.xml`, uncomment the lines:

```
1.<Connector port="8443" maxHttpHeaderSize="8192"
2.                           maxThreads="150" minSpareThreads="25" maxSpareThreads="75"
3.                           enableLookups="false" disableUploadTimeout="true"
4.                           acceptCount="100" scheme="https" secure="true"
5.                           clientAuth="false" sslProtocol="TLS" SSLEnabled="true"
6.                           URIEncoding="UTF-8" keystorePass="<MY_CERTIFICATE_PASSWORD>"/>
```

   Or for users of Confluence 2.2 to 2.9.2:
   Open `conf/server.xml`, uncomment the lines:

```
1.<Connector port="8443" maxHttpHeaderSize="8192"
2.                                           maxThreads="150" minSpareThreads="25"
maxSpareThreads="75"
3.                                           enableLookups="false" disableUploadTimeout=
"true"
4.                                           acceptCount="100" scheme="https" secure=
"true"
5.                                           clientAuth="false" sslProtocol="TLS"
6.                                           URIEncoding="UTF-8"
keystorePass="<MY_CERTIFICATE_PASSWORD>" />
```


If your Confluence server is running off Apache Tomcat version 6.0.0 or later, you should ensure that the parameter-value pair `SSLEnabled="true"` has been added to the `Connector` tag above.

Establishing a CA-issued Certificate

In preparation for a production instance, an official CA-issued key pair is required. Find instructions in the Tomcat documentation.

Verify the Certificate is in the Correct Location

By default, Tomcat will look for the certificates in the file `C:\Documents and Settings\\#CURRENT_USER#\.keystore` on Windows or `~/.keystore` on Unix. If your Certificate is not in this location, you will need to update your `<INSTALL>/conf/server.xml` file as outlined below, so that Tomcat can find it.

1. For users of Confluence 2.2 or later:
   Open `conf/server.xml`, add the `keystoreFile="<MY_CERTIFICATE_LOCATION>"` parameter to the Connector tag as shown below:

```
1.<Connector port="8443" maxHttpHeaderSize="8192"
2.                                           maxThreads="150" minSpareThreads="25"
maxSpareThreads="75"
3.                                           enableLookups="false" disableUploadTimeout=
"true"
4.                                           acceptCount="100" scheme="https" secure=
"true"
5.                                           clientAuth="false" sslProtocol="TLS"
6.                                           URIEncoding="UTF-8"
keystorePass="<MY_CERTIFICATE_PASSWORD>" keystoreFile="<MY_CERTIFICATE_LOCATION>" />
```

Make sure to change your Server Base URL to https.

Specifying URL Patterns to be Redirected

Restart Tomcat and access your instance on `https://<MY_BASE_URL>:8443/`.

For more detailed information on setting up SSL with Tomcat (including additional configuration options), have a look at Tomcat 5.5 SSL or Tomcat 6 SSL.

Although HTTPS is now activated and available, the old HTTP URLs (http://localhost:8080) are still available. In most situations one wants these URLs to continue working, but for some to redirect to their HTTPS equivalent.

If you have changed the port that the SSL connector is running on from the preconfigured 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.

Site-Wide Protection

If using the RSS macro, you may need to configure this with a firewall rule, rather than on Tomcat. See RSS Feed Macro for further information.

If security is a concern, we recommend using SSL encryption site wide, for the reasons listed here: CONF-4116. To do this:

Edit the `confluence/WEB-INF/web.xml` file and add the following declaration to the end, **before** the `</web-app>` tag:

```
01.<security-constraint>
02.     <web-resource-collection>
03.          <web-resource-name>Restricted URLs</web-resource-name>
04.          <url-pattern>/</url-pattern>
05.     </web-resource-collection>
06.     <user-data-constraint>
07.          <transport-guarantee>CONFIDENTIAL</transport-guarantee>
08.     </user-data-constraint>
09.</security-constraint>
```

Once this change is made, restart Confluence and access http://localhost:8080. You should be redirected to https://localhost:8443/login.action.

Please note this will prevent IE from downloading attachments correctly if you are using Tomcat! In order to get this to work with IE you will need to modify server.xml or the relevant context file and add

```
1.<Valve className="org.apache.catalina.authenticator.NonLoginAuthenticator"
2.                                                    disableProxyCaching=
"true" securePagesWithPragma="false" />
```

within the <Context element

Protection for Login Action Only

If you'd like to protect login.action only:

Please note that redirecting login.action only will leave the rest of the site open to http connections, but will not redirect the user back to http once they're on https. To do that, you'll need to add a rewrite rule in Apache or IIS.

```
01.<security-constraint>
02.     <web-resource-collection>
03.          <web-resource-name>Login and Restricted Space URLs</web-resource-name>
04.          <url-pattern>/login.action</url-pattern>
05.     </web-resource-collection>
06.     <user-data-constraint>
07.          <transport-guarantee>CONFIDENTIAL</transport-guarantee>
08.     </user-data-constraint>
09.</security-constraint>
```

Note that the example above specifies a `url-pattern` for the login URL */login.action*. This means that whenever a user tries to access the unprotected version of the login page, they will be redirected automatically to the secured version of it.

Protection for Individual Spaces

If you want to protect individual spaces, there isn't a complete way of doing this at the moment. You can add a pattern like this:

```
<security-constraint>
    <web-resource-collection>
        <web-resource-name>Login and Restricted Space URLs</web-resource-name>
        <url-pattern>/login.action</url-pattern>
        <url-pattern>/display/SALARIES/*</url-pattern>
    </web-resource-collection>
    <user-data-constraint>
        <transport-guarantee>CONFIDENTIAL</transport-guarantee>
    </user-data-constraint>
</security-constraint>
```

Troubleshooting

Check the Confluence Knowledge Base articles at Troubleshooting SSL.

Anonymous Access to Remote API

Sites may wish to disable anonymous access to the remote API to make it harder for malicious users to write 'bots' that perform bulk changes to the site. If you wish to enable the Remote APIs but do not want anonymous users to access Confluence remotely, you can disable anonymous access from the **Administration Console**.

To disable anonymous access to Remote APIs,

1. Go to the Confluence '**Administration Console**'. To do this:
    - Open the '**Browse**' menu and select '**Confluence Admin**'. The 'Administration Console' view will open.
2. Select '**General Configuration**' in the left panel.
3. Click '**Edit**' at the bottom of the '**Options and Settings**' screen.
4. Select '**Off**' beside '**Anonymous Access to API**'.
5. '**Save**' your changes.

RELATED TOPICS

- Hiding the People Directory
- Anonymous Access to Remote API
- Spam Prevention via Captcha
- Hiding External Links From Search Engines
- User Email Visibility
- Ignoring External Referrers
- Managing External Referrers
- Enabling or Disabling Public Signup
- Hiding external referrers
- Excluding external referrers
- Adding SSL for Secure Logins and Page Security

Administrators Guide Home Confluence Documentation Home

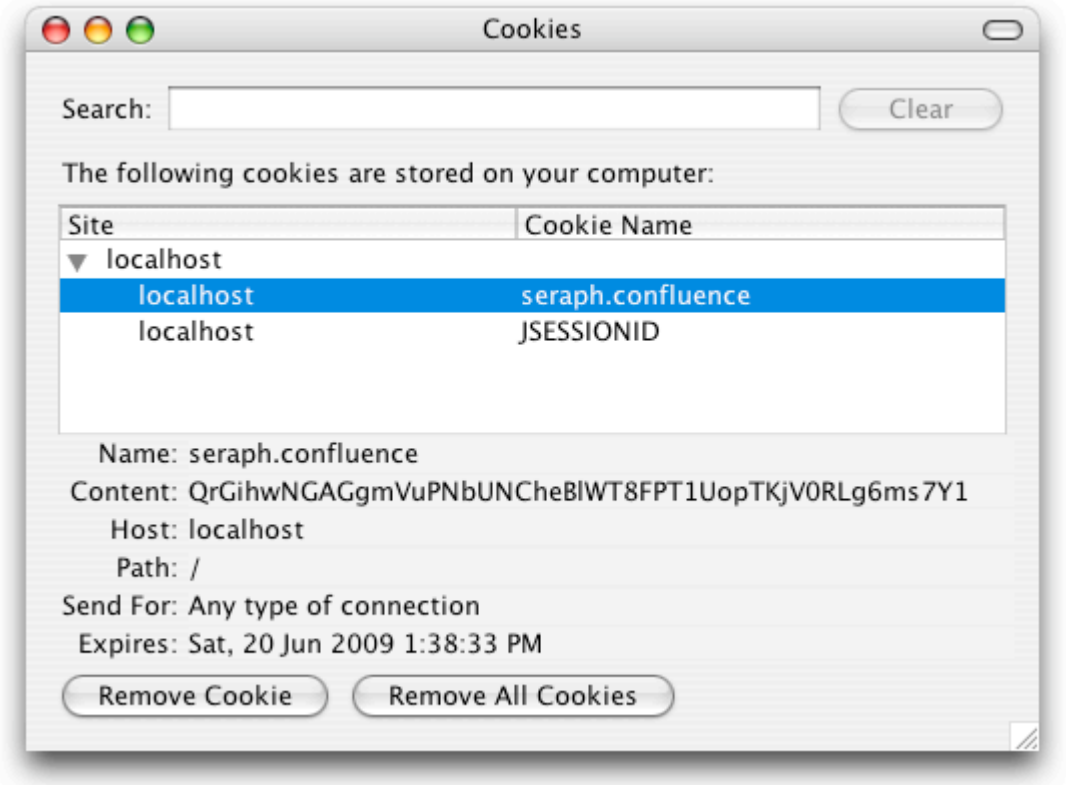
Confluence Cookies

Confluence uses Seraph, an open source framework, for HTTP cookie authentication.

Cookies

Confluence uses two cookies. The first, a JSESSIONID cookie, is created by the Application Server and used for session tracking purposes. The second, the 'Remember my login on this computer' cookie, is generated by Confluence when the user selects the check-box on the

log-in page.



Safe Information Transit

The cookie information is always encoded by the server before it is given to a client. A cookie that has been tampered with will be considered to be not valid.

Session and Cookie Logic

Essentially, the cookie contains encrypted username and the user's password. To be more precise confluence uses PBE (password based encryption) with MD5 and DES, where password(also known as 'private key') is configurable in the `seraph.xml` file. The user's password in the cookie is necessary to ensure that the cookie is no longer valid if the user changes their password. The username must be retrievable by the server to identify the user solely from the cookie, which is what the 'Remember my login on this computer' feature does.



The private key for confluence is stored at *confluence-install/confluence/WEB-INF/Classes/seraph.xml* under the *cookie.encoding* parameter. Please change this to something other than the default.

Is it Possible to Disable the 'Remember my login on this computer' Feature?

At the moment there is no available option for disabling "Remember My Login on this computer" feature via the Admin console. See the workaround here.

The user login *Auto Completion* functionality is a browser feature, and there is nothing Confluence can enable or disable.

RELATED TOPICS

 Hiding the People Directory

 Revoking Space Permissions

 Security Overview

 Confluence Security Advisory 2006-01-23

 Confluence Cookies

 Assigning Space Permissions

 Confluence Security

Edit in Word Link Macro

Space Permissions Overview

View File Macro

HTML Macro

How to Hide the Referrer

Administrators Guide Home Confluence Documentation Home

Enabling or Disabling Public Signup

Enabling '**Public Signup**' allows users to sign themselves up to the site.

If you want to restrict your site to a particular set of users, you may want to disable 'Public Signup'. In this instance, administrators can add new users from the **Administration Console**.

To enable or disable public signup,

1. Go to the Confluence '**Administration Console**'. To do this:
    - Open the '**Browse**' menu and select '**Confluence Admin**'. The 'Administration Console' view will open.
2. Select '**General Configuration**' in the left-hand panel.
3. This will display the '**General Configuration**' screen. Click '**Edit**'.
4. Beside '**Public Signup**', select '**On**' to enable Public Signup. Select '**Off**' to disable it.
5. Click '**Save**'.

RELATED TOPICS

Disabling the Built-In User Management
User Management
Security

Administrators Guide Home Confluence Documentation Home

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. Go to the Confluence '**Administration Console**'. To do this:
    - Open the '**Browse**' menu and select '**Confluence Admin**'. The 'Administration Console' view will open.
2. Select '**General Configuration**' in the left panel.
3. This will display the '**General Configuration**' screen. Click 'Edit'.
4. Select 'On' beside '**Hide External Links From Search Engines**'.
5. '**Save**' your changes.

Background to the nofollow attribute
As part of the effort to combat the spamming of wikis and blogs (Confluence being both), Google came up with some markup which instructs search engines not to follow links. By removing the main benefit of wiki-spamming it's hoped that the practice will stop being cost-effective and eventually die out.

RELATED TOPICS

Hiding the People Directory

Anonymous Access to Remote API

Spam Prevention via Captcha

Hiding External Links From Search Engines
User Email Visibility
Ignoring External Referrers
Managing External Referrers
Enabling or Disabling Public Signup
Hiding external referrers
Excluding external referrers
Adding SSL for Secure Logins and Page Security

Administrators Guide Home Confluence Documentation Home

Hiding the People Directory

The People Directory provides a list of all users in your Confluence system.

If you need to disable the People Directory set the following system properties on your application server command line:

- **To disable the People Directory for anonymous users,**

```
-Dconfluence.disable.peopledirectory.anonymous=true
```

- **To disable the People Directory entirely,**

```
-Dconfluence.disable.peopledirectory.all=true
```

To remove the link on the dashboard:

This only applies to Confluence 2.5.2 to 2.9.x. Confluence 2.10.x or later only needs to configure system properties using the above.
Edit the <confluence-install>/confluence/decorators/global.vmd:
Comment out line 37:

```
1.<!--                    <img src="$req.contextPath/images/icons/people_directory_32.gif"
align='absmiddle' height="32" width="32"> <b><a class="fontSizeDefault" href=
"$req.contextPath/peopledirectory.action"> $action.getText("people.directory.title"
)</a></b><span class="smalltext"> -    $action.getText("people.directory.description"
)</span><br> -->
```

RELATED TOPICS

Hiding the People Directory
Anonymous Access to Remote API
Spam Prevention via Captcha
Hiding External Links From Search Engines
User Email Visibility
Ignoring External Referrers
Managing External Referrers
Enabling or Disabling Public Signup
Hiding external referrers

Excluding external referrers

Adding SSL for Secure Logins and Page Security

Administrators Guide Home Confluence Documentation Home

Managing External Referrers

An external referrer is any site that links to your Confluence instance. Each time someone clicks on the external link, your Confluence site can record the click as a referral.

By default, external referrers for a page are listed under '**Hot Referrers**' on the '**Info**' screen of the page. (See *Screenshot 1* below.) Confluence shows a maximum of 10 referrers. If there are more than 10, confluence shows the 10 with the highest number of hits.

Note that you do *not* need to enable trackback in order to have external referrers enabled.

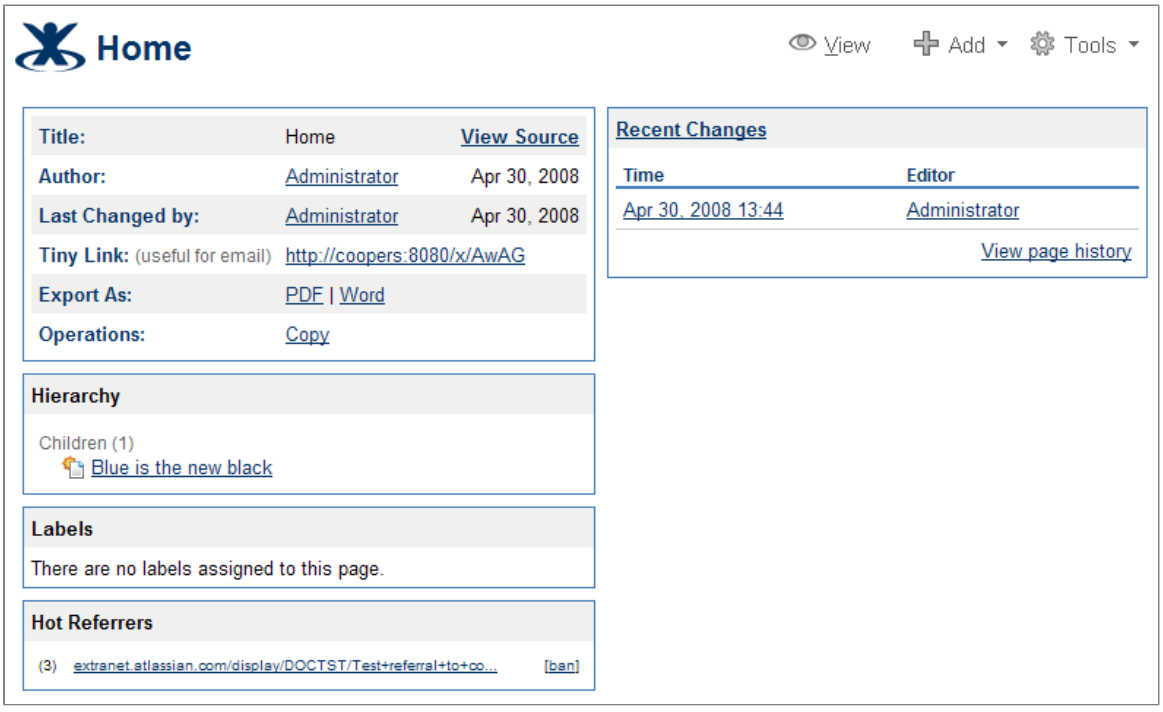
To manage your external referrers,

1. Go to the Confluence '**Administration Console**'. To do this:
    - Open the '**Browse**' menu and select '**Confluence Admin**'. The 'Administration Console' view will open.
2. Select the '**Manage Referrers**' option (See *Screenshot 2* below.).

The following actions will be available:

- Record or ignore all external referrers: By default, Confluence records the number of hits made to a page from the link on the external site. If you turn this option off, Confluence will not record the hits.
- Show or hide all external referrers: By default, Confluence lists the external referrers as '**Hot Referrers**' on the '**Info**' screen of a page, as shown below. If you turn this option off, external referrers will not be listed on the page.
- Specify which external referrers to exclude: You can decide which referrers you want to exclude from being displayed on your site.

Screenshot 1: Hot Referrers showing on a page's Info screen



Screenshot 2: Managing external referrers



RELATED TOPICS

- Hiding the People Directory (Confluence Docs 3.2)
- Anonymous Access to Remote API (Confluence Docs 3.2)
- Spam Prevention via Captcha (Confluence Docs 3.2)
- Hiding External Links From Search Engines (Confluence Docs 3.2)
- User Email Visibility (Confluence Docs 3.2)
- Ignoring External Referrers (Confluence Docs 3.2)
- Managing External Referrers (Confluence Docs 3.2)
- Enabling or Disabling Public Signup (Confluence Docs 3.2)
- Hiding external referrers (Confluence Docs 3.2)
- Excluding external referrers (Confluence Docs 3.2)
- Adding SSL for Secure Logins and Page Security (Confluence Docs 3.2)

Administrators Guide Home Confluence Documentation Home

Excluding external referrers

An external referrer is any site that links to your Confluence instance. Each time someone clicks on the external link, your Confluence site can record the click as a referral.

You can exclude external referrers to prevent them from being recorded or displayed anywhere on your site. Once you have specified your list of blocked URLs, any incoming links from URLs that match the list will no longer be recorded. Referrer URLs are blocked if they start with any of the URLs in the exclusion list. So http://evilspamsite.blogspot.com will also match http://evilspamsite.blogspot.com/nastypage.html

There are two instances where you may want to do this:

1. If you are running a Confluence installation that is open to public:
   In a site that is open to public, one unfortunate problem is that malicious sites can spam the display of a page's incoming links statistics. This is usually done to get the site's URL to appear in the sidebar. By adding these sites to the 'excluded referrers' list, you can prevent them from being listed on your site.
2. If Confluence is installed on a server with multiple domain names or IP addresses:
   Confluence will consider any URL originating from the domain name where Confluence is installed as an internal link. However, if Confluence is installed on a server with multiple domain names or IP addresses, you will need to add the other domain name prefixes to this list to let Confluence know that any links from these domains should not be considered external links.

You need to be a Confluence administrator and to know the URL of the site to add it to the excluded referrers list.

To add a URL to the excluded referrers list,

1. Go to the Confluence '**Administration Console**'. To do this:
    - Open the '**Browse**' menu and select '**Confluence Admin**'. The 'Administration Console' view will open.
2. Select '**Manage Referrers**' in the left-hand panel.
3. Add the URL to the '**Excluded External Referrer Prefixes**' section.
    - You must include 'http://' at the front of the URL.
    - You can add more than one URL by putting each URL on a new line.

Screenshot: Excluding external referrers



RELATED TOPICS

- Hiding the People Directory
- Anonymous Access to Remote API
- Spam Prevention via Captcha
- Hiding External Links From Search Engines
- User Email Visibility
- Ignoring External Referrers
- Managing External Referrers
- Enabling or Disabling Public Signup
- Hiding external referrers
- Excluding external referrers
- Adding SSL for Secure Logins and Page Security

Administrators Guide Home Confluence Documentation Home

Hiding external referrers

By default, Confluence lists the external referrers as '**Hot Referrers**' on the '**Info**' screen of a page. If you turn this option off, external referrers will not be listed on the page.

To hide external referrers,

1. Go to the Confluence '**Administration Console**'. To do this:
    - Open the '**Browse**' menu and select '**Confluence Admin**'. The 'Administration Console' view will open.
2. Select '**Manage Referrers**' in the left-hand panel.
3. Click '**Off**' beside '**Show Referrers in Page Info**'.

Screenshot: Managing external referrers



RELATED TOPICS

- Hiding the People Directory
- Anonymous Access to Remote API
- Spam Prevention via Captcha
- Hiding External Links From Search Engines
- User Email Visibility
- Ignoring External Referrers
- Managing External Referrers
- Enabling or Disabling Public Signup
- Hiding external referrers
- Excluding external referrers
- Adding SSL for Secure Logins and Page Security

Administrators Guide Home Confluence Documentation Home

Ignoring External Referrers

An external referrer is any site that links to your Confluence instance. Each time someone clicks on the external link, your Confluence site can record the click as a referral. By default, Confluence records the number of hits made to a page from any link on an external site. If you turn this option off, Confluence will not record the hits.

To ignore external referrers,

1. Go to the Confluence '**Administration Console**'. To do this:
    - Open the '**Browse**' menu and select '**Confluence Admin**'. The 'Administration Console' view will open.
2. Select '**Manage Referrers**' in the left-hand panel.
3. Click '**Off**' beside '**Record External Referrers**'.

Screenshot: Managing external referrers

Record External Referrers:	On | Off
Show Referrers in Page Info:	On | Off
Excluded External Referrer Prefixes:	Add Purge All

RELATED TOPICS

- Hiding the People Directory
- Anonymous Access to Remote API
- Spam Prevention via Captcha
- Hiding External Links From Search Engines
- User Email Visibility
- Ignoring External Referrers
- Managing External Referrers
- Enabling or Disabling Public Signup
- Hiding external referrers
- Excluding external referrers
- Adding SSL for Secure Logins and Page Security

Administrators Guide Home Confluence Documentation Home

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. Go to the Confluence '**Administration Console**'. To do this:
    - Open the '**Browse**' menu and select '**Confluence Admin**'. The 'Administration Console' view will open.
2. Select '**General Configuration**' in the left-hand panel.
3. This will display the '**General Configuration**' screen. Click '**Edit**'.
4. Select one of the options next to '**User email visibility**': '**public**', '**masked**', or '**only visible to site administrators**'.
5. '**Save**' your changes.

Screenshot : email visibility

User email visibility:	○ public ○ masked (i.e. user at example dot com) ◉ only visible to site administrators

RELATED TOPICS

Hiding the People Directory
Anonymous Access to Remote API
Spam Prevention via Captcha
Hiding External Links From Search Engines
User Email Visibility
Ignoring External Referrers
Managing External Referrers
Enabling or Disabling Public Signup
Hiding external referrers
Excluding external referrers
Adding SSL for Secure Logins and Page Security

Administrators Guide Home Confluence Documentation Home

Spam Prevention via Captcha

You need to be a Confluence administrator to enable Captcha.

If your Confluence site is open to the public you may find that automated spam is being added, in the form of comments or new pages.

You can configure Confluence to deter automated spam by asking users to prove that they are human before they are allowed to:

- Sign up for an account
- Add a comment
- Create a page
- Edit a page



Captcha is the technical term for a test that can distinguish a human being from an automated agent such as a web spider or robot.

When Captcha is switched on, users will need to recognise a distorted picture of a word, and must type the word into a text field. This is easy for humans to do, but very difficult for computers.

You can configure Confluence to enforce Captcha for certain types of users. You can exempt logged-in users (they will have completed a Captcha when they signed up), or members of particular groups.

By default, Captcha images will not be shown to logged-in users. Only anonymous users will have to perform the Captcha test when creating comments or editing pages.

To enable Captcha for Confluence,

1. Go to the Confluence '**Administration Console**'. To do this:
    - Open the '**Browse**' menu and select '**Confluence Admin**'. The 'Administration Console' view will open.
2. Select '**Spam Prevention**' from the 'Configuration' menu on the left.
3. Turn on Captcha by clicking the '**ON**' link.
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 the '**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. Click the '**Save**' button.

Working with Confluence Logs

On this page:

- Logs Location
- Background
- Finding the Log Configuration File
- Changing the Destination of the Log Files
- Changing the Logging Levels
- Using Some Specific Confluence Logging Options

Logs Location

This section describes Confluence's default logging behaviour, assuming that you have not changed the destination of the logs. So as to unify logging across different application servers, Confluence uses the **atlassian-confluence.log**, not the application server log, as it's primary log.

> Atlassian Support will almost always ask for the **atlassian-confluence.log** from the `confluence-home/logs` directory. **In Confluence 3.1 and later, you can access the logs from the Confluence Administration Console in the application, via the Support Utility**. Otherwise, the easiest way to find this location is to look for the "Confluence Home" setting from `Administration > System Information`. If you can't access `Administration > System Information`, check `<confluence-install>/confluence/WEB-INF/classes/confluence-init.properties` and look for the confluence.home setting, then find the logs in that directory.

For **Confluence 2.6.x and earlier**, the default behaviour is:

- For Confluence Standalone, log entries are written to `<confluence_install>/logs`. The main log file is called `atlassian-confluence.log`.
- For Confluence EAR/WAR, log entries are written to the application server logs, i.e. the default log files of the application container.

For **Confluence 2.7.x and later**, both Standalone and EAR/WAR distributions follow the same default behaviour:

- When you start Confluence, log entries will be sent to the application server logs until Confluence has completed its initial bootstrap. Any log entries written to the console will be repeated into the `<confluence-home>` log described below.
- Once the initial startup sequence is complete, all logging will be to `<confluence-home>/logs/atlassian-confluence.log`. For example: `c:/confluence/data/logs/atlassian-confluence.log`.

  Note that the default location is now the Confluence **home directory** instead of the **application server's log file**. The home directory is specified in <confluence-install>/confluence/WEB-INF/classes/confluence-init.properties.

Background

Confluence uses Apache's log4j logging service. This allows a developer or administrator to control the logging behavior and the log output file by editing a configuration file, without touching the application binary. There are six known log4j logging levels.

Finding the Log Configuration File

Confluence's logging behaviour is defined in the following properties file:
`<CONFLUENCE-INSTALL>/confluence/WEB-INF/classes/log4j.properties`

This file is a standard log4j configuration file, as described in the Apache log4j documentation.

Changing the Destination of the Log Files

Terminology: In log4j, an output destination is called an 'appender'.

To change the destination of the log files, you need to stop Confluence and then change the settings in the '**Logging Location and Appender**' section of the `log4j.properties` file. The location of this file is described above.

In the standard properties file supplied with **Confluence 2.7 and later**, you will find entries for two appenders:

- `com.atlassian.confluence.logging.ConfluenceHomeLogAppender` – This is a custom appender which controls the default logging destination described above. This appender allows the following settings:
  - MaxFileSize
  - MaxBackupIndex
- `org.apache.log4j.RollingFileAppender` – If you want to log to a different location, uncomment the `RollingFileAppender` line and change the destination file in the line below it. Comment out the previous lines referring to the `ConfluenceHomeLogAppender`.

Confluence ships with the full suite of appenders offered by log4j. Read more about appenders in the log4j documentation.

Changing the Logging Levels

See Configuring Logging for instructions on how to change the logging configuration of Confluence.

Using Some Specific Confluence Logging Options

This section contains some pointers to specific log configurations you may need.

Log the Details of SQL Requests made to the Database

You may want to increase Confluence's logging so that it records individual SQL requests sent to the database. This is useful for troubleshooting specific problems.

You can enable detailed SQL logging in two ways:

- At runtime – see instructions above.
- Via the logging properties file – see the detailed instructions.

Log the Details of Users Viewing/Accessing each Confluence Page

You can configure the log to show which users are accessing which pages in Confluence. This can only be done via the logging properties file – see the detailed instructions.

Where are my Thread Dumps?

Thread dumps are logged to the application server log file.

RELATED TOPICS

- Important Directories and Files
- Enabling detailed SQL logging
- Enabling user access logging
- Generating a Thread Dump
- Enabling Page Request Profiling

log4j Logging Levels

Logging Levels

- **DEBUG** - designates fine-grained informational events that are most useful to debug an application (*what is going on*)
- **INFO** - announcements about the normal operation of the system - scheduled jobs running, services starting and stopping, user-triggered processes and actions
- **WARN** - any condition that, while not an error in itself, may indicate that the system is running sub-optimally
- **ERROR** - a condition that indicates something has gone wrong with the system
- **FATAL** - a condition that indicates something has gone wrong so badly that the system can not recover
- **TRACE** - n/a within confluence



There are two ways to modify the logging levels, as described in Working with Confluence Logs.

1. Modifying the runtime log levels via the **Administration Console**.
2. Manually modifying the `<Confluence-Install>\confluence\WEB-INF\classes\log4j.properties` file.

Default Log Level

The standard Confluence log level **WARN** is a way for Confluence to communicate with the server administrator. Logging at WARN level and higher should be reserved for situations that require some kind of attention from the server administrator, and for which corrective action is possible.

Reference : log4j manual

Configuring OAuth

On this page:

- What is OAuth?

- Accessing and Using Confluence's OAuth Administration Page

What is OAuth?

OAuth is a protocol that allows one application to share a finite set of its private resources and data (through gadgets, for example) with another application. These applications could be a Confluence or JIRA site, or a website such as iGoogle. However, all applications involved must be OAuth-compliant.

Using OAuth, you can access data within a Confluence installation externally, via a Confluence gadget published on a JIRA site's dashboard, another Confluence site's page, or a website like iGoogle. While some data in Confluence may be accessible anonymously on the external application, other data may be restricted to a specific user account within the Confluence installation. OAuth provides the facility to access this restricted data.

The key security advantage of OAuth is that Confluence's user-restricted resources can be shared without Confluence having to hand out user authentication details. Instead, access to these private resources is handled via an '**access token**'. Access tokens define what Confluence resources (which are typically based on access privileges) can be accessed by another application and the duration of this access. However, access tokens are dissociated from a user's authentication details, since authentication to gain access to these resources is handled separately.

In OAuth terminology, an application that shares its resources is known as a **service provider** and an application that accesses a service provider's resources is known as a **consumer**.

For more information about OAuth, please refer to the OAuth protocol workflow section of our Gadgets and Dashboards documentation. It is important to understand this workflow first before establishing OAuth relationships between your Confluence installation and other external web applications (either Atlassian or non-Atlassian ones).

Important information about establishing OAuth relationships for gadgets

If you wish to use a gadget served by any Atlassian application *and* require this gadget to access data which is restricted to a user account on that application, then either a Trusted Application or OAuth relationship between the service provider and consumer application *must be established first*.

Alternatively, if the gadget is served by an Atlassian application which supports Atlassian's Trusted Applications feature (for example, JIRA, Confluence or Bamboo), you can establish a Trusted Applications relationship instead of an OAuth one. Bear in mind that in Trusted Application relationships, you can only access data restricted to a user account on the service provider if:

1. The usernames of user accounts on the service provider and consumer applications match.
2. The user has already logged in to the consumer application.

Unlike Trusted Application relationships, OAuth relationships provide the ability to access restricted data on the service provider when an individual's usernames on the service provider and consumer applications are different. This is because authentication is part of the OAuth protocol workflow.

Not all external gadgets used in Confluence require the establishment of an OAuth relationship. If the gadget does not need to access restricted resources on the service provider, then there should be no need to establish an OAuth relationship.

The instructions in this section provide information on how to establish an OAuth relationship between your Confluence site and another web application's site. This could even apply to situations where Confluence is either the consumer or service provider in the OAuth relationship.

Accessing and Using Confluence's OAuth Administration Page

Confluence's OAuth Administration section, which handles the establishment of OAuth relationships between consumer and service provider web applications, is found in the Administration Console area of Confluence.

To access Confluence's OAuth Administration page,

1. Go to the Confluence '**Administration Console**'. To do this:
    - Open the '**Browse**' menu and select '**Confluence Admin**'. The 'Administration Console' view will open.
2. Click '**OAuth**' from the 'Administration' section in the left navigation panel to open the 'OAuth Administration' page. On this page:
    - Click the '**Consumers**' tab to configure consumer applications that will be accessing your Confluence installation's resources (such as Confluence's gadgets). Refer to Configuring OAuth Consumers for more information.
    - Click the '**Consumer Info**' tab to view or edit your Confluence installation's Consumer information. Refer to Configuring OAuth Consumer Information for more information.
    - Click the '**Service Providers**' tab to configure service providers whose resources your Confluence installation will be using. Refer to Configuring OAuth Service Providers for more information.

      In the documentation links above, 'your Atlassian application' refers to your Confluence installation.

User Management

- Confluence User Management
  - Searching For and Managing Users
  - Adding a Group
  - Adding a New User
  - Adding or Removing Users in Groups
  - Changing Usernames
  - Editing User Details
  - Global Groups Overview
  - Global Permissions Overview
  - Migrating to new User Management
  - Removing a Group
  - Removing a User
  - Setting up Anonymous Access
  - Viewing members of a group
  - How to Improve User Search Performance — If your Confluence instance contains thousands of user accounts and you are experiencing performance issues when searching for users, the following migration guide is for you.
  - Restoring Passwords To Recover Admin User Rights
- Integrating with Crowd
- JIRA User Management
  - Delegate user management to use JIRA logins
    - Revert from JIRA to internal user management
- LDAP User Management
  - Add LDAP Integration
    - Automatically Add LDAP users to the confluence-users Group
    - Customising atlassian-user.xml
    - Migrate to LDAP User Management From OsUser
  - Add LDAP Integration For User Authentication Only
  - atlassian-user.xml reference — || XML tag || Default value || Description |
  - Changes in osuser.xml from 1.0.3a to 1.1.x
  - Configuring multiple LDAP repositories
  - Connect to LDAP, JIRA or Other Services Via SSL
  - Disabling the Built-In User Management
  - Legacy User Management Documentation
    - LDAP Authentication with OSUser
  - Troubleshooting LDAP User Management
  - Troubleshooting the "Not Permitted" Screen under LDAP Integration
- Migrating users from Confluence to JIRA — There is currently no way to delegate user management from JIRA to Confluence. So, if you are in a situation where your users are defined in Confluence and would like to take advantage of Confluence's ability to use JIRA user management, you will need to transfer all of your existing Confluence users into JIRA. You can do this manually, or if you have a large number of users, you can use the attached XML-RPC script.
- Requesting External User Management Support
  - Paddle
- Understanding User Management in Confluence
- User Management Frequently Asked Questions

Confluence User Management

- Searching For and Managing Users
- Adding a Group
- Adding a New User
- Adding or Removing Users in Groups
- Changing Usernames
- Editing User Details
- Global Groups Overview
- Global Permissions Overview
- Migrating to new User Management
- Removing a Group
- Removing a User
- Setting up Anonymous Access
- Viewing members of a group
- How to Improve User Search Performance
- Restoring Passwords To Recover Admin User Rights

Searching For and Managing Users

If you are a Confluence Administrator, you can add users, assign them to groups and edit their user details.

On this page:

- Accessing the User Management Screen
- Listing All Users
- Using the Simple User Search
- Using the Advanced User Search

Accessing the User Management Screen

To search for and manage users,

1. Go to the user management screen for the user concerned. There are two ways to do this:
   - Either,
     - Go to the user's Profile and click the '**Administer User**' link on the user's profile screen. (This link is available in Confluence 2.8.2 and later.)
   - Or,

     Go to the Confluence '**Administration Console**'. To do this:

     - Open the '**Browse**' menu and select '**Confluence Admin**'. The 'Administration Console' view will open.
     - Select '**Manage Users**' in the left-hand panel.
     - The '**Manage Users**' screen appears, as shown below. You can now choose to list all users or you can search for a specific user.

Screenshot: Manage users

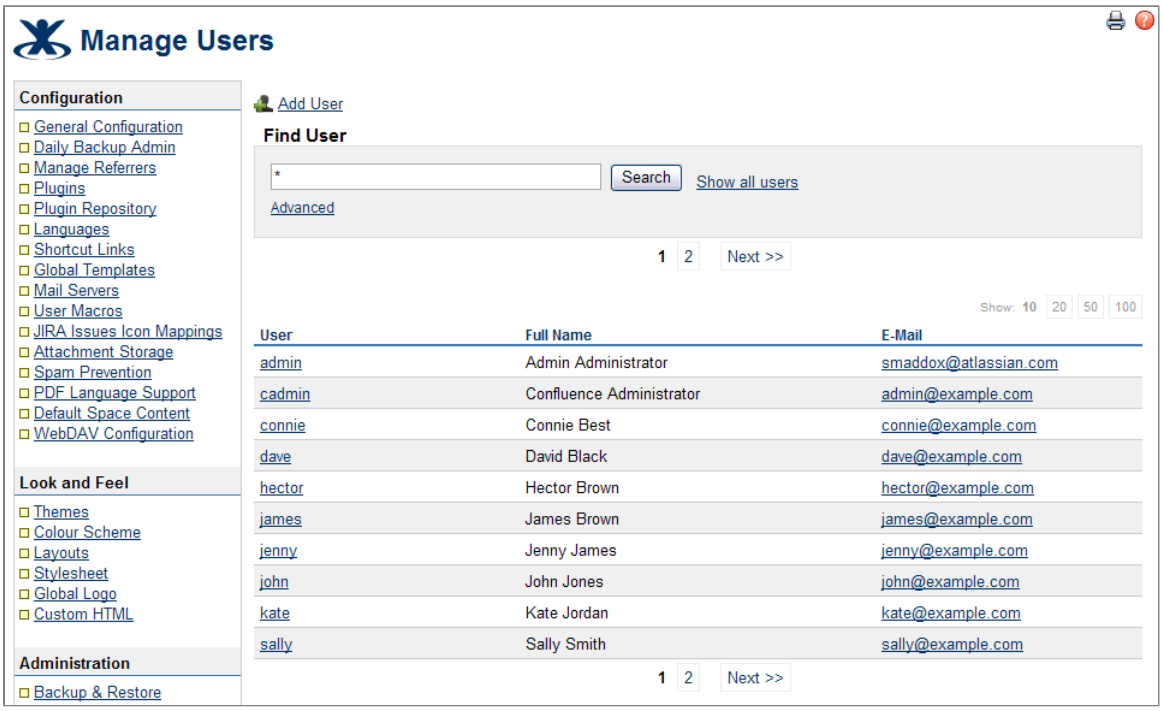


Listing All Users

To list all users,

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

Screenshot: List all users



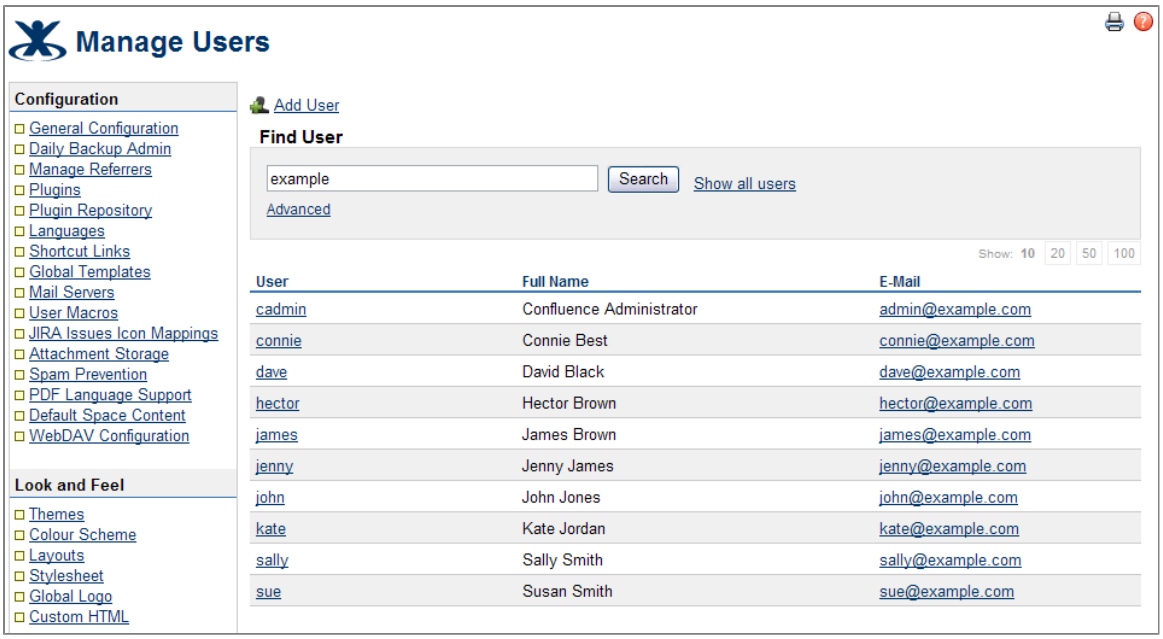
Using the Simple User Search

> **Crowd and the User Search**
> If you are using Atlassian's Crowd for user management, you will need Crowd 1.5.1 or later to use the 'Simple' option in the user search. If your version of Crowd does not support the simple user search, you will see only the 'Advanced' search form.

To search for a specific user via the simple user search,

1. Open the '**Manage Users**' screen as described above.
2. If the '**Simple**' link is showing, click it. (If you see the 'Advanced' link and no 'Simple' link, then you're fine. The simple search is already active.)
3. The simple user search screen will appear, as shown below.
4. Type some information about the user into the 'Search' textbox. You can type all or part of their username, full name or email address.
5. Click the '**Search**' button.
6. Confluence will display a list of matching users. Click the link on a username to see and edit the details for that user.

Screenshot: Simple user search



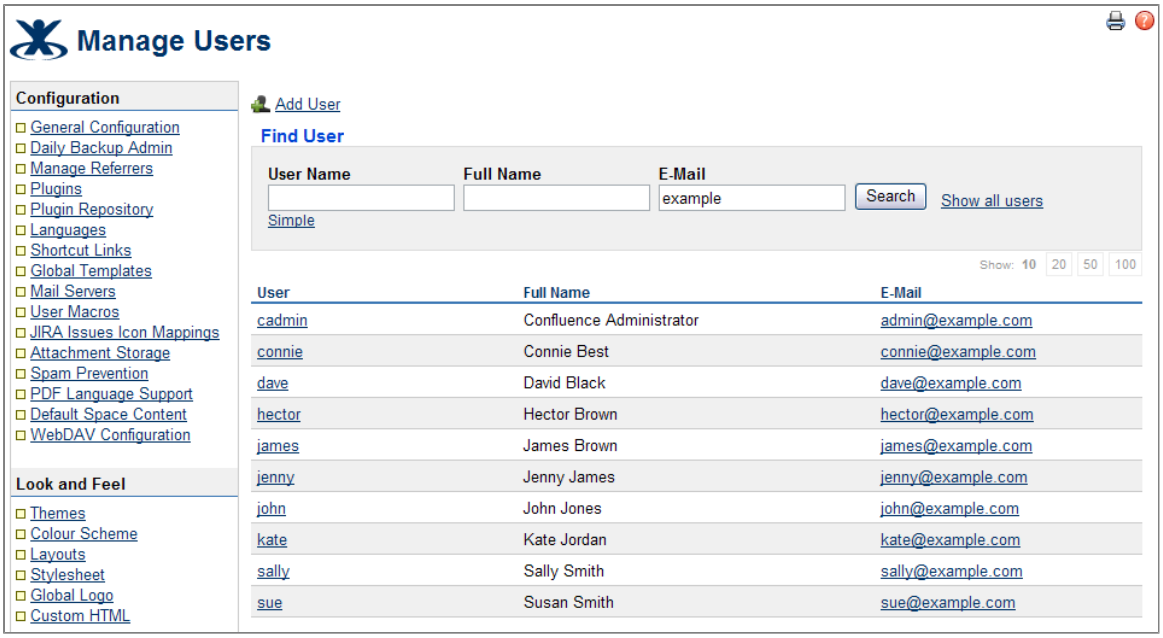
Using the Advanced User Search

The advanced user search allows you to specify the field in which your search term appears, i.e. username, full name or email address. You may find this useful if you need to limit the number of users appearing in the search results.

To search via the advanced user search,

1. Open the '**Manage Users**' screen as described above.
2. If the '**Advanced**' link is showing, click it. (If you see the 'Simple' link and no 'Advanced' link, then you're fine. The advanced search is already active.)
3. The advanced user search screen will appear, as shown below.
4. Complete one or more of the following fields:
    - **User Name** — Enter all or part of the person's username i.e. their login id, e.g. 'joe', or 'bloggs'.
    - **Full Name** — Enter all or part of the person's name, e.g. 'joe bloggs', or 'bloggs', or 'joe'.
    - **E-Mail** — Enter all or part of the person's email address, e.g. 'acme'
5. Click the '**Search**' button.
6. Confluence will display a list of matching users. Click the link on a username to see and edit the details for that user.

Screenshot: Advanced user search



RELATED TOPICS

- Searching For and Managing Users
- Security Overview
- Adding or Removing Users in Groups
- Troubleshooting LDAP User Management
- Changing Usernames
- Setting up Anonymous Access
- Enabling or Disabling Public Signup
- Removing a Group
- Global Groups Overview
- Adding a Group
- Adding a New User
- Editing User Details
- Removing a User
- Global Permissions Overview
- Disabling the Built-In User Management

Showing first 15 of 16 results

[!Administration Guide Attachments directory^adminhome.gif!] Confluence Documentation Home

Adding a Group

To add a new group,

1. Go to the Confluence '**Administration Console**'. To do this:
    - Open the '**Browse**' menu and select '**Confluence Admin**'. The 'Administration Console' view will open.
2. Select '**Manage Groups**' in the left panel.
3. Enter a name for your group in the '**Add Group**' input field and click '**Save**'.

You are now ready to start adding users to the group.

RELATED TOPICS

- Searching For and Managing Users
- Adding or Removing Users in Groups
- Removing a Group
- Global Groups Overview
- Viewing members of a group

[!Administration Guide Attachments directory^adminhome.gif!] Confluence Documentation Home

Adding a New User

There are two ways a new user can be added to Confluence:

Public Signup: Enabling public signup from the **Administration Console** allows users to sign themselves up to the site.

By Confluence Administrators : If you want to restrict your site to a select group of users, you may want to disable 'Public Signup'. In this instance, administrators with Confluence Administrator or System Administrator permissions can add new users from the **Administration Console**.

To add a new user to Confluence from the Administration Console,

1. Go to the Confluence '**Administration Console**'. To do this:
    - Open the '**Browse**' menu and select '**Confluence Admin**'. The 'Administration Console' view will open.
2. Select '**Manage Users**' in the left-hand panel.
3. Click the link '**Add new user**' at the top of the page.
4. In the form displayed, enter the user's details: username, password, name and email address.
5. Click '**Create**' to add the user.

RELATED TOPICS

- Searching For and Managing Users
- Security Overview
- Adding or Removing Users in Groups
- Troubleshooting LDAP User Management
- Changing Usernames
- Setting up Anonymous Access
- Enabling or Disabling Public Signup
- Removing a Group
- Global Groups Overview
- Adding a Group
- Adding a New User
- Editing User Details
- Removing a User
- Global Permissions Overview
- Disabling the Built-In User Management

Showing first 15 of 16 results

Administrators Guide Home Confluence Documentation Home

Adding or Removing Users in Groups

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

This page tells you how to add a user to a group or remove a user from a group. For an overview of users and groups, please refer to Users and Groups and Confluence User Management.

You can edit group membership in two places:

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

Both methods are described below.

On this page:

- Adding and Removing Members via the Group Management Screen
- Editing Group Membership from the User Management Screen
- A Note about External User Directories

Adding and Removing Members via the Group Management Screen

This is the recommended method, available in **Confluence 2.10** and later. It allows you to manage the group membership for a number of users at the same time.

To add members to a group,

1. Go to the Confluence '**Administration Console**'. To do this:
    - Open the '**Browse**' menu and select '**Confluence Admin**'. The 'Administration Console' view will open.
2. Select '**Manage Groups**' in the left-hand panel.
3. The '**Manage Groups**' screen appears, showing a list of groups. Select the group to which you want to add users.
4. The '**Group Members**' screen appears, showing the users who belong to the selected group. (See screenshot below.) Click the '**Add Members**' link.
5. The '**Add Members**' screen appears, as shown below. Type in the usernames of the people you want to add to the group. You can also search for and select users by clicking the icon, as described in Searching for Users.
6. When you have added the required username(s), click the '**Add**' button to add the member(s) to the group.

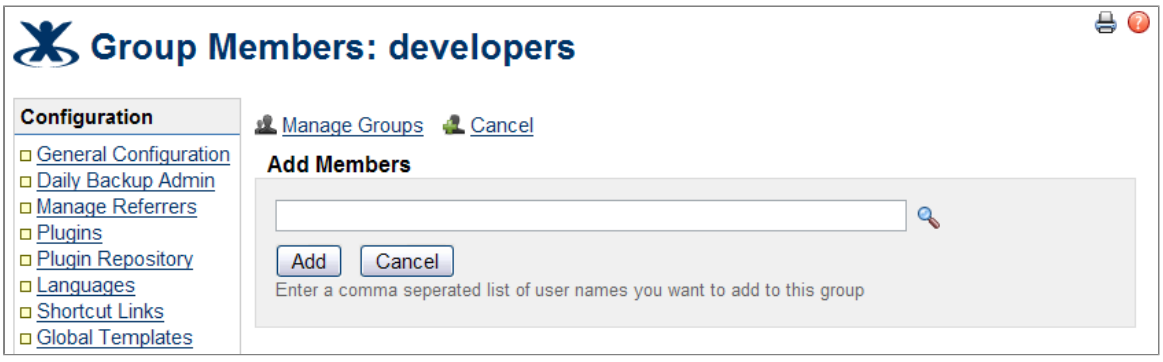
To remove members from a group,

1. Go to the Confluence '**Administration Console**'. To do this:
    - Open the '**Browse**' menu and select '**Confluence Admin**'. The 'Administration Console' view will open.
2. Select '**Manage Groups**' in the left-hand panel.
3. The '**Manage Groups**' screen appears, showing a list of groups. Select the group from which you want to remove the user.
4. The '**Group Members**' screen appears, showing the users who belong to the selected group. (See screenshot below.) Click the 'Remove user from group' icon next to the user whose group membership you want to remove.

Screenshot: Group Members



Screenshot: Add Members



Editing Group Membership from the User Management Screen

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

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

1. Go to the user management screen for the user concerned. There are two ways to do this:
    - Either,
        - Go to the user's Profile and click the '**Administer User**' link on the user's profile screen. (This link is available in Confluence 2.8.2 and later.)
    - Or,

      Go to the Confluence '**Administration Console**'. To do this:

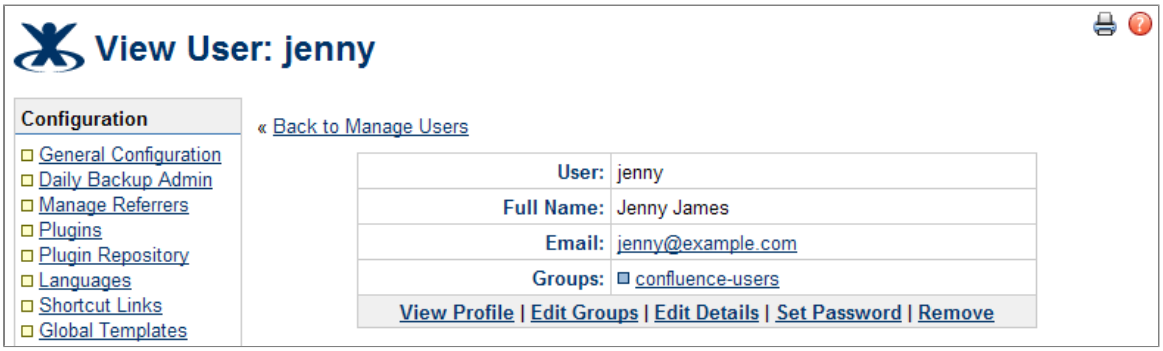
        - Open the '**Browse**' menu and select '**Confluence Admin**'. The 'Administration Console' view will open.
        - Select '**Manage Users**' in the left-hand panel.
        - The '**Manage Users**' screen appears, as shown below. You can now choose to '**Show all users**' or you can search for a specific user by entering all or part of the person's username, full name or email address. (For more details about the user search, see Searching For and Managing Users.)
        - Click the link on the username you want to edit.
2. Now you should be able to see the user's current details, with links allowing you to edit the user's details and groups. See the screenshot showing a user's details below.

3. Click '**Edit Groups**'. This will display two lists of groups, as shown in the screenshot below. Update the user's group membership as follows:
    - '**Not a member of groups**' — This box shows all groups to which the user does not belong. To add the user to a group, select a group and click '**Join**'. Hold the Ctrl key down and click to select more than one group.
    - '**Member of groups**' — This box shows all groups to which the user belongs. Select a group and click '**Leave**' to remove the user from the group.

Screenshot: Manage users



Screenshot: User details



Screenshot: Editing a user's groups



A Note about External User Directories

If you are using external user management, you cannot use the Confluence administration screens to add or remove users in external groups, as Confluence access to LDAP and JIRA groups is read only.

If you have Confluence integrated with Crowd and external user management turned **off** in Confluence, then your group membership changes in Confluence will be passed through to the external directory (e.g. LDAP) managed by Crowd (assuming that Crowd has read-write access to the directory).

RELATED TOPICS

- Searching For and Managing Users
- Security Overview
- Adding or Removing Users in Groups
- Troubleshooting LDAP User Management
- Changing Usernames
- Setting up Anonymous Access
- Enabling or Disabling Public Signup
- Removing a Group
- Global Groups Overview
- Adding a Group
- Adding a New User
- Editing User Details
- Removing a User
- Global Permissions Overview
- Disabling the Built-In User Management
- Viewing members of a group

Administrators Guide Home Confluence Documentation Home

Changing Usernames

A **username** is the name used to log into Confluence, eg. `jsmith`.

Currently, there is no straightforward method for changing a username and its associated content, to that of another user. The only practicable method currently available is to execute direct SQL queries on your database. There is a feature request to facilitate this process via a web interface and you can vote for it to improve its chances of being implemented. Be aware, however, that no matter what method you use to change usernames in Confluence, there is **no support** provided for this process. The instructions below provide suggested guidelines on how to change a username via SQL queries, although this may vary depending on your database.

Instructions For Changing Usernames

The following SQL commands are only tested for MYSQL and POSTGRES Databases. If you have any other database please contact your DBA to determine the equivalent queries.

Usernames can only be changed through direct update to the Confluence database.

1. If you have a database administrator, request that they approve the database-related steps described below
2. If you are using JIRA user management, Revert from JIRA To Internal User Management
3. Backup Confluence
4. Creating a `usermigration` table:

```
1.create table usermigration
2.(
3.oldusername varchar,
4.newusername varchar
5.)
```

5. Usernames that will be changed must be placed in the `usermigration` table with their current and planned usernames:

```
1.insert into usermigration (oldusername, newusername)
2.values ('oldusername', 'newusername');
```

6. If your DB administration tool does not support multiple SQL queries, these must be entered individually.

```
001.update attachments
002.set creator = newusername from usermigration u
003.where creator = u.oldusername;
004.
005.update attachments
006.set lastmodifier = newusername from usermigration u
007.where lastmodifier = u.oldusername;
008.
009.update content
010.set lastmodifier = newusername from usermigration u
011.where lastmodifier = u.oldusername;
012.
013.update content
014.set creator = newusername from usermigration u
015.where creator = u.oldusername;
016.
017.update content
018.set username = newusername from usermigration u
019.where username = u.oldusername;
020.
021.update extrnlnks
022.set creator = newusername from usermigration u
023.where creator = u.oldusername;
024.
025.update extrnlnks
026.set lastmodifier = newusername from usermigration u
027.where lastmodifier = u.oldusername;
028.
029.update label
030.set owner = newusername from usermigration u
031.where owner = u.oldusername;
032.
033.update content_label
034.set owner = newusername from usermigration u
035.where owner = u.oldusername;
036.
037.update links
038.set lastmodifier = newusername from usermigration u
039.where lastmodifier = u.oldusername;
040.
041.update links
```

```
042.set creator = newusername from usermigration u
043.where creator = u.oldusername;
044.
045.update notifications
046.set lastmodifier = newusername from usermigration u
047.where lastmodifier = u.oldusername;
048.
049.update notifications
050.set creator = newusername from usermigration u
051.where creator = u.oldusername;
052.
053.update pagetemplates
054.set lastmodifier = newusername from usermigration u
055.where lastmodifier = u.oldusername;
056.
057.update pagetemplates
058.set creator = newusername from usermigration u
059.where creator = u.oldusername;
060.
061.update spaces
062.set creator = newusername from usermigration u
063.where creator = u.oldusername;
064.
065.update spaces
066.set lastmodifier = newusername from usermigration u
067.where lastmodifier = u.oldusername;
068.
069.update spacepermissions
070.set permusername = newusername from usermigration u
071.where permusername = u.oldusername;
072.
073.update spacepermissions
074.set creator = newusername from usermigration u
075.where creator = u.oldusername;
076.
077.update spacepermissions
078.set lastmodifier = newusername from usermigration u
079.where lastmodifier = u.oldusername;
080.
081.update contentlock
082.set creator = newusername from usermigration u
083.where creator = u.oldusername;
084.
085.update contentlock
086.set lastmodifier = newusername from usermigration u
087.where lastmodifier = u.oldusername;
088.
089.update os_user
090.set username = newusername from usermigration u
091.where username = u.oldusername;
092.
093.update trackbacklinks
094.set creator = newusername from usermigration u
095.where creator = u.oldusername;
096.
097.update trackbacklinks
098.set lastmodifier = newusername from usermigration u
099.where lastmodifier = u.oldusername;
```

7. If using Confluence 2.1 or newer, run the following command:

```
1.update users
2.set name = newusername from usermigration u
3.where name = u.oldusername;
```

8. Reassign personal spaces and content associated with the old username to the new username.

```
1.update content_label
2.set owner = 'newusername'
3.where owner = 'oldusername';
```

For the two queries below, the tilda (~) is required as it is prepended to the space key of all personal spaces.

```
1.update spaces
2.set spacekey = '~newusername'
3.where spacekey = '~oldusername';
4.
5.update bandana
6.set bandanacontext = '~newusername'
7.where bandanacontext = '~oldusername';
```

9. Each username is associated with a full name. For example, username 'jsmith' may have a full name of 'John M Smith'. If this fullname needs to be changed, modify the fullname in the `users` or `os_user` table.

All old usernames in Confluence should now be replaced with the new usernames from the `usermigration` table.

RELATED TOPICS

- Searching For and Managing Users
- Security Overview
- Adding or Removing Users in Groups
- Troubleshooting LDAP User Management
- Changing Usernames
- Setting up Anonymous Access
- Enabling or Disabling Public Signup
- Removing a Group
- Global Groups Overview
- Adding a Group
- Adding a New User
- Editing User Details
- Removing a User
- Global Permissions Overview
- Disabling the Built-In User Management

Showing first 15 of 16 results

Administrators Guide Home Confluence Documentation Home

Editing User Details

To update a user's details,

1. First, go to the user management screen for the user concerned. There are two ways to do this:
   - Either,
     - Go to the user's Profile and click the '**Administer User**' link on the user's profile screen. (This link is available in Confluence 2.8.2 and later.)
   - Or,

     Go to the Confluence '**Administration Console**'. To do this:

     - Open the '**Browse**' menu and select '**Confluence Admin**'. The 'Administration Console' view will open.
     - Select the link '**Manage Users**' in the left-hand panel.
     - Locate the user by doing a search on the username or the groups to which they belong.
     - Click the user link.
2. Now you should be able to see the user's current details and links allowing you to edit them.
   - **View Profile** — View the user's profile.
   - **Edit Groups** — Add or remove this user from a group.
   - **Edit Details** — Edit the user's name and email address. Changing a user's username is not supported through the application, see Changing Usernames for other solutions.
   - **Set Password** — Edit the user's password details.
   - **Deactivate** — Deactivation is an alternative to removing a user. It is not possible to remove a user if the user is responsible for content on the site. The reason is that Confluence will need the user information to maintain a history of the pages they have updated. In this case, you can deactivate the user so that they can no longer log in to Confluence.
   - **Remove** — You can remove a user permanently if the user has not added or edited any content on the site.

Deactivating users (unavailable) : This functionality was removed in Confluence 2.1.x. See Removing a User for more information.

Screenshot : User Details



RELATED TOPICS

- Searching For and Managing Users
- Spam Prevention via Captcha
- Adding or Removing Users in Groups
- Setting up Anonymous Access
- Adding a New User
- Editing User Details
- Removing a User
- Global Permissions Overview

Global Groups Overview

There are two special default groups in Confluence:

1. **confluence-administrators**: This is a group of 'super-users' who can access the '**Administration Console**' and perform site-wide administration. Members of this group can also see all spaces in the Confluence instance.
   Any user who is a member of this group has site-wide administration powers, regardless of any other setting. The settings on the Global Permissions screen do **not** affect the powers allowed to members of this group.

   There is an outstanding request to remove the 'confluence-administrators' group from a future version of Confluence (see CONF-4616).

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

Other user groups : A Confluence administrator can also group users together into user groups for more convenient administration. Once created, groups become available at the space and page levels to allow for flexible access control. A user in one of these groups will automatically be granted all permissions granted to the group.

Anonymous users : Confluence treats all users who do not log in when they access Confluence as being 'anonymous'. You can grant anonymous '**Use Confluence**' permission via the Global Permissions screen. This will allow non-registered users to access pages and spaces in Confluence. A space administrator can then further control anonymous access per space via the space permissions.

RELATED TOPICS

- Searching For and Managing Users
- Security Overview
- Adding or Removing Users in Groups
- Troubleshooting LDAP User Management
- Changing Usernames
- Setting up Anonymous Access
- Enabling or Disabling Public Signup
- Removing a Group

Global Groups Overview

Adding a Group

Showing first 10 of 16 results

Administrators Guide Home Confluence Documentation Home

Global Permissions Overview

Permissions determine the actions which a user is allowed to perform within Confluence. Global permissions are one of the levels of permission provided by Confluence.

In order to assign these permissions, you must already have the global '**Confluence Administrator**' or '**System Administrator**' permission (described below). You can then assign global permissions to groups, individual users and anonymous users. Further permissions are granted from the space administration screens.

On this page:

- Overview of the Global Permissions
- Comparing the System Administrator with the Confluence Administrator Permission
- Comparing the Administrator Permissions with the confluence-administrators Group
- Updating Global Permissions

Overview of the Global Permissions

Global permissions control access across the whole Confluence site. Here is a list:

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

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

Comparing the System Administrator with the Confluence Administrator Permission

New with **Confluence 2.7 and later** comes the ability to have two levels of administrator in Confluence:

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

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

The following functions are excluded from the '**Confluence Administrator**' permission:

Administration Screen	Excluded Function
General Configuration	The following functionality is disallowed: • Server Base URL • Remote API plugin • External user management • Public Signup
Daily Backup Admin	This function is disallowed entirely.
Plugins	This function is disallowed entirely.
Plugin Repository	This function is disallowed entirely.
Mail Servers	This function is disallowed entirely.
User Macros	This function is disallowed entirely.
Attachment Storage	This function is disallowed entirely.
Layouts	This function is disallowed entirely.
Custom HTML	This function is disallowed entirely.
Backup & Restore	This function is disallowed entirely.
SnipSnap Import	This function is disallowed entirely.
Logging and Profiling	This function is disallowed entirely.
Cluster Configuration	This function is disallowed entirely.

Comparing the Administrator Permissions with the confluence-administrators Group

The '**confluence-administrators**' group defines a set of 'super-users' who can access the Administration Console and perform site-wide administration. Members of this group can also see all pages and spaces in the Confluence instance except pages for which they are excluded by page restrictions (restrictions can be removed by members of the confluence-administrators group in the Space Admin screen if need be). The settings on the 'Global Permissions' screen do not affect the powers allowed to members of this group.

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

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

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

There is an outstanding request to remove the 'confluence-administrators' group from a future version of Confluence (see CONF-4616).

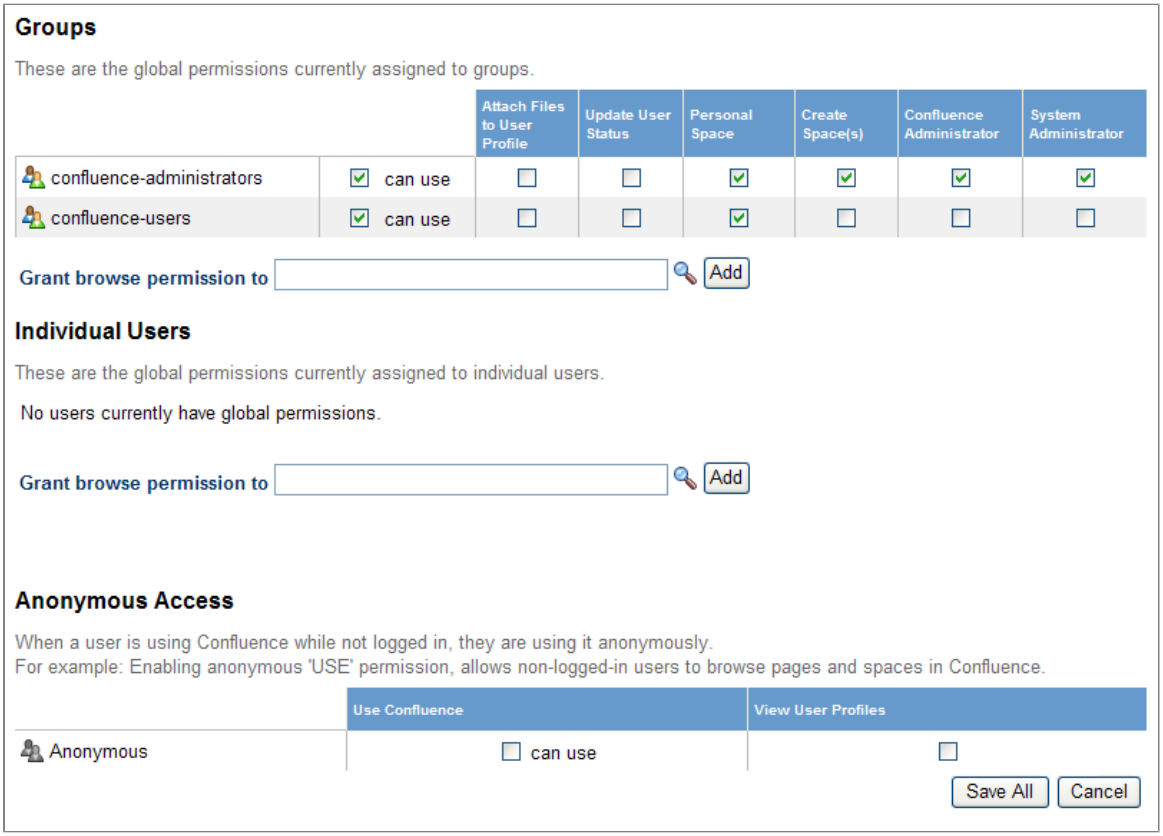
Read more about global groups.

Updating Global Permissions

To edit the global permissions for a group or user,

1. Go to the Confluence '**Administration Console**'. To do this:
    - Open the '**Browse**' menu and select '**Confluence Admin**'. The 'Administration Console' view will open.
2. Select '**Global Permissions**' in the '**Security**' section of the left-hand panel.
3. The '**View Global Permissions**' screen appears. Click the '**Edit Permissions**' button.
4. The '**Edit Global Permissions**' screen appears, as shown below. Add or edit group and user permissions as follows:
    - To add permissions for a group:
        a. First add the group to Confluence, if you have not already done so.
        b. Now on the 'Edit Global Permissions' screen, enter the group name in the text box labelled '**Grant browse permission to**' in the '**Groups**' section. You can click the magnifying glass to search for the group name.
        c. Click the '**Add**' button.
        d. The group will appear in the list and you can now edit its permissions.
    - To add permissions for a specific user:
        a. First add the user to Confluence, if you have not already done so.
        b. Consider adding the user to a group and then assigning the permissions to the group, as described above, instead of assigning permissions to the specific user.
        c. To assign permissions to a specific user on the 'Edit Global Permissions' screen, enter the username in the text box labelled '**Grant browse permission to**' in the '**Individual Users**' section. You can click the magnifying glass to search for the username.
        d. Click the '**Add**' button.
        e. The username will appear in the list and you can now edit its permissions.
    - To add or edit the permissions for a user or group:
        - Select the check box under the relevant permission and next to the relevant user/group. A tick in the check box indicates that the permission is granted. Click again to clear the check box and deny the permission.
    - To allow anonymous access to your Confluence site, select the 'Use Confluence' and 'View User Profile' options in the 'Anonymous Access' section.
      For more information about these permissions, refer to Setting up Anonymous Access.
    - Click the '**Save All**' button to save your changes.

Screenshot: Editing global permissions



About some error messages you may see
In **Confluence 2.7.2** and later, Confluence will let you know if there is a problem with some permissions. In rare situations, you may see the following error messages below a permission:

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

RELATED TOPICS

- Searching For and Managing Users
- Security Overview
- Adding or Removing Users in Groups
- Troubleshooting LDAP User Management
- Changing Usernames
- Setting up Anonymous Access
- Enabling or Disabling Public Signup
- Removing a Group
- Global Groups Overview
- Adding a Group
- Adding a New User
- Editing User Details
- Removing a User
- Global Permissions Overview
- Disabling the Built-In User Management

Showing first 15 of 16 results Administrators Guide Home Confluence Documentation Home

Migrating to new User Management

In order to support advanced forms of user management, Confluence now uses the AtlassianUser management framework, which can store users in the database (through Hibernate) instead of delegating the user management to OSuser.

Manual migration no longer required from Confluence 2.7.0
This page describes how to perform a manual migration of your users from OSuser to AtlassianUser. For **Confluence 2.7.0 and later**, there is no need to perform a manual migration of your users to the AtlassianUser framework. If you are installing Confluence 2.7.0 or later for the first time, you will automatically receive the AtlassianUser framework. If you are upgrading from an earlier version to Confluence 2.7.0 or later and have not changed the default user management configuration, your users will be automatically migrated. Refer to the details in the Confluence 2.7 Upgrade Guide.

If going from OsUser to LDAP, use Migrate to LDAP User Management From OsUser rather than this page. This page is intended for upgrading internal user management repositories.

For **Confluence 2.6.x** and earlier, Confluence delegates user management to OSuser by default. However, you may wish to migrate your users away from OSuser for one or more of the following reasons:

- In preparation to use external user management.
- To take advantage of the native AtlassianUser's much more efficient searching and user administration.
- Because you are having problems with OSuser, such as CONF-5218.

The migration instructions below are valid for Confluence version 2.2 and later.

Step 1 - Upgrade Confluence

Please check that you are running the latest version of Confluence. If not, we strongly recommend that you consider upgrading Confluence according to this guide. Confirm that you have upgraded successfully before trying to add LDAP to the new version.

Step 2 - Confluence User Migration

1. Find your Confluence base URL. To check this from Confluence, go to `Administration > General Configuration > Base Url`. Record this for later in the process.
2. Make a backup of your:
   - database
   - Confluence home directory
   - `confluence/WEB-INF/classes/atlassian-user.xml` (only if you have made changes)

   If you do not create a backup, you cannot roll back to the old version if the migration is unsuccessful.
3. Download ldap_hibernate_osuser_atlassian-user.xml, rename it to `atlassian-user.xml` and copy it to your `confluence/WEB-INF/classes` directory. (You can overwrite the one that is there).
   If you have already set up LDAP in your `osuser.xml` file and wish to migrate to atlassian-user LDAP, then you need to uncomment the ldap section, and fill in the correct details (as described in Customising atlassian-user.xml). This will prevent users in your osuser table that exist in LDAP from being migrated over. If you haven't already set up LDAP in `osuser.xml` please do NOT uncomment the ldap section.
4. Restart Confluence.
5. Log in as a System Administrator, copy the address `http://<BASEURL>/<contextpath>/admin/osuser2atluser.jsp` and paste it into your browser's address bar. Change <BASEURL> to your actual base URL and <contextpath> to your context path (usually 'confluence') and follow the link.
6. Click the link **Begin migration**. You will know the migration has been successful if you see this reported:

   ```
   1.Migrating users and groups ...
   2.Users and groups migrated successfully!
   ```

   If you encounter errors, please create a support ticket at http://support.atlassian.com and attach your application server logs.
7. Stop Confluence.
8. Start up Confluence and check that you can log in using the admin account you first set up when running through the Confluence Setup Wizard. If not, re-examine your steps and repeat from the point where you may have gone wrong.
9. Download hibernate_cache_atlassian-user.xml, rename it to `atlassian-user.xml` then copy to your `<INSTALL>/confluence/WEB-INF/classes` directory. It should overwrite the previous `atlassian-user.xml`.
10. Restart Confluence. Check that your users can still log in.

RELATED TOPICS

Understanding User Management in Confluence
Confluence 2.7 Upgrade Guide

Removing a Group

To remove a group,

1. Go to the Confluence '**Administration Console**'. To do this:
   - Open the '**Browse**' menu and select '**Confluence Admin**'. The 'Administration Console' view will open.
2. Select '**Manage Groups**' in the left panel. A list of all existing groups is displayed along with links to remove them.
3. Click '**Remove**' beside the group you want to remove. You will need to confirm your action before the group is deleted.

RELATED TOPICS

Configuring Character Encoding
Configuring Time and Date Formats
Configuring Attachment Size
Number of Ancestors to Show in Breadcrumbs

Thumbnail Settings

Configuring Number Formats

Configuring HTTP Timeout Settings

Configuring Indexing Language

Recognised System Properties

Removing a User

It is not possible to remove a user if the user is responsible for content on the site (another words, if a user has contributed content. For example, edit, create, or commented within Confluence), because Confluence will need the user information to maintain a history of pages. If you wish to prevent such a user from accessing Confluence, you can deactivate a user so that they can no longer log in to Confluence. Deactivating a user will not remove the content created by them from the site.

You can remove a user if the user has not added or edited any content on the site.

To deactivate or remove a user,

1. First, go to the user management screen for the user concerned. There are two ways to do this:
   - Either,
     - Go to the user's Profile and click the '**Administer User**' link on the user's profile screen. (This link is available in Confluence 2.8.2 and later.)
   - Or,

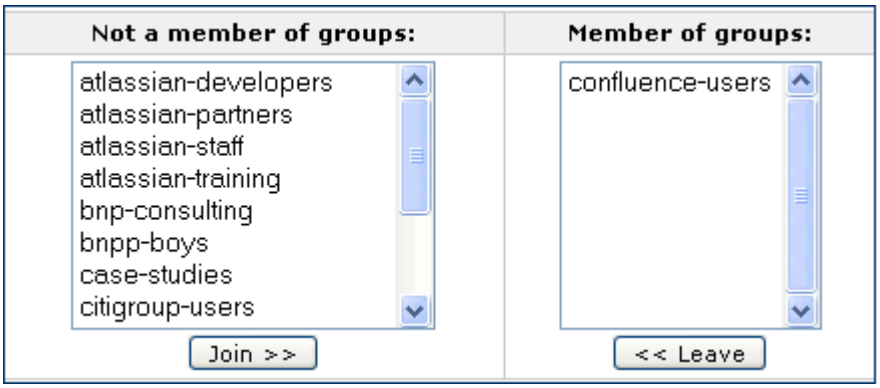
     Go to the Confluence '**Administration Console**'. To do this:

     - Open the '**Browse**' menu and select '**Confluence Admin**'. The 'Administration Console' view will open.
2. Select '**Manage Users**' in the left-hand panel.
3. Find the user by searching on the username, full name or email address. You can also click '**Show all users**' to browse the list of all users.
4. Click the username.
5. Now you should be able to see the user's current details and links allowing you to edit them.
6. Click '**Remove**' if the user is not responsible for any content on the site.
7. If the user is responsible for content on the site, you will need to deactivate the user (also known as disabling the user):
   - Remove the user from all groups that have the global '**Can Use**' permission. Click '**Edit Groups**' to remove the person from the group(s).
   - If the specific user has the global '**Can Use**' permission, you will also need to remove this permission.

Number of users and your license
The License Details page tells you how many users your Confluence instance is licensed to support, and how many are currently registered The number of registered users includes only users who have the 'Can Use' global permission. Deactivated users, as described above, are not included.

Deactivating users in Confluence 2.0.x and earlier
Earlier Confluence releases provide a '**Deactivate**' link. If you click this link, the user will be prevented from logging in to the Confluence site. This function was removed from Confluence 2.1.x and later for performance reasons.

Screenshot: Adding and removing users to/from groups



RELATED TOPICS

- Searching For and Managing Users
- Security Overview
- Adding or Removing Users in Groups
- Troubleshooting LDAP User Management
- Changing Usernames
- Setting up Anonymous Access
- Enabling or Disabling Public Signup
- Removing a Group
- Global Groups Overview
- Adding a Group
- Adding a New User
- Editing User Details
- Removing a User
- Global Permissions Overview
- Disabling the Built-In User Management
- Viewing members of a group

Administrators Guide Home Confluence Documentation Home

Setting up Anonymous Access

You can enable anonymous access (also known as public access) to your site by granting the 'Use Confluence' permission to 'Anonymous' users from the 'Administration Console'.

This user category has been created for convenient administration of users who have not logged into the site. Permissions assigned to this group apply to all anonymous users of the site.

To enable public access to your site,

1. Go to the Confluence '**Administration Console**'. To do this:
    - Open the '**Browse**' menu and select '**Confluence Admin**'. The 'Administration Console' view will open.
2. Select '**Global Permissions**' in the left panel.
3. Click '**Edit Permissions**'.
4. In the 'Anonymous Access' section, select the '**can use**' check box to enable basic public access to the content on your site.
5. If you selected the '**can use**' check box in the previous step and want to allow public access to user profile views, select the check box in the '**View User Profiles**' section.
   You cannot grant the 'View User Profiles' permission independently of the 'Use Confluence' permission.
6. Click the '**Save All**' button to save your changes.

Once you grant this permission, further permissions can be granted from the space administration screens to control the viewing and editing privileges of anonymous users. See Space Permissions Overview.

To disable public access to your site, clear the **can use** check box. If you also have the '**View User Profiles**' check box selected, it must be cleared before saving the configuration changes.

RELATED TOPICS

- Searching For and Managing Users
- Spam Prevention via Captcha
- Adding or Removing Users in Groups
- Setting up Anonymous Access
- Adding a New User
- Editing User Details
- Removing a User
- Global Permissions Overview

Administrators Guide Home Confluence Documentation Home

Viewing members of a group

To view the members of a group,

1. Go to the Confluence '**Administration Console**'. To do this:
    - Open the '**Browse**' menu and select '**Confluence Admin**'. The 'Administration Console' view will open.
2. Click '**Manage Groups**' in the left-hand panel. This will list all the existing groups on the site.
3. Click a group name to display all the users in the group.

RELATED TOPICS

- Searching For and Managing Users
- Adding or Removing Users in Groups
- Removing a Group
- Global Groups Overview
- Viewing members of a group

Administrators Guide Home Confluence Documentation Home

How to Improve User Search Performance

If your Confluence instance contains thousands of user accounts and you are experiencing performance issues when searching for users, the following migration guide is for you.

Background

In Confluence 2.1, we introduced a new system for user management inside Confluence (atlassian-user) that was more powerful than the previous system (OSUser). However, to avoid potential upgrade issues, we continued to use OSUser when storing users in the local Confluence database.

The native atlassian-user storage format provides much more efficient searching, and greatly improves the performance of user administration and Confluence's 'user picker' pop-up. We plan on migrating all Confluence instances to the new format around version 2.6 or 2.7, but until then Confluence instances with large numbers of users can still take advantage of these performance improvements by performing the migration manually.

Migration procedure

Do not use this procedure if you have LDAP user management enabled.

This guide assumes that you are using Confluence's local users and groups. If you have already configured Confluence for LDAP user/group management and are experiencing user management slowness, please follow the guide for Requesting External User Management Support.

Manual migration no longer required from Confluence 2.7.0
This page describes how to perform a manual migration of your users from OSuser to AtlassianUser. For **Confluence 2.7.0 and later**, there is no need to perform a manual migration of your users to the AtlassianUser framework. If you are installing Confluence 2.7.0 or later for the first time, you will automatically receive the AtlassianUser framework. If you are upgrading from an earlier version to Confluence 2.7.0 or later and have not changed the default user management configuration, your users will be automatically migrated. Refer to the details in the Confluence 2.7 Upgrade Guide.

For details of the procedure, refer to Migrating to new User Management.

Restoring Passwords To Recover Admin User Rights

Use this document if you are unable to login as administrator, to manually replace administrator passwords or give users administration rights.

New Confluence User Management

From Confluence 2.7 onwards the user management is handled by AtlassianUser. Hence in the database, Confluence will refer to 'USERS' table to store and refer to its users. When you imported your backup on upgrade, what should happen is the users in the 'OS_USER' table should get copied into 'USERS' table.

If you are still using OSUser please refer to our older document.

Learn more about the algorithm Confluence is using.

Stage One - Identify Administrator

To find out which usernames have admin privileges, connect to your database using a database admin tool such as DBVisualiser. Please download a database admin tool now if you do not have one installed already. Once installed, connect to your database and retrieve the list of administrator usernames with:

```
1.select name from users u, local_members l, groups g where g.groupname = 'confluence-administrators'
and g.id=l.groupid and u.id=l.userid;
```

Stage Two - Replace Administrator Password

Confluence does not store passwords in plain text in the database, but uses hashes computed from the original password. You instead cut and a paste a hash, rather than the plain password, over the existing password. Below is the hash for the password `admin`

```
1.x61Ey612Kl2gpFL56FT9weDnpSo4AV8j8+qx2AuTHdRyY036xxzTTrw10Wq3+4qQyB+XURPWx1ONxp3Y3pB37A==
```

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

1. Shutdown Confluence
2. Connect to your database.
3. The SQL to run is:

```
1.update users set password =
2.'x61Ey612Kl2gpFL56FT9weDnpSo4AV8j8+qx2AuTHdRyY036xxzTTrw10Wq3+4qQyB+XURPWx1ONxp3Y3pB37A=='
3.where name='<USERNAME>';
```

4. Start Confluence

5. Login with your modified username and use password `admin`

For the evaluation embedded database

1. Shut down Confluence.
2. Open <confluence-home>/database/confluence-db.script. Search for:

```
1.INSERT INTO USERS VALUES(
```

3. Replace the password for the appropriate user, you can copy and paste the hash value above.
4. Save the file, and restart.

If No Local Users Exist

In rare circumstances, when local users are deleted and only LDAP users exist, it may be required to insert a user. Here's how to do that:

```
01.insert into users (id, name, password, email, created, fullname) values (1212121, 'admin',
'x61Ey612Kl2gpFL56FT9weDnpSo4AV8j8+qx2AuTHdRyY036xxzTTrw10Wq3+4qQyB+XURPWx1ONxp3Y3pB37A==',
'a@a.com', '2009-11-26 17:42:08', 'admin');
02.
03.// Then find out the ID's of Groups
04.SELECT * FROM groups;
05.SELECT * FROM users;
06.
07.// Add group memberships into local_users
08.insert into local_members (userid, groupid) values (<from select above for
user>,<from_select_above_for_conf_users_group>);
09.insert into local_members (userid, groupid) values (<from select above for
user>,<from_select_above_for_conf_admin_group>);
```

Integrating with Crowd

Atlassian's Crowd identity management system can be integrated with Confluence. Please refer to the Crowd documentation on Integrating Crowd with Confluence.

JIRA User Management

- Delegate user management to use JIRA logins
  - Revert from JIRA to internal user management

Additional Information

- JIRA User Management FAQs

Delegate user management to use JIRA logins

If you already have a significant user base set up inside JIRA, it makes sense to connect Confluence to JIRA so that user management is centralised and not duplicated. This document outlines how to delegate Confluence's user authentication and group management to JIRA so that you can use your JIRA users to log in to Confluence.



Known Issues
Before attempting the steps below, please refer to the known issues listed in the troubleshooting section below.

On this page:

- Read Before Proceeding
- Technical Overview
- Step One: Installing Confluence
- Step Two: Setting up a Datasource to JIRA's Database
- Step Three: Installing the JDBC Driver
- Step Four: Modifying osuser.xml
- Step Five: Customising osuser.xml
- Step Six: Modifying atlassian-user.xml
- Step Seven: Creating Confluence Groups in JIRA
- Step Eight: Activating External User Management

Read Before Proceeding

1. The examples used in this document are based on **Tomcat Application Server** and the **MySQL database**. The same concepts (but

not the verbatim examples) can be applied to other application servers or databases.

2. If JIRA is using LDAP for authentication, you should not use JIRA for Confluence user management. Use Add LDAP Integration instead.

3. Always install Confluence with a new database. Do **not** attempt to use the existing JIRA database, with either JDBC or data source.

4. If you have existing users or groups in Confluence, these users will not be available once you switch to using JIRA's user management. Any existing content will no longer be associated with valid users. Do not add any spaces or content once Confluence installation is complete, apart from verifying that your Confluence instance is up and working, until you have completed the procedure on this page. Users in Confluence will no longer be valid once you switch over to using your JIRA users.

5. If you run into a problem, check the Troubleshooting section below.

Technical Overview

In the configuration described below, Confluence will use JIRA's database for its user and group information. The Confluence application will have two database connections:

1. A connection to the primary database, set up during Confluence installation. This database stores all the normal Confluence data: spaces, pages, comments, etc.
2. A read-only data source connection to JIRA's database, set up after Confluence is installed. Confluence reads information about users and groups from this database.

The reason this works is because both JIRA and Confluence use the same user management library, OSUser. The OSUser database schema is the same in JIRA and Confluence, so Confluence can easily read from JIRA's tables to get the user and group information.

Step One: Installing Confluence

Skip this step if you have already installed Confluence and completed the Setup Wizard.

1. Install Confluence and run the Setup Wizard:
    - If you are running *JIRA standalone* please follow these instructions for installing Confluence.
    - If you have JIRA deployed under your own Tomcat server, please follow these instructions.

2. Ensure that Confluence is running and has been set up, that is, you have completed the Confluence Setup Wizard and verified that you can create pages.

3. Shut down Confluence.

Step Two: Setting up a Datasource to JIRA's Database

In order to delegate all user authentication attempts and group membership queries to JIRA, Confluence needs to be aware of JIRA's database, and hence the user tables in JIRA's database.

In Tomcat this is achieved by specifying JIRA's database as a resource. You will need to declare it inside the `<context>` descriptor you set up in *Step One*.

Notes:

- If there is an existing block of `<Resource>` in the `<context>` descriptor, please do not replace it. Rather, just add the following `<Resource>` block inside the `<context>` descriptor.
- If you are running Confluence EAR/WAR distribution separate to JIRA, or under JIRA standalone 3.3 and later, your Confluence context will be in the `confluence.xml` file.
  If you are running Confluence standalone separate to JIRA, or Confluence WAR/WAR distribution under an older version of JIRA, your Confluence context will be in the `server.xml` file. You should never have a Confluence context in both.
- If you are running Confluence standalone (or Confluence inside a JIRA standalone) and are not sure which version of Tomcat you are using, check your log files. You'll see *INFO: Starting Servlet Engine: Apache Tomcat/5.5.nn* if you are using Tomcat 5.5.
- The DataSource configuration below uses MySQL as an example. You will need to modify these settings according to the database that you are using.
- You should add the appropriate validation check for the connection pool to avoid intermittent problems authenticating.

Sample context descriptor for Tomcat 4.x and 5.0.x

```
<Context path="/confluence" docBase="C:/programs/confluence" swallowOutput="true">
          <Resource name="jdbc/JiraDS" auth="Container" type="javax.sql.DataSource"/>
          <ResourceParams name="jdbc/JiraDS">
          <parameter>
          <name>username</name>
                    <value>your_db_username</value>
          </parameter>
          <parameter>
          <name>password</name>
                    <value>your_db_password</value>
          </parameter>
          <parameter>
          <name>driverClassName</name>
                    <value>com.mysql.jdbc.Driver</value>
          </parameter>
          <parameter>
          <name>url</name>
                    <value>jdbc:mysql://your.domain.com/jira_database_name?autoReconnect=true&lt
;/value>
          </parameter>
          <parameter>
                    <name>factory</name>
                    <value>org.apache.commons.dbcp.BasicDataSourceFactory</value>
          </parameter>
          </ResourceParams>
</Context>
```

Tomcat 5.5.x

This version of Tomcat has a new syntax for specifying resources. Note that you don't add a new context to `server.xml`, just add the Resource to your existing Context:

```
<Context path="/confluence" docBase="C:/programs/confluence" swallowOutput="true">
     <Resource name="jdbc/JiraDS" auth="Container" type="javax.sql.DataSource"
               username="your_db_username"
               password="your_db_password"
               driverClassName="com.mysql.jdbc.Driver"
               url="jdbc:mysql://your.domain.com/jira_database_name?autoReconnect=true"/>
</Context>
```



Some customers have encountered periodic login failures after delegating user management to JIRA. This usually occurs after a network failure or after the database server reboots. If you do experience such problems after completing this procedure, please refer to Surviving Connection Closures for further information on resolving them. As explained in this document, resolving these issues typically involves adding a `validationQuery` parameter to your JIRA datasource connection definition (such as the example `jdbc/JiraDS` definition above).

Step Three: Installing the JDBC Driver

Ensure that your JDBC driver is on the classpath of your application server. In this example, a jar for the mysql driver should be in the `. . . /common/lib` folder (or potentially `. . . /lib` for Tomcat version 6 and beyond).

1. Download the mysql driver from here.
2. Copy the jar file into the `. . . /common/lib` folder (or `. . . /lib`).

Step Four: Modifying osuser.xml



Please make sure you have completed the Confluence Setup Wizard before performing this step.

1. Find the `osuser.xml` file in the `/confluence/WEB-INF/classes` folder and open it in a text editor. Comment out the following block of code:

```
<provider class="bucket.user.providers.CachingCredentialsProvider">
                    <property name="chain.classname"
>com.opensymphony.user.provider.hibernate.HibernateCredentialsProvider</property>
                    <property name="chain.configuration.provider.class"
>bucket.user.BucketHibernateConfigProvider</property>
          </provider>
          <provider class="bucket.user.providers.CachingAccessProvider">
                    <property name="chain.classname"
>com.opensymphony.user.provider.hibernate.HibernateAccessProvider</property>
                    <property name="chain.configuration.provider.class"
>bucketuser.BucketHibernateConfigProvider</property>
          </provider>
          <provider class="bucket.user.providers.CachingProfileProvider">
                    <property name="chain.classname"
>com.opensymphony.user.provider.hibernate.HibernateProfileProvider</property>
                    <property name="chain.configuration.provider.class"
>bucket.user.BucketHibernateConfigProvider</property>
          </provider>
```

2. Uncomment this block:

```
<provider class="bucket.user.providers.CachingCredentialsProvider">
                    <property name="chain.classname"
>com.atlassian.confluence.user.providers.jira.JiraJdbcCredentialsProvider</property>
                    <property name="chain.datasource"
>java:comp/env/jdbc/JiraDS</property>
          </provider>
          <provider class="bucket.user.providers.CachingAccessProvider">
                    <property name="chain.classname"
>com.atlassian.confluence.user.providers.jira.JiraJdbcAccessProvider</property>
                    <property name="chain.datasource"
>java:comp/env/jdbc/JiraDS</property>
          </provider>
          <provider class="bucket.user.providers.CachingProfileProvider">
                    <property name="chain.classname"
>com.atlassian.confluence.user.providers.jira.JiraJdbcProfileProvider</property>
                    <property name="chain.datasource"
>java:comp/env/jdbc/JiraDS</property>
                    <property name="chain.configuration.provider.class"
>bucket.user.BucketHibernateConfigProvider</property>
          </provider>
```

Your `osuser.xml` should now look like this:

```
<opensymphony-user>
        <!--
                Authenticators can take properties just like providers.

                This smart authenticator should work for 'most' cases - it dynamically
looks up
                the most appropriate authenticator for the current server.
        -->
        <authenticator class="com.opensymphony.user.authenticator.SmartAuthenticator"/>

        <!-- JIRA User management (with caching) -->
        <!-- Note: Do not add any line breaks or spaces when specifying the
chain.classname, otherwise a ClassNotFoundException will be thrown -->

        <provider class="bucket.user.providers.CachingCredentialsProvider">
                <property name="chain.classname"
>com.atlassian.confluence.user.providers.jira.JiraJdbcCredentialsProvider</property>
                <property name="chain.datasource"
>java:comp/env/jdbc/JiraDS</property>
        </provider>
        <provider class="bucket.user.providers.CachingAccessProvider">
                <property name="chain.classname"
>com.atlassian.confluence.user.providers.jira.JiraJdbcAccessProvider</property>
                <property name="chain.datasource"
>java:comp/env/jdbc/JiraDS</property>
        </provider>
        <provider class="bucket.user.providers.CachingProfileProvider">
                <property name="chain.classname"
>com.atlassian.confluence.user.providers.jira.JiraJdbcProfileProvider</property>
                <property name="chain.datasource"
>java:comp/env/jdbc/JiraDS</property>
                <property name="chain.configuration.provider.class"
>bucket.user.BucketHibernateConfigProvider</property>
        </provider>

        <!--
        <provider class="bucket.user.providers.CachingCredentialsProvider">
                <property name="chain.classname"
>com.opensymphony.user.provider.hibernate.HibernateCredentialsProvider</property>
                <property name="chain.configuration.provider.class"
>bucket.user.BucketHibernateConfigProvider</property>
        </provider>
        <provider class="bucket.user.providers.CachingAccessProvider">
                <property name="chain.classname"
>com.opensymphony.user.provider.hibernate.HibernateAccessProvider</property>
                <property name="chain.configuration.provider.class"
>bucket.user.BucketHibernateConfigProvider</property>
        </provider>
        <provider class="bucket.user.providers.CachingProfileProvider">
                <property name="chain.classname"
>com.opensymphony.user.provider.hibernate.HibernateProfileProvider</property>
                <property name="chain.configuration.provider.class"
>bucket.user.BucketHibernateConfigProvider</property>
        </provider>
        -->

        <!--<provider class=
"com.opensymphony.user.provider.memory.MemoryCredentialsProvider" />
        <provider class="com.opensymphony.user.provider.memory.MemoryAccessProvider" />
        <provider class="com.opensymphony.user.provider.memory.MemoryProfileProvider"
/>-->
</opensymphony-user>
```

In this example, **JiraDS** is the name of the JIRA datasource you are sharing with Confluence. If you have changed the name in *Step Two* of this documentation, you will need change all occurrences of the value here too.

You can also download the already configured file here.

Step Five: Customising osuser.xml

In some cases you may need to customise the behaviour of the JiraJdbc classes. You can do this by setting properties within the `osuser.xml` file.

This process is documented here.

Step Six: Modifying atlassian-user.xml

 This step is **only applicable for Confluence 2.7 and later** .

Please comment out or remove the following line from your `<Confluence-Install>/confluence/WEB-INF/classes/atlassian-user.xml` file:

```
<hibernate name="Hibernate Repository" key="hibernateRepository"  description="Hibernate
Repository" cache="true"/>
```

and add this line instead:

```
<osuser name="OSUser Repository" key="osuserRepository"/>
```

Step Seven: Creating Confluence Groups in JIRA

1. Add the `confluence-users` and `confluence-administrators` groups in JIRA.
2. Add yourself to both these groups.
3. To give your existing JIRA users access to Confluence, you have two options.
   - Option 1: Manually edit the groups of these users inside JIRA and give them membership to one or both of these confluence groups.
   - Option 2: Start up Confluence. Log in using your JIRA account and go to **Administration** and then **Global Permissions**. Now assign the '**can use**' permission to your desired JIRA groups.

> In order to use Confluence, users must be a member of the `confluence-users` group (or have Confluence 'can use' permission).

Step Eight: Activating External User Management

Since user management is now conducted in JIRA and outside of Confluence, you will need to switch external user management on.

> **Activating external user management will remove user and group management options from Confluence.**
> Your users will also no longer be able to edit their **full name** or **email address** inside Confluence. (If they want to, they would have to do so in JIRA).

To switch external user management on:

1. Log into Confluence using your JIRA account.
2. Go to the **Administration Console** and click **General Configuration** in the left-hand panel
3. Click '**Edit**' at the bottom of the '**Options and Settings**' screen.
4. Select '**ON**' beside '**External User Management**'.

For troubleshooting, see the JIRA Integration FAQ.

RELATED TOPICS

- Migrating users from Confluence to JIRA
- Revert from JIRA to internal user management
- Delegate user management to use JIRA logins

Administrators Guide Home Confluence Documentation Home

Revert from JIRA to internal user management

> Check out Crowd for a fully featured user management solution.

Administrators can revert a Confluence instance that uses JIRA for user management back to internal user management. With few users, it is easier to manually recreate the JIRA users and groups in Confluence. For more users, migrate JIRA users and groups into the Confluence database instead.

Option A - Manually Recreate Users In Confluence

This option is too time consuming for hundreds or thousands of users. After completing the reversion, links to users who created or updated Confluence content may go to error screens.

To manually recreate the users, you must first have an instance of Confluence with internal user management and your data.

- If you have made limited customisations to Confluence and migrating would be desirable - follow the upgrade guide and import your data to a new installation.
- Alternatively, if you have made extensive customisations or do not wish to migrate - go to delegating user management to JIRA and remove your JIRA user management by undoing the instructions in revese order. These steps are specific to your instance so cannot be covered here.

Then manually create JIRA's groups and users in Confluence. If you have assigned permissions in Confluence to a group which exists in JIRA, you must create a group in Confluence with the same name. If a user who exists in JIRA has created content or has had permissions assigned to them, you must also create that user in Confluence.

Option B - Transfer JIRA Users & Groups To Confluence

This option manually migrates JIRA users into the Confluence database, but requires knowledge of SQL.



Users not using mySQL
Users of non-mySQL databases must be experienced enough to modify the SQL to work in their database as examples are provided for mySQL only. If you adapt the SQL to another database, please consider posting the SQL you used to the comments.



Users of Confluence 2.0 or older
Pre-Confluence 2.0 users may need to modify the instructions to your older schema, or upgrade Confluence. For example, on Confluence 2.1.5 and older, SQL references to the property table must be updated to be called OS_PROPERTYENTRY in all upper-case.

Stage One - Create Backups

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

1. From Confluence, create a full XML backup including attachments.
2. Stop Confluence.
3. Take a backup copy of the Confluence home and install directories.
4. Repeat the above steps for JIRA.
5. From your mySQL admin tool, create a database backup for the JIRA and Confluence databases.

Stage Two - Replace Confluence User Management

Replace the Confluence user and group permissions with JIRA by transferring table content. The SQL provided is specific to mySQL and must be modifed for other databases. For each SQL statement, do a find and replace on the JIRA and Confluence table names to your table names. In the examples, they are called confluence224 and jira364.

1. Login to a DBA tool that can execute SQL on your DB.
2. Erase user and group content from the Confluence DB:

```
1.delete from confluence224.os_propertyentry where entity_name='OSUser_user';
2.delete from confluence224.os_user_group;
3.delete from confluence224.os_group;
4.delete from confluence224.os_user;
```

3. Copy JIRA's groupbase table into Confluence's os_group table:

```
1.insert into confluence224.os_group (id, groupname)
2.select *
3.from jira364.groupbase;
```

4. Copy JIRA's userbase table into Confluence's os_user table:

```
1.insert into confluence224.os_user (id, username, passwd)
2.select *
3.from jira364.userbase;
```

5. Copy JIRA's membershipbase table into Confluence's os_user_group table.

```
1.insert into confluence224.os_user_group (group_id, user_id)
2.select distinct groupbase.id as "group_id", userbase.id as "user_id"
3.from jira364.groupbase, jira364.membershipbase, jira364.userbase
4.where membershipbase.user_name = userbase.username and membershipbase.group_name =
groupbase.groupname;
```

6. Merge relevant content from JIRA's propertyentry and propertystring tables into Confluence's os_propertyentry table. Some versions of SQL use "0" instead of "false" for boolean values.

```
1.insert into confluence224.os_propertyentry (entity_name, entity_id, entity_key, key_type,
boolean_val, double_val, string_val, text_val, long_val, int_val, date_val)
2.select 'OSUser_user', propertyentry.entity_id, propertyentry.property_key, 5, false, 0,
propertystring.propertyvalue, '', 0, 0, null
3.from jira364.propertyentry, jira364.propertystring
4.where propertyentry.entity_name='OSUser' and propertyentry.id=propertystring.id;
```

Stage Three - Revert To Local Management

- If you have made limited customisations to Confluence and migrating would be desirable - Install a new instance of Confluence using the upgrade guide.
- Alternatively, if you have made extensive customisations or do not wish to migrate - go to Delegating User Management to JIRA and remove your JIRA user management by undoing the instructions in revese order. These steps are specific to your instance so cannot be covered here.
- If you wish to migrate to LDAP user management at this point, you can follow the instructions to Migrate to LDAP User Management From OsUser.

Done! Note that the original administrator may not display their groups correctly, however their groups are still present.

LDAP User Management

Overview

Confluence integrates with LDAP user repositories in a variety of ways. Start with the Overview of external user management to learn more, then choose your preferred LDAP connectivity.

There are three choices for LDAP integration:

LDAP Configuration	Internal Users	Internal Groups	LDAP Authentication	LDAP Users	LDAP Groups
LDAP for Authentication Only	✓	✓	✓	✓	✗
LDAP for Group and User Management	✓	✓	✓	✓	✓
LDAP Authentication with OSUser (not supported after 2.7)	✓	✓	✓	✗	✗

In addition to these three options, you may also choose to delegate user management to JIRA or use Crowd for a full-featured solution for multiple applications, directories, and single signon. JIRA user management is a good solution if you already manage users in JIRA.

Additional Information

- Help with troubleshooting external user management
- LDAP FAQ

Check the listing of children pages below for other useful resources.

Add LDAP Integration

Try Atlassian Crowd for powerful LDAP integration
Atlassian's Crowd is a web-based single sign-on (SSO) tool that simplifies application provisioning and identity management.

Confluence can delegate user authentication to LDAP and use LDAP group memberships to set the user's Confluence access permissions. This also allows Active Directory (AD) integration. This guide is for both users enabling LDAP, and those upgrading their LDAP scheme to support group management. It applies to LDAP over HTTP and SSL/HTTPS.

Once the LDAP is enabled and LDAP users are using Confluence, you cannot revert to local user management without those users being disabled. However, you can create new local users while using LDAP integration.

Who is this guide for?

If you are using local user management in a version prior to Confluence 2.7, or os_user with authentication-only or jira user management, follow the guide to Migrate to LDAP User Management From OsUser. Otherwise, this is the correct guide for you.



Integrate only after completing Setup
If you are doing an LDAP integration as part of a new install, do not integrate until after you complete the initial setup. You can add LDAP integration after you create the admin user for your instance.

Step 1 - Upgrade Confluence

Please check that you are running the latest version of Confluence. If not, we strongly recommend that you consider upgrading Confluence according to this guide. Confirm that you have upgraded successfully before trying to add LDAP to the new version.

Step 2 - Contact your LDAP/AD Administrator

Integration can only be setup by an administrator confident with running user queries against their LDAP directory. You should request assistance from your LDAP or Active Directory administrator for the following steps.

Step 3 - Check your LDAP server

Confirm this information about your LDAP server.

1. Check your server LDAP version. Supported versions are v2 and v3. Supported LDAP servers include OpenLDAP, Microsoft Active Directory, Novell eDirectory, and any server that uses Java JNDI-LDAP mapping.
2. Your LDAP or Active Directory server must support static groups. This means that the *user DNs* must be stored against a membership attribute inside an LDAP groups. An example of a static group is shown below:

```
01.Dn: CN=Sales and Marketing,CN=Users,DC=ad,DC=atlassian,DC=com
02.objectClass: top; group;
03.cn: Sales and Marketing;
04.distinguishedName: CN=Sales and Marketing,CN=Users,DC=ad,DC=atlassian,DC=com;
05.name: Sales and Marketing;
06....
07.member: CN=John Smith,CN=Users,DC=ad,DC=atlassian,DC=com
08.member: CN=Sally Smith,CN=Users,DC=ad,DC=atlassian,DC=com
09....
```

   The membership attribute in this case is **member**, but this is not required. Note that the full DNs of *John and Sally Smith* are listed. If the values against **member** are not full DNs, but are just usernames, then you need to add the flag

```
1.<useUnqualifiedUsernameForMembershipComparison>true
</useUnqualifiedUsernameForMembershipComparison>
```

   to your LDAP tag in atlassian-user.xml. Open Directory on OS X uses this configuration.
3. You must not have LDAP groups called 'confluence-users' or 'confluence-administrators'.
4. You must have at least one existing Confluence administrator with System Administrator permissions, whose username does not exist in the LDAP server (see Step 4).

Step 4 - Check the System Administrator account

This step assumes that you have at least one Confluence user account which has System Administrator permissions for your Confluence site. For this account, please check that there isn't an account on your LDAP system that has the *exact same username*.

If there ***is*** an LDAP account with the exact same username, and you do not have another local Confluence account that has System Administrator permissions rights, then you should perform *one* of the following:

- create another account, that *doesn't* exist on LDAP, to act as the administrator
  OR:
- rename your local Confluence administrator account to use another username that doesn't exist in LDAP
  OR:
- rename your LDAP account

This will ensure that you will have an account *that has sufficient rights* to administer your site after you migrate your users.

Step 5 - Configure your LDAP repository

1. Follow Customising atlassian-user.xml
2. Start up Confluence and check that you can log in using the System Administrator account you first set up when running through the Confluence Setup Wizard. If not, re-examine your steps and repeat where necessary.
3. If you can't successfully log in with this account, please check that the **username** of this account does not **already exist** in your LDAP server. If *usernames are the same, Confluence recognises LDAP accounts over local Confluence accounts.*
4. If you were using OS user previously, run the user migration. After the migration has run, remove the os user tag from atlassian-user.xml and restart Confluence.

Step 6 - Grant access to LDAP users and groups

To grant Confluence login access to your LDAP groups and users,

1. Go to the Confluence '**Administration Console**'. To do this:
    - Open the '**Browse**' menu and select '**Confluence Admin**'. The 'Administration Console' view will open.
2. Select '**Global Permissions**' in the left panel.
3. Click to **Edit Permissions for Groups**.
4. In the textbox to '**Grant Browse Permission**', enter the name of an LDAP group that should have Confluence access. Click '**Add**'.
5. Tick the Can Use box for the LDAP group. If the group is not found, it was not present in your LDAP server.
6. For other LDAP groups that need access to Confluence, add them using the same method.
7. If you are integrating LDAP with Confluence for authentication only, no LDAP groups will appear in Confluence. All the individual LDAP users will have to be manually added to an internal Confluence group with Can Use permissions enabled before they can have access to Confluence.
8. Set up your Confluence page and space permissions for these LDAP groups and users.

Installation complete!

Related Pages

- Confluence LDAP Documentation Index

Troubleshooting

Local user management not retained

If you run into this problem, you may be experiencing this bug.

Check your Confluence version

This documentation applies to the latest version of Confluence. There are a couple of key bugs that have been resolved in Confluence 2.6 or 2.6.1, but that pertain to 2.5.6 and 2.5.7.

1. http://jira.atlassian.com/browse/CONF-9434 relates to hibernate cache=true;
   The xml file supplied here has the hibernate cache set to "true".
2. http://jira.atlassian.com/browse/CONF-9195 relates to the migration step.
   Version 2.6.1 corrects this problem.

More information

- Browse the LDAP FAQ.
- If LDAP users or groups are not displayed in Confluence, try the External User Test tool.
- Check the list of known, unresolved LDAP bugs
- See the comments on this page, from other users who may have left some useful information.
- The 'External User Management' setting in the Confluence Administration Console should be set to OFF. This setting is for using JIRA or Crowd for External User Management.

Support

Failing all else, lodge a support request. Be sure to attach your `atlassian-user.xml`, a copy of the output from the External User Test tool, and a zip of your Confluence logs.

Automatically Add LDAP users to the confluence-users Group

Users in your LDAP repository will not automatically have the right group membership to be able to access your Confluence instance. In most cases, this is due to Confluence access being limited to those in the `confluence-users` group. This page describes options for automatically adding users into this group, when they log in for the first time.

There are three Authenticators you can use to simplify the integration with LDAP.

Authenticator	Purpose	Configuration
ConfluenceAuthenticator	This authenticator will not add any users to `confluence-users` automatically	`1.<authenticator class= "com.atlassian.confluence.user.ConfluenceAuthenti`
ConfluenceGroupJoiningAuthenticator	This authenticator will add **ALL** users to `confluence-users` automatically	`1.<authenticator class= "com.atlassian.confluence.user.ConfluenceGroupJoi`

ConfluenceLDAPGroupJoiningAuthenticator	This authenticator will only add **LDAP** users to `confluence-users` automatically	`1.<authenticator class= "com.atlassian.confluence.user.ConfluenceLDAPGrou />`

To use these authenticators, find the `authenticator` element in your `seraph-config.xml` file and replace this element with the appropriate XML snippet from the table above.

In both of the joining cases above, users will only be added to the `confluence-users` group if they at least have `VIEW` permissions to your application.



Issue CONF-17366
There is currently a problem in that the above procedure does not work for users who are already authorised to use Confluence. The cause is that the authenticator does not add LDAP users to the `confluence-users` group if they already have the 'Can Use' permission via another group or individual global permission. Please follow CONF-17366 to see the progress of this issue.



Issue CONF-13754
If you are experiencing performance problems when logging in, it may be due to CONF-13754. You can download the following files and put them in `<confluence-installation-directory>/confluence/WEB-INF/classes/com/atlassian/confluence/user` : ConfluenceGroupJoiningAuthenticator.class and ConfluenceGroupJoiningAuthenticator$1.class. This will override the built-in version of ConfluenceGroupJoiningAuthenticator in Confluence 3.0.1 and earlier versions of Confluence.

Customising atlassian-user.xml

The LDAP server connection is specified by manually editing the file `atlassian-user.xml`. Confluence 2.3 onwards supports multiple LDAP servers by repeating the instructions below for each server so that there are multiple repositories defined.

Stage 1 - Determine if you are migrating from os_user or atlassian_user

Check your current `.../confluence/WEB-INF/classes/atlassian-user.xml` file. If your os user tag is active, you're using os user in some fashion. Make sure to run the user migration after doing your LDAP details, as described in Migrating to new User Management.

This usually only applies to versions of Confluence prior to 2.7.

Stage 2 - Configure Connection Details



To make life easier, use Paddle and Apache Directory Studio to test your LDAP connections without restarting Confluence.

1. Edit the file `.../confluence/WEB-INF/classes/atlassian-user.xml` and configure the connection for either AD or LDAP.
   - Connections in Active Directory
   - Connections in other LDAP servers
2. If your Active Directory Server allows anonymous searches, then you *do not* need to specify a securityPrincipal and securityCredential at all. For an example of how you would configure Confluence to allow anonymous authentication, see Enable Anonymous Authentication in LDAP or Active Directory
3. To connect to LDAP over SSL, see Connect to LDAP via SSL
4. Check your configuration against the example connection details shown below.



If you change your ldap key to a different name, you will need to change the cache name to pick it up. This is described in CONFKB181536872.

```
01.<ldap key="ldapRepository" name="LDAP Repository@hecate.atlassian.com" cache="true">
02.          <host>hecate.atlassian.com</host>
03.          <port>389</port>
04.          <securityPrincipal>cn=admin,dc=atlassian,dc=private</securityPrincipal>
05.          <securityCredential>secret</securityCredential>
06.          <securityProtocol>plain</securityProtocol>
07.          <securityAuthentication>simple</securityAuthentication>
08.          <baseContext>dc=atlassian,dc=private</baseContext>
09.
10.....
```

Stage 3 - Map LDAP Data Tree

1. Configuring the mappings in `atlassian-user.xml` for either AD or LDAP.
    - Mapping Active Directory
    - Mapping other LDAP servers
2. Check your configuration against the example connection details shown below.

```
01....
02.
03.          <baseUserNamespace>dc=staff,dc=perftest,dc=atlassian,dc=private</
baseUserNamespace>
04.          <baseGroupNamespace>dc=groups,dc=perftest,dc=atlassian,dc=private</
baseGroupNamespace>
05.          <usernameAttribute>cn</usernameAttribute>
06.          <userSearchFilter>(objectClass=inetorgperson)</userSearchFilter>
07.          <firstnameAttribute>givenname</firstnameAttribute>
08.          <surnameAttribute>sn</surnameAttribute>
09.          <emailAttribute>mail</emailAttribute>
10.          <groupnameAttribute>cn</groupnameAttribute>
11.          <groupSearchFilter>(objectClass=groupOfNames)</groupSearchFilter>
12.          <membershipAttribute>member</membershipAttribute>
13.</ldap>
```

Stage 4 - Directory Search Depth Settings

These are the default settings:

```
1....
2.<userSearchAllDepths>false</userSearchAllDepths>
3.<groupSearchAllDepths>false</groupSearchAllDepths>
```

The above settings configure the search depth on users and groups. If you set either attribute to:

- false - Confluence will search only for users/groups directly defined in <baseUserNamespace> and <baseGroupNamespace>.
- true - Confluence will search for users/groups defined in the above namespaces and also in namespaces nested within them. For example, if your users are distributed across multiple namespaces, you should set this option to **true**.

Setting the value to true may have a high cost in performance for large directories, because Confluence will search the whole tree and not just the immediate namespace.

Stage 5 - Optional LDAP Settings

The following settings are the default values for all the options under <ldap> in the `atlassian-user.xml` file. Some of them do not appear in the file normally, but can be added if you need to customise them:

XML tag	Default value	Description
LDAP connection properties		
`securityProtocol`	plain ssl	Allow encrypted (SSL) connections. Can be omitted if anonymous access is available.
`authentication`	simple	Plain-text password transmission. Can be 'none' if anonymous access to the LDAP server is available.
`initialContextFactory`	com.sun.jndi.ldap.LdapCtxFactory	Class name of LDAP provider (default: Sun JNDI)
`batchSize`	100	Size of pages in search results
`poolingOn`	true	Use connection pooling
`connectTimeout`	30000	Timeout in milliseconds when opening new server connections. Default: 30 seconds.
`readTimeout`	60000	Timeout in milliseconds for search and other read operations. Default: 60 seconds.
LDAP connection pool properties		
`initSize`	1	Initial size of connection pool, e.g. number of connections to open at start-up. [1]
`prefSize`	10	Preferred size of connection pool. [1]

maxSize	0	Maximum size of connection pool. Zero means no maximum size. [1]
timeout	300000	Idle time in milliseconds for a connection before it is removed from the pool. Default: 5 minutes. [1]
debugLevel	none	Debug level for logging. [1]
poolAuthentication	simple	Authentication for pool connections. [1]
LDAP search properties		
timeToLive	0	Time limit on searches in milliseconds. Zero means no limit. [2]
userSearchAllDepths	false	Whether user searches should search through the LDAP tree or only for direct children of the DN specified by the userSearchFilter. [3]
groupSearchAllDepths	false	Whether group searches should search through the LDAP tree or only for direct children of the DN specified by the groupSearchFilter. [3]
useUnqualifiedUsernameForMembershipComparison	false	If set to true, Confluence will use the value of the usernameAttribute on the user to check for group membership comparisons instead of the complete distinguished name.

Notes

1. The connection pool properties provided by Atlassian-User correspond with the connection pooling properties in JNDI. See this documentation for further information.
2. More information on this time limit is available on Sun's JNDI tutorial.
3. The "searchAllDepths" parameters toggle between SearchControls.SUBTREE_SCOPE (true) and SearchControls.ONELEVEL_SCOPE (false). See Sun's JNDI tutorial on scope configuration for more information.

To override the default values listed above, you can add the value inside your <ldap> tag after the rest of your configuration:

```
1.<ldap ...>
2....
3.          <groupnameAttribute>cn</groupnameAttribute>
4.          <groupSearchFilter>(objectClass=groupOfNames)</groupSearchFilter>
5.          <membershipAttribute>member</membershipAttribute>
6.          <initSize>20</initSize>      <!-- bigger initial connection pool! -->
7.</ldap>
```

Stage 6 - Optional: Configure LDAP for User Authentication Only

If you'd like to configure LDAP for user management and authentication only, start by configuring LDAP integration with groups then follow the instructions for authentication-only LDAP to disable the group management functionality.

Optional: Configuring Multiple LDAP Repositories

For some LDAP servers, it might be necessary to configure Confluence to connect to multiple LDAP servers. This functionality is available in Confluence 2.3 and above, and has a separate guide: Configuring multiple LDAP repositories.

RELATED TOPICS

Add LDAP Integration
Configuring multiple LDAP repositories
atlassian-user.xml reference
How to write a LDAP search filter, in the *Atlassian Development* space.

Migrate to LDAP User Management From OsUser

Is this guide for me?

Use Add LDAP Integration instead if:

- You are setting up Confluence for the first time.
- You do not need to retain group membership for existing users.
- You are using hibernate user management. To find out, check

`<confluence-home>/confluence/WEB-INF/classes/atlassian-user.xml`. If there is a hibernate tag but not an os_user tag, you're using hibernate user management.

Background Information

Confluence has three types of user management: os_user (deprecated in confluence 2.7 and later), hibernate (also known as atlassian-user), and LDAP integration.

These correspond to the three tags in atlassian-user.xml:

- <ldap key=...>
- <hibernate name=...>
- <osuser key=...>

In Confluence 2.7, during the upgrade task, users are migrated from os_user to hibernate when Confluence is first started. This migration can also be run manually.

When to Run the User Migration

Under most conditions, you'll want to migrate from OsUser to Hibernate User Management. However, there is one condition under which you'll want to hold off on doing the automatic user management migration when upgrading.

As described in this jira issue, local groups are not retained when migrating from hibernate user management to LDAP. However, if you are using OsUser for user management and want to switch to LDAP, you can retain group membership by following this procedure:



If you want to upgrade to 2.7 or later and do LDAP integration, we suggest you do your LDAP integration first, then run the upgrade. If you want to run the upgrade first, you can prevent the user migration from occurring by copying your <confluence-home>/confluence/WEB-INF/classes/atlassian-user.xml file and your osuser.xml files into place before starting Confluence. To check whether the migration has run, look in the 'users' table on the database.



Make sure your usernames match between os_users and LDAP.

1. Create a test environment. This will create an exact replica of your current version. Use your current version. This will ensure that the automatic user migration does not occur when you restart.
2. Download ldap_hibernate_osuser_atlassian-user.xml, rename it to `atlassian-user.xml` then copy to your `<INSTALL>/confluence/WEB-INF/classes` directory. It should overwrite the previous `atlassian-user.xml`. Your `atlassian-user.xml` file should have LDAP, hibernate and osuser repositories (in this order) enabled.
3. Follow Customising atlassian-user.xml
4. Restart Confluence. Login as an Administrator, and go to this URL:

```
1.<BASEURL>/admin/osuser2atluser.jsp
```

   Replace `<BASEURL>` with the URL you currently use to access Confluence. For example, http://confluence.atlassian.com or http://foobar.com/confluence.



The comment "Once the LDAP repository is configured, this migration will ignore users who have the same username as an LDAP user. This will ensure users are not duplicated in Confluence when you have both LDAP and local Confluence users enabled" is meant to describe avoiding CONF-8098. Your users, if they match in LDAP, will be migrated to the LDAP user repository on the database. See the Testing section below to confirm your results.

5. Click the link **Begin migration**. You will know the migration has been successful if you see this reported:

```
1.Migrating users ... Users migrated successfully!
2.Migrating propertyset data ... Propertyset data migrated successfully!
3.Migrating groups ... Groups migrated successfully\!
```

   If you encounter errors, please create a support ticket at http://support.atlassian.com and attach your application server logs.

6. Stop Confluence.
7. Edit `atlassian-user.xml` file and comment out the `<osuser>` repository.
   Change this line:

```
1.<osuser key="osuserRepository" name="OSUser Repository"/>
```

   to this:

```
1.<!-- <osuser key="osuserRepository" name="OSUser Repository"/> -->
```

8. Start up Confluence and check that you can login using the admin account you first set up when running through the Confluence Setup Wizard. If not, re-examine your steps and repeat from there.

Grant access to LDAP users and groups

To grant Confluence login access to your LDAP groups and users:

1. From Confluence, go to Administration > Global Permissions
2. Click to Edit Permissions for Groups
3. In the textbox to Grant Browse Permission, enter the name of an LDAP group that should have Confluence access. Click Add.
4. Tick the Can Use box for the LDAP group. If the group is not found, it was not present in your LDAP server.
5. For other LDAP groups that need access to Confluence, add them using the same method.
6. If you are integrating LDAP with Confluence for authentication only, no LDAP groups will appear in Confluence. All the individual LDAP users will have to be manually added to an internal Confluence group having with Can Use permissions enabled before they can have access to Confluence.
7. Setup your Confluence page and space permissions for these LDAP groups and users.

To setup all LDAP users as members of particular Confluence internal groups, use the LDAP Dynamic Groups Plugin.

Testing

1. Check that groups are associated by visiting a user from the User Browser and logging in.
2. Check your external_entites and external_members tables on your database. If done correctly, these tables should contain the users who were migrated from os_users and matched in LDAP. You should be able to see their group associations in the external_members table.

```
1.SELECT * FROM external_entities;
2.SELECT * FROM external_members;
```

3. Check your users table. These are the users who were *not* matched in LDAP. It might include the 'admin' user that you made when you originally created the wiki, and perhaps others who've signed up who are not in LDAP. If you need to migrate users from this table into LDAP, check the utility attached to CONF-10654.

Related Pages

- Confluence LDAP Documentation Index
- Add LDAP Integration For User Authentication Only

More information

- LDAP FAQ
- If LDAP users or groups are not displayed in Confluence, download the Paddle diagnostic tool
- List of known, unresolved LDAP bugs
- Comments on this page.

Support

Failing all else, lodge a support request. Be sure to attach your atlassian-user.xml, Paddle logs and a zip of your Confluence logs.

Add LDAP Integration For User Authentication Only

Explanation

If you'd like to configure LDAP for user management and authentication only, start by configuring LDAP integration with groups, and continue to this point to remove group management.

Applies For

- Enabling LDAP for the first time
- Upgrading existing LDAP without enabling group management

Important Points

- LDAP users will be mapped to Confluence. If the Confluence username coincides with the LDAP username, the password lookup is done against LDAP. Group management will happen in Confluence.
- Each LDAP user must be added to an internal Confluence group having Can Use permissions in order for those LDAP users to access Confluence. If a password is created for an LDAP user in Confluence, it will be ignored as the LDAP password will override it.

Instructions

If you do not wish Confluence to retrieve any of your LDAP groups and display them inside Confluence then you can do this by specifying a dummy value for the groupSearchFilter filter in your atlassian-user.xml file. That is, update your atlassian-user.xml file with the following:

```
1.<groupSearchFilter>(objectClass=dummyValue)</groupSearchFilter>
```

An example atlassian-user.xml file:

```
<baseUserNamespace>cn=users,dc=ad,dc=atlassian,dc=com</baseUserNamespace>
<baseGroupNamespace>ou=groups,dc=ad,dc=atlassian,dc=com</baseGroupNamespace>
<usernameAttribute>sAMAccountName</usernameAttribute>
<userSearchFilter>(objectClass=user)</userSearchFilter>
<firstnameAttribute>givenname</firstnameAttribute>
<surnameAttribute>sn</surnameAttribute>
<emailAttribute>mail</emailAttribute>
<groupnameAttribute>cn</groupnameAttribute>
<groupSearchFilter>(objectClass=dummyValue)</groupSearchFilter>
<membershipAttribute>member</membershipAttribute>
<userSearchAllDepths>false</userSearchAllDepths>
<groupSearchAllDepths>false</groupSearchAllDepths>
```

Please note: You will still have to provide a *valid* LDAP DN for `baseGroupNamespace`. Confluence still performs a search for groups quite frequently, so you should use a DN without many child nodes, like an individual user DN.

atlassian-user.xml reference

This page describes the function of each of the tags in an atlassian-user.xml file. These can be added as child tags of the `<ldap>` tag in your atlassian-user.xml file to configure each option.

Developer note: this information is derived from atlassian-user-defaults.xml, which can be found in the source of Atlassian-User under `src/main/resources/`. It also ships in atlassian-user.jar.

Core settings

These settings do not have a default value and must be provided to configure an LDAP connection.

XML tag	Description
LDAP connection properties	
`host`	The host name of the machine running the LDAP server. This must resolve from the machine running Confluence.
`port`	The port number that the LDAP server is running on. This should usually be 389.
`securityPrincipal`	The distinguished name (DN) of a user who is allowed to browse the entire LDAP repository. This can be omitted if the repository has anonymous access enabled.
`securityCredential`	The password for the user configured as the securityPrincipal. This can be omitted if the repository has anonymous access enabled. [1]
`baseContext`	The DN of the top of the LDAP tree that contains both users and groups.
LDAP user mapping properties	
`baseUserNamespace`	The DN at the top of the LDAP tree which contains users. For example: `ou=users,dc=example,dc=com`.
`userSearchFilter`	An LDAP search filter which matches only users under the baseUserNamespace. For example: `(objectClass=person)`. [1]
`usernameAttribute`	The attribute on a user in LDAP which contains the Confluence username. It must be unique across all users. For example: `cn` (OpenLDAP), `sAMAccountName` (AD).
`firstnameAttribute`	The attribute on a user in LDAP which contains the first name of the user. For example: `givenName`.
`surnameAttribute`	The attribute on a user in LDAP which contains the last name of the user. For example: `sn`.
`emailAttribute`	The attribute on a user in LDAP which contains the email address of the user. For example: `mail`.
LDAP group mapping properties	
`baseGroupNamespace`	The DN at the top of the LDAP tree which contains groups. For example: `ou=groups,dc=example,dc=com`.
`groupSearchFilter`	An LDAP search filter which matches only group entities under the baseGroupNamespace. For example: `(objectClass=group)`. [1]
`groupnameAttribute`	The attribute on a group in LDAP which contains the Confluence group name. It must be unique across all groups. For example: `cn`.
`membershipAttribute`	The attribute on a group in LDAP which contains the DN of each member in the group. [2] For example: `member`.

Notes

1. If these values contain ampersands, they must be escaped in the XML file. This is a common situation with LDAP search filters and passwords containing ampersands. For example, the LDAP search filter `(&(objectClass=user)(mail=*@example.com))` would be put in the XML as:
   `<userSearchFilter>(&(objectClass=user)(mail=*@example.com))</userSearchFilter>`.
2. This can also contain the user name of each member of the group. See the `useUnqualifiedUsernameForMembershipComparison` optional configuration setting below.

Optional settings

XML tag	Default value	Description
LDAP connection properties		
`securityProtocol`	plain ssl	Allow encrypted (SSL) connections. Can be omitted if anonymous access is available.
`authentication`	simple	Plain-text password transmission. Can be 'none' if anonymous access to the LDAP server is available.
`initialContextFactory`	com.sun.jndi.ldap.LdapCtxFactory	Class name of LDAP provider (default: Sun JNDI)
`batchSize`	100	Size of pages in search results
`poolingOn`	true	Use connection pooling
`connectTimeout`	30000	Timeout in milliseconds when opening new server connections. Default: 30 seconds.
`readTimeout`	60000	Timeout in milliseconds for search and other read operations. Default: 60 seconds.
LDAP connection pool properties		
`initSize`	1	Initial size of connection pool, e.g. number of connections to open at start-up. [1]
`prefSize`	10	Preferred size of connection pool. [1]
`maxSize`	0	Maximum size of connection pool. Zero means no maximum size. [1]
`timeout`	300000	Idle time in milliseconds for a connection before it is removed from the pool. Default: 5 minutes. [1]
`debugLevel`	none	Debug level for logging. [1]
`poolAuthentication`	simple	Authentication for pool connections. [1]
LDAP search properties		
`timeToLive`	0	Time limit on searches in milliseconds. Zero means no limit. [2]
`userSearchAllDepths`	false	Whether user searches should search through the LDAP tree or only for direct children of the DN specified by the userSearchFilter. [3]
`groupSearchAllDepths`	false	Whether group searches should search through the LDAP tree or only for direct children of the DN specified by the groupSearchFilter. [3]
`useUnqualifiedUsernameForMembershipComparison`	false	If set to true, Confluence will use the value of the usernameAttribute on the user to check for group membership comparisons instead of the complete distinguished name.

Notes

1. The connection pool properties provided by Atlassian-User correspond with the connection pooling properties in JNDI. See this documentation for further information.
2. More information on this time limit is available on Sun's JNDI tutorial.

3. The "searchAllDepths" parameters toggle between SearchControls.SUBTREE_SCOPE (true) and SearchControls.ONELEVEL_SCOPE (false). See Sun's JNDI tutorial on scope configuration for more information.

RELATED TOPICS

Customising atlassian-user.xml
Add LDAP Integration
LDAP User Management

Changes in osuser.xml from 1.0.3a to 1.1.x

If you have setup Confluence to:

- delegate user management to JIRA (as per our documentation here) or
- use LDAP/Active directcory authentication (as per our documentation here)

Then you will need to make the following changes to your `osuser.xml` file:

1. replace

```
1.com.atlassian.confluence.user.providers.Caching
```

   with

```
1.bucket.user.providers.Caching
```

   Doing this will effectively convert the following:

```
1.com.atlassian.confluence.user.providers.CachingCredentialsProvider
2.com.atlassian.confluence.user.providers.CachingAccessProvider
3.com.atlassian.confluence.user.providers.CachingProfileProvider
```

   to

```
1.bucket.user.providers.CachingCredentialsProvider
2.bucket.user.providers.CachingAccessProvider
3.bucket.user.providers.CachingProfileProvider
```

2. and replace

```
1.com.atlassian.confluence.user.ConfluenceHibernateConfigProvider
```

   with

```
1.bucket.user.BucketHibernateConfigProvider
```

Alternatively

You can just reconfigure the new osuser.xml with your changes.

Configuring multiple LDAP repositories

In this document::

- Prerequisites
- Configuration
- Side effects
- Cache configuration
- Two connections to the same server
- Single Sign-On Alternative
- Related pages

Prerequisites

Before reading this, ensure you understand the configuration details outlined in Customising atlassian-user.xml. This describes how to configure a single LDAP repository in Confluence, and is prerequisite knowledge for following the instructions below.

For brevity, all examples on this page are partial examples. A complete `atlassian-user.xml` LDAP configuration can be found in Customising atlassian-user.xml.

You will need at least Confluence 2.3 or higher to be able to use these instructions.

Configuration

To configure multiple LDAP repositories in Confluence, put multiple `<ldap>...</ldap>` entries into `confluence/WEB-INF/classes/atlassian-user.xml`.

The order of the entries in the file will be the order that the repositories are searched for users. That is, if a user tries to log in with the

username *jsmith*, the first repository in `atlassian-user.xml` will be searched for the user with the username *jsmith*. If no user is found in that repository, the second repository specified in `atlassian-user.xml` will be searched.

Here is a partial configuration that connects Confluence to two different LDAP servers. They are given the identifiers *ldap1* and *ldap2*, and connect to the servers `ldap-sf.example.org` and `ldap-nyc.example.org` respectively.

If you change your ldap key to a different name, you will need to change the cache name to pick it up. This is described in CONFKB181536872.

```
<atlassian-user>
        <repositories>
                <ldap key="ldap1" name="San Francisco Example Repository" cache="true">
                        <host>ldap-sf.example.org</host>
                        <port>389</port>

                        <!-- ... remainder of server configuration ... -->

                </ldap>

                <ldap key="ldap2" name="New York City Example Repository" cache="true">
                        <host>ldap-nyc.example.org</host>
                        <port>389</port>

                        <!-- ... remainder of server configuration ... -->

                </ldap>

                <hibernate key="hibernate" name="Hibernate Repository" description="Hibernate Repository" />
        </repositories>
</atlassian-user>
```

Points to note:

- each server must have a unique *key* attribute
- each server must include the full LDAP configuration, including `baseUserNamespace`, `baseGroupNamespace` and so on
- Confluence's internal repository, the `<hibernate>` repository, must be specified last
- you can include more than two LDAP repositories, but please read the *Side effects* section below.

Side effects

The main side effect of configuring multiple LDAP servers is degrading performance. There are many activities in Confluence where user or group information is retrieved:

- logging in
- user/group searches
- permission checks when viewing or editing a page.

Confluence tries to cache as much information as possible from the LDAP queries, but **almost certainly adding multiple LDAP servers will degrade the performance of the application**. This is especially true if any of the LDAP servers are geographically distant from Confluence, where any LDAP query has a significant latency (> 50 ms roundtrip).

Cache configuration

You will need to configure your `<Confluence-Home-Directory>/config/confluence-coherence-cache-config.xml` or `<Confluence-Home-Directory>/config/confluence-coherence-cache-config-clustered.xml` file to add LDAP related caches for every additional LDAP repository being added to `atlassian-user.xml`.

To do this, please add the following block of lines to respective cache configuration file, for each additional LDAP repository being configured:

```
<!-- second LDAP repository -->

                <cache-mapping>
                        <cache-name
>com.atlassian.user.impl.ldap.LDAPUserManagerReadOnly.ldapRepository2.users</
cache-name>
                        <scheme-name>user</scheme-name>
                </cache-mapping>

                <cache-mapping>
                        <cache-name
>com.atlassian.user.impl.ldap.LDAPUserManagerReadOnly.ldapRepository2.users_ro</
cache-name>
                        <scheme-name>user</scheme-name>
                </cache-mapping>
                <cache-mapping>
                        <cache-name
>com.atlassian.user.impl.ldap.LDAPUserManagerReadOnly.ldapRepository2.repository</
cache-name>
                        <scheme-name>user</scheme-name>
                </cache-mapping>

                <cache-mapping>
                        <cache-name
>com.atlassian.user.impl.ldap.LDAPGroupManagerReadOnly.ldapRepository2.groups</
cache-name>
                        <scheme-name>user</scheme-name>
                </cache-mapping>
                <cache-mapping>
                        <cache-name
>com.atlassian.user.impl.ldap.LDAPGroupManagerReadOnly.ldapRepository2.groups_hasMembership</
cache-name>
                        <scheme-name>user</scheme-name>
                </cache-mapping>
                <cache-mapping>
                        <cache-name
>com.atlassian.user.impl.ldap.LDAPGroupManagerReadOnly.ldapRepository2.groups_getGroupsForUser</
cache-name>
                        <scheme-name>user</scheme-name>
                </cache-mapping>
                <cache-mapping>
                        <cache-name
>com.atlassian.user.impl.ldap.LDAPGroupManagerReadOnly.ldapRepository2.repositories</
cache-name>
                        <scheme-name>user</scheme-name>
                </cache-mapping>

                <!-- END second LDAP Repository -->
```

If you are using a Standard Edition of Confluence, please refer to this knowledge base article for details on implementing these LDAP configurations into your `<Confluence-Home-Directory>/config/ehcache.xml` file.

Please replace the example LDAP key above `ldapRepository2`, with your relevant LDAP key, such that it matches the one defined in your `atlassian-user.xml` file.

Two connections to the same server

It also possible, but not usually recommended, for Confluence to connect twice to the same server. When connecting twice to the same server, **you must not have overlapping group or user namespaces in the LDAP tree**.

Here is an partial configuration, retrieving two separate LDAP user branches, but only one LDAP group branch. To configure only a single group branch, the group filter in the second LDAP repository searches for a non-existent value so it will not return any results. (This is generally fast as long as your LDAP server has an index on objectClass for the given tree section.)

```
<atlassian-user>
          <repositories>
                    <ldap key="ldap1" name="Example Repository, SF user tree" cache="true">
                              <host>ldap.example.org</host>
                              <port>389</port>
                              <!-- ... remainder of connection configuration ... -->

                              <!-- user search filter -->
                              <baseUserNamespace>cn=San Francisco,dc=ldap,dc=example,dc=org</
baseUserNamespace>
                              <userSearchFilter>(objectClass=user)</userSearchFilter>

                              <!-- ... remainder of user configuration ... -->

                              <!-- group search filter -->
                              <baseGroupNamespace>cn=Groups,dc=ldap,dc=example,dc=org</
baseGroupNamespace>
                              <groupSearchFilter>(objectClass=group)</groupSearchFilter>

                              <!-- ... remainder of server configuration ... -->
                    </ldap>
                    <ldap key="ldap2" name="Example Repository, NYC user tree" cache="true">
                              <host>ldap.example.org</host>
                              <port>389</port>
                              <!-- ... remainder of connection configuration ... -->

                              <!-- user search filter -->
                              <baseUserNamespace>cn=New York City,dc=ldap,dc=example,dc=org</
baseUserNamespace>
                              <userSearchFilter>(objectClass=user)</userSearchFilter>

                              <!-- ... remainder of user configuration ... -->

                              <!-- group search filter -->
                              <baseGroupNamespace>cn=Groups,dc=ldap,dc=example,dc=org</
baseGroupNamespace>
                              <groupSearchFilter>(objectClass=nothing)</groupSearchFilter>

                              <!-- ... remainder of server configuration ... -->
                    </ldap>
                    <hibernate key="hibernate" name="Hibernate Repository" description="Hibernate
Repository" />
          </repositories>
</atlassian-user>
```

Points to note:

- each repository will have its own connection pool, so **Confluence will use twice as many connections to the LDAP server**
- performance will typically be degraded, as discussed in *Side effects* above
- each server must have a unique *key* attribute
- each server must include the full LDAP configuration, including `baseUserNamespace`, `baseGroupNamespace` and so on
- Confluence's internal repository, the `<hibernate>` repository, must be specified last.

Single Sign-On Alternative

Rather than configuring multiple LDAP repositories, you're able to reduce the overhead of having multiple LDAP servers by setting up an SSO solution as an intermediate user manager. Confluence only needs to lookup the SSO tool once, and that tool then looks up both servers on behalf of Confluence. The SSO tool essentially federates your LDAP servers into a single service and also caches the results across all applications that use single sign-on.

The main advantage is that if a user is already logged into any applications that have single sign-on enabled, then their account details will already be cached by the SSO tool and can normally be served from cache. Only the first application request needs to wait for the LDAP response, with subsequent requests from that or other applications able to use the SSO cache until it expires. If your LDAP servers are high-latency, then you can also reduce latency on cache lookups by locating the SSO tool closer to your Confluence server. Atlassian Crowd is an example SSO solution that already integrates with LDAP and all Atlassian tools.

Related pages

Customising atlassian-user.xml
Add LDAP Integration

Connect to LDAP, JIRA or Other Services Via SSL

Confluence and SSL Services

To serve Confluence pages over SSL, refer to Adding SSL for Secure Logins and Page Security. This document describes connections to SSL Services. The most common are LDAP over SSL and Trusted Applications when JIRA is running over SSL.

Importing SSL Certificates

> The following commands are in reference to JDK 1.5. For commands/syntax relevant to JDK 1.6, please refer to this document.

Atlassian User LDAP supports connecting to an LDAP server over SSL/HTTPS. The Trusted Communications protocol requires configuration if JIRA is running over SSL.

1. Add the root certificate to your default Java keystore with the following command. This is the certificate that was used to authorise the LDAP server's certificate; it will be either the one that was used for signing it, or will come from further up in the trust chain, possibly the root certificate. This is often a self-signed certificate, when both ends of the SSL connection are within the same network. Again, the exact alias is not important.

```
keytool -import -alias serverCert -file RootCert.crt -keystore
%JAVA_HOME%/jre/lib/security/cacerts (Windows)
keytool -import -alias serverCert -file RootCert.crt -keystore
$JAVA_HOME/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, and you can use any alias of your choosing in place of "JIRAorLDAPServer.crt".

```
keytool -import -alias ldapCert -file JIRAorLDAPServer.crt -keystore
%JAVA_HOME%/jre/lib/security/cacerts (Windows)
keytool -import -alias ldapCert -file JIRAorLDAPServer.crt -keystore
$JAVA_HOME/jre/lib/security/cacerts (Linux/Unix/Mac)
```

3. Edit the file in your Confluence install directory, `confluence\WEB-INF\classes\atlassian-user.xml`. Change the value of `securityProtocol` from "plain" to "ssl":

```
<securityProtocol>ssl</securityProtocol>
```

   Switch the LDAP connection to the SSL port, if it's different from the default LDAP port. If you're using the most common LDAPS port, set

```
<port>636</port>
```

   The keytool will ask you for a password. The default password is "changeit" without the quotes.

4. 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)
```

5. Ensure that you've updated JAVA_OPTS to specify the path to the keystore, as specified in Connecting to SSL services, before restarting tomcat/Confluence.
   There's 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 Unable to Connect to SSL Services due to PKIX Path Building Failed sun.security.provider.certpath.SunCertPathBuilderException or the more general SSL Troubleshooting articles for troubleshooting information.

RELATED ARTICLES

JIRA Connecting to SSL Services
Confluence Unable to Connect to SSL Services
Configure Web Proxy Support for Confluence
Adding SSL for Secure Logins and Page Security
Troubleshooting SSL

Disabling the Built-In User Management

You only need to follow the instructions on this page if you are using **JIRA** for user management. It disables all the group and user management screens in Confluence.

You can choose to enable this setting for **Crowd** or **LDAP** user management, if you are happy to manage users and groups outside Confluence.

Enabling the "external user management" setting will make user and group management screens **read-only** within Confluence.

This will also prevent users from signing up to the site, as well as editing their name, email and password particulars from within Confluence.

You need to have System Administrator permissions in order to perform this function.

To disable management of users and groups within Confluence,

1. Go to the Confluence '**Administration Console**'. To do this:
    - Open the '**Browse**' menu and select '**Confluence Admin**'. The 'Administration Console' view will open.
2. Select '**General Configuration**' in the left-hand panel.
3. In the '**General Configuration**' screen, click '**Edit**'.
4. Select '**On**' beside '**External User Management**'.
5. Click '**Save**'.

Confluence installations using Crowd for user and group management:
If Crowd's directory permissions are configured so that Confluence cannot update the Crowd directories, then Confluence's external user management setting *must* be turned on. Otherwise, a 'System Error' will occur when Confluence attempts to write data into Crowd. For more information about integrating Crowd with Confluence, please refer to Integrating Crowd with Atlassian Confluence.

RELATED TOPICS

Changes in osuser.xml from 1.0.3a to 1.1.x

Add LDAP Integration For User Authentication Only

Disabling the Built-In User Management

Administrators Guide Home Confluence Documentation Home

Legacy User Management Documentation

This section contains legacy user management documentation.

- LDAP Authentication with OSUser

LDAP Authentication with OSUser

Deprecated for Confluence 2.7.0 and later
For Confluence 2.7.0 and above, all user management is performed by AtlassianUser as described in Understanding User Management in Confluence. OSUser user management, described in this document, has been deprecated.

For up-to-date instructions on configuring LDAP authentication, see Add LDAP Integration.

Overview

To configure Confluence to connect to LDAP for user management only, you have two options:

1. Use LDAP for users and authentication, not groups. That's described in Add LDAP integration, and includes an optional step to remove group management. This method is supported in Confluence versions after version 2.7.

   or

2. Create and manage users and groups in Confluence, and use LDAP for authentication only. This is the deprecated OSUser management approach, as described in this document. If a username exists in both Confluence and LDAP, the user uses their LDAP password to log in. This method has been deprecated after Confluence version 2.7

Important Points about Option 2 Above

- Only the password lookup is done against LDAP and *only* if the Confluence username is the same as the LDAP username. Users and user profiles are still managed in Confluence. See technical explanation below.
- A Confluence account must be created for each LDAP user, as users do not automatically have access. For an LDAP user to access Confluence, a site administrator will still need to create an account for them. The password in this Confluence account will be ignored as the LDAP password will override it.

Instructions

Step 1: Modify atlassian-user.xml to use OSUser Management

For Confluence 2.7 and above, the default user repository is the hibernate repository (Atlassian User). To revert to OSUser, you will need to put the OSUser repository tag in the top position so it's the primary user management option. Modify /confluence/WEB-INF/classes/atlassian-user.xml to contain this:

```
1.<atlassian-user>
2.<repositories>
3.<osuser key="osuserRepository" name="OSUser Repository"/>
4.<hibernate name="Hibernate Repository" key="hibernateRepository" description="Hibernate
Repository" cache="true"/>
5.</repositories>
6.</atlassian-user>
```

Note: For Confluence version prior to 2.7, if you have delegated your user management to JIRA, LDAP or any other external user management system, copy the following files from your old Confluence installation to your new Confluence installation:

- <Installation-Directory>/confluence/WEB-INF/classes/osuser.xml.
- <Installation-Directory>/confluence/WEB-INF/classes/atlassian-user.xml (if you are upgrading from Confluence 2.2 or later).

Step 2: Open the osuser.xml file located in your home directory under WEB-INF/classes

In the osuser.xml file, the **CredentialsProviders** are responsible for authenticating passwords. The default CachingCredentialsProvider looks in the Confluence database. To enable LDAP aunthentication, you will need to add a **LDAPCredentialsProvider**, so that LDAP users can also be authenticated:

Here's what the default osuser.xml contains:

```
01.<provider class="bucket.user.providers.CachingCredentialsProvider">
02.<property name="chain.classname"
>com.opensymphony.user.provider.hibernate.HibernateCredentialsProvider</property>
03.<property name="chain.configuration.provider.class">bucket.user.BucketHibernateConfigProvider</
property>
04.</provider>
05.<provider class="bucket.user.providers.CachingAccessProvider">
06.<property name="chain.classname"
>com.opensymphony.user.provider.hibernate.HibernateAccessProvider</property>
07.<property name="chain.configuration.provider.class">bucket.user.BucketHibernateConfigProvider</
property>
08.</provider>
09.<provider class="bucket.user.providers.CachingProfileProvider">
10.<property name="chain.classname"
>com.opensymphony.user.provider.hibernate.HibernateProfileProvider</property>
11.<property name="chain.configuration.provider.class">bucket.user.BucketHibernateConfigProvider</
property>
12.</provider>
```

Step 3: Edit the osuser.xml file as shown below

For Confluence version 2.1 and later:

```
01.<provider class="com.atlassian.confluence.user.ConfluenceLDAPCredentialsProvider">
02.<property name="java.naming.factory.initial">com.sun.jndi.ldap.LdapCtxFactory</property>
03.<property name="java.naming.provider.url">ldap://localhost:389</property>
04.<property name="searchBase">dc=atlassian,dc=com</property>
05.<property name="uidSearchName">cn</property>
06.<!--
07.<property name="java.naming.security.principal">cn=Manager,dc=atlassian,dc=com</property>
08.<property name="java.naming.security.credentials">secret</property>
09.<property name="exclusive-access">true</property>
10.-->
11.</provider>
12.
13.<provider class="bucket.user.providers.CachingCredentialsProvider">
14.<property name="chain.classname"
>com.opensymphony.user.provider.hibernate.HibernateCredentialsProvider</property>
15.<property name="chain.configuration.provider.class">bucket.user.BucketHibernateConfigProvider</
property>
16.</provider>
17.<provider class="bucket.user.providers.CachingAccessProvider">
18.<property name="chain.classname"
>com.opensymphony.user.provider.hibernate.HibernateAccessProvider</property>
19.<property name="chain.configuration.provider.class">bucket.user.BucketHibernateConfigProvider</
property>
20.</provider>
21.<provider class="bucket.user.providers.CachingProfileProvider">
22.<property name="chain.classname"
>com.opensymphony.user.provider.hibernate.HibernateProfileProvider</property>
23.<property name="chain.configuration.provider.class">bucket.user.BucketHibernateConfigProvider</
property>
24.</provider>
```

For older verisons of Confluence

```
01.<provider class="com.opensymphony.user.provider.ldap.LDAPCredentialsProvider">
02.<property name="java.naming.factory.initial">com.sun.jndi.ldap.LdapCtxFactory</property>
03.<property name="java.naming.provider.url">ldap://localhost:389</property>
04.<property name="searchBase">dc=atlassian,dc=com</property>
05.<property name="uidSearchName">cn</property>
06.<!--
07.<property name="java.naming.security.principal">cn=Manager,dc=atlassian,dc=com</property>
08.<property name="java.naming.security.credentials">secret</property>
09.<property name="exclusive-access">true</property>
10.-->
11.</provider>
12.
13.<provider class="bucket.user.providers.CachingCredentialsProvider">
14.<property name="chain.classname"
>com.opensymphony.user.provider.hibernate.HibernateCredentialsProvider</property>
15.<property name="chain.configuration.provider.class">bucket.user.BucketHibernateConfigProvider</
property>
16.</provider>
17.<provider class="bucket.user.providers.CachingAccessProvider">
18.<property name="chain.classname"
>com.opensymphony.user.provider.hibernate.HibernateAccessProvider</property>
19.<property name="chain.configuration.provider.class">bucket.user.BucketHibernateConfigProvider</
property>
20.</provider>
21.<provider class="bucket.user.providers.CachingProfileProvider">
22.<property name="chain.classname"
>com.opensymphony.user.provider.hibernate.HibernateProfileProvider</property>
23.<property name="chain.configuration.provider.class">bucket.user.BucketHibernateConfigProvider</
property>
24.</provider>
```

- **Update the following properties to suit your LDAP server:**
    - **url** (currently set to **ldap://localhost:389**)
    - **searchBase** (currently set to **dc=atlassian,dc=com**)
    - **uidSearchName** (currently set to **cn**)
- **If your LDAP server is not configured to allow anonymous lookups, you need to:**
    - remove comment tags
    - enter the username; including searchBase (currently set to *cn=Manager,dc=atlassian,dc=com*)
    - enter password (currently set to *secret*)

The Credentials (password) checking is a separate operation from user-profile lookups. The profile can be loaded from the Confluence database, but the password is looked up from LDAP. Furthermore, multiple credentials providers can be specified (here, LDAP and OSUser), and if one fails, the other will be used. This allows non-LDAP users to log in with their Confluence password.

How this works

It is useful to have a general idea of how this setup works. This section outlines some consequences of this OSUser implementation and provides some help for people experiencing LDAP connection problems.

Only password-checking for LDAP users is done in Confluence

User profiles are still managed in Confluence (by the CachingProfileProvider in osuser.xml). Only the password lookup is performed against LDAP and only if the Confluence username coincides with a LDAP username. This is because Credentials (password) checking is a separate operation to user-profile lookups. The profile can be loaded from the Confluence database, but the password is looked up from LDAP.

Not all LDAP users have Confluence access

Another effect of this implementation is that LDAP users do not automatically have access to Confluence. A Confluence account must be created for each user wishing to use Confluence.

This is because each Confluence user has a set of groups (for example, 'confluence-users') stored in their profile. Without an associated group, that user can do nothing; not even browse Confluence (that is, they lack the 'use' permission).

Thus, for an LDAP user to use Confluence, a Confluence admin must create an account for them and assign them to a group (typically 'confluence-user'). The password in this Confluence account will be ignored, as the LDAP password will override it.

RELATED TOPICS

Changes in osuser.xml from 1.0.3a to 1.1.x

Add LDAP Integration For User Authentication Only

Disabling the Built-In User Management

Confluence Documentation Home

Troubleshooting LDAP User Management

Confluence supplies an LDAP connectivity tool (also called '**Paddle**') that will test the LDAP settings in your `atlassian-user.xml` file. This will help you to diagnose problems with LDAP user management, such as:

- LDAP server not responding.
- LDAP settings incorrectly configured in atlassian-user.
- Other issues reported from the LDAP queries run by the testing tool.

LDAP is the protocol used for user management by Active Directory and other LDAP directories.

> **This document applies to Confluence 2.8 and later**
> From Confluence version 2.8, the LDAP User Test tool is integrated into the Administration Console. In previous versions of Confluence, this tool was available as a separate utility called '**Paddle**'. If you are running an earlier version of Confluence, please refer to the Confluence 2.7 documentation or choose the relevant Confluence version from the list of previous versions on the documentation home page.

On this page:

- Running the LDAP Connectivity Tool via the Administration Console
- Running the LDAP Connectivity Test Tool outside Confluence
  - Steps in Detail
  - Parameters
- Sample Output
  - Output from a Successful Test
  - Output showing Failure to Connect
  - Output showing No LDAP Configuration
    - RELATED TOPICS

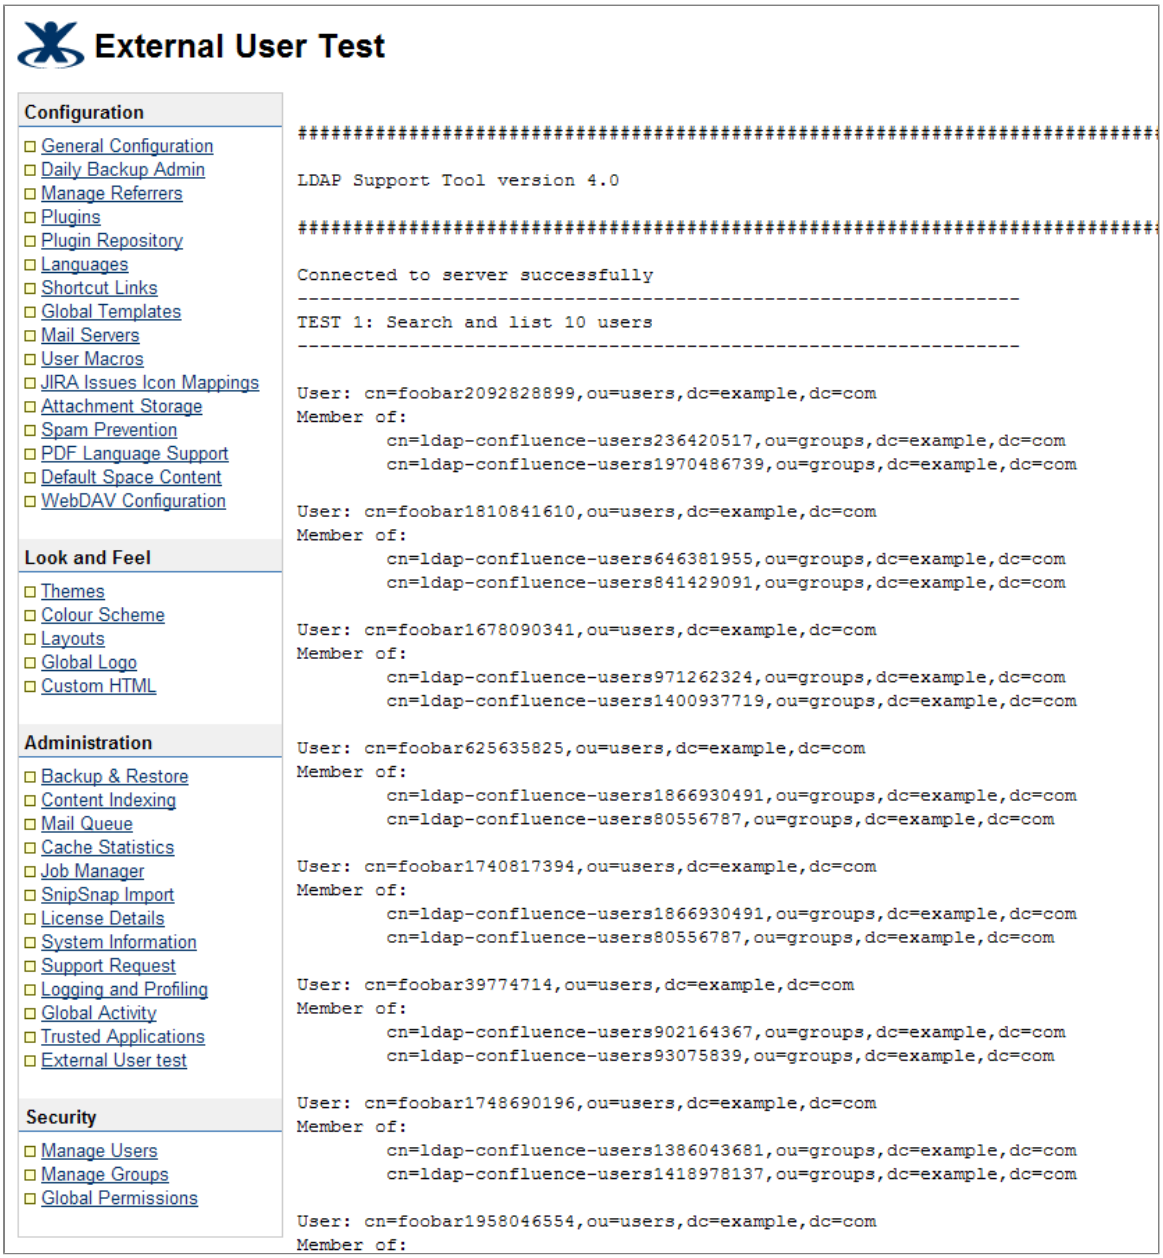
Running the LDAP Connectivity Tool via the Administration Console

Note: You will need to restart your Confluence server each time you make changes to your `atlassian-user.xml` file. To avoid this, consider running the tool outside of Confluence.

To test your LDAP connection settings,

1. First verify that your atlassian-user.xml contains at least one LDAP repository with a key of "ldapRepository".
2. Go to the Confluence '**Administration Console**'. To do this:
    - Open the '**Browse**' menu and select '**Confluence Admin**'. The 'Administration Console' view will open.
3. Select '**LDAP Connectivity Test**' in the '**Administration**' section of the left-hand panel.
   The 'LDAP Connectivity Test' will only appear if your atlassian-user.xml contains an ldap repository with a key of "ldapRepository", as described in the preceding step.
4. The results of the test will appear in the main panel of the screen. Refer to the screenshot below for an example, and to the output samples lower down on this page.

Screenshot: External User Test



Running the LDAP Connectivity Test Tool outside Confluence

You do not need to have Confluence running to run this tool. Instead, you can download this jar and perform the following steps.

Steps in Detail

1. Install the Java SDK from Sun. See Setting the JAVA_HOME Variable in Windows for tips on a Windows installation. On a Mac, Java comes pre-installed.
2. Copy the Paddle jar file into a directory where you have permission to create files.
3. Copy your `atlassian-user.xml` file into the same directory. You will find this file at the following location in your Confluence Installation Directory: `<Confluence-Installation>/confluence/WEB-INF/classes/atlassian-user.xml`. If your instance is hosted, you can download one here to supply to our hosted team.
4. Configure atlassian-user.xml according to Customising atlassian-user.xml.
5. CD into this directory from a command prompt
    a. Mac users: Open Terminal from `Applications >> Utilities >> Terminal`. Use the 'cd' command to change to the proper directory.
    b. Windows users: Open a command prompt from `Start >> Run >> cmd`.
6. Run `java -jar paddle-6.jar`.

The output will appear both in the command console window and in a paddle output log file that gets written to the local directory.

Parameters

Paddle supports the following parameters:

Name	Example	Purpose
debug	`java -jar paddle-x.x.jar debug`	Prints DEBUG messages to the console as well as paddle.log.
limit	`java -jar paddle-x.x.jar limit=100`	Sets the limit on the number of results returned by user and group queries. Defaults to 10.

Sample Output

Below are some examples of the output from the user test tool. The output is the same, whether you run the tool via the Administration Console or outside Confluence.

Output from a Successful Test

This is an example of a successful run:

```
001.
################################################################################################################
002.
003.LDAP Support Tool version 4.0
004.
005.
################################################################################################################
006.
007.Connected to server successfully
008.----------------------------------------------------------------
009.TEST 1: Search and list 10 users
010.----------------------------------------------------------------
011.
012.User: cn=foobar2092828899,ou=users,dc=example,dc=com
013.Member of:
014.          cn=ldap-confluence-users236420517,ou=groups,dc=example,dc=com
015.          cn=ldap-confluence-users1970486739,ou=groups,dc=example,dc=com
016.
017.User: cn=foobar1810841610,ou=users,dc=example,dc=com
018.Member of:
019.          cn=ldap-confluence-users646381955,ou=groups,dc=example,dc=com
020.          cn=ldap-confluence-users841429091,ou=groups,dc=example,dc=com
021.
022.User: cn=foobar1678090341,ou=users,dc=example,dc=com
023.Member of:
024.          cn=ldap-confluence-users971262324,ou=groups,dc=example,dc=com
025.          cn=ldap-confluence-users1400937719,ou=groups,dc=example,dc=com
026.
027.User: cn=foobar625635825,ou=users,dc=example,dc=com
028.Member of:
029.          cn=ldap-confluence-users1866930491,ou=groups,dc=example,dc=com
030.          cn=ldap-confluence-users80556787,ou=groups,dc=example,dc=com
031.
032.User: cn=foobar1740817394,ou=users,dc=example,dc=com
033.Member of:
034.          cn=ldap-confluence-users1866930491,ou=groups,dc=example,dc=com
035.          cn=ldap-confluence-users80556787,ou=groups,dc=example,dc=com
036.
037.User: cn=foobar39774714,ou=users,dc=example,dc=com
038.Member of:
039.          cn=ldap-confluence-users902164367,ou=groups,dc=example,dc=com
040.          cn=ldap-confluence-users93075839,ou=groups,dc=example,dc=com
041.
```

```
042.User: cn=foobar1748690196,ou=users,dc=example,dc=com
043.Member of:
044.          cn=ldap-confluence-users1386043681,ou=groups,dc=example,dc=com
045.          cn=ldap-confluence-users1418978137,ou=groups,dc=example,dc=com
046.
047.User: cn=foobar1958046554,ou=users,dc=example,dc=com
048.Member of:
049.          cn=ldap-confluence-users78102241,ou=groups,dc=example,dc=com
050.          cn=ldap-confluence-users220323310,ou=groups,dc=example,dc=com
051.
052.User: cn=foobar1263319669,ou=users,dc=example,dc=com
053.Member of:
054.          cn=ldap-confluence-users283478949,ou=groups,dc=example,dc=com
055.          cn=ldap-confluence-users1866930491,ou=groups,dc=example,dc=com
056.
057.User: cn=foobar1541805698,ou=users,dc=example,dc=com
058.Member of:
059.          cn=ldap-confluence-users1736591783,ou=groups,dc=example,dc=com
060.          cn=ldap-confluence-users91103896,ou=groups,dc=example,dc=com
061.
062.----------------------------------------------------------------
063.TEST 2: Search and list 10 groups
064.----------------------------------------------------------------
065.
066.Group: cn=ldap-confluence-users16368779,ou=groups,dc=example,dc=com
067.Members:
068.          cn=foobar1588244344,ou=users,dc=example,dc=com
069.          cn=foobar997638232,ou=users,dc=example,dc=com
070.          cn=foobar1474202427,ou=users,dc=example,dc=com
071.          cn=foobar1371091481,ou=users,dc=example,dc=com
072.          cn=foobar794627990,ou=users,dc=example,dc=com
073.          cn=foobar904159761,ou=users,dc=example,dc=com
074.          cn=foobar1992607123,ou=users,dc=example,dc=com
075.          cn=foobar839977082,ou=users,dc=example,dc=com
076.          cn=foobar725629285,ou=users,dc=example,dc=com
077.          cn=foobar1317948366,ou=users,dc=example,dc=com
078.
079.Group: cn=ldap-confluence-users1721354031,ou=groups,dc=example,dc=com
080.Members:
081.          cn=foobar991592891,ou=users,dc=example,dc=com
082.          cn=foobar1710975716,ou=users,dc=example,dc=com
083.          cn=foobar1505673129,ou=users,dc=example,dc=com
084.          cn=foobar1924966176,ou=users,dc=example,dc=com
085.          cn=foobar799226338,ou=users,dc=example,dc=com
086.          cn=foobar1999491022,ou=users,dc=example,dc=com
087.          cn=foobar14075397,ou=users,dc=example,dc=com
088.          cn=foobar539697111,ou=users,dc=example,dc=com
089.          cn=foobar112108820,ou=users,dc=example,dc=com
090.          cn=foobar1369772211,ou=users,dc=example,dc=com
091.
092.Group: cn=ldap-confluence-users1418978137,ou=groups,dc=example,dc=com
093.Members:
094.          cn=foobar1588906497,ou=users,dc=example,dc=com
095.          cn=foobar713515675,ou=users,dc=example,dc=com
096.          cn=foobar111198746,ou=users,dc=example,dc=com
097.          cn=foobar1413209068,ou=users,dc=example,dc=com
098.          cn=foobar1748690196,ou=users,dc=example,dc=com
099.          cn=foobar1650270608,ou=users,dc=example,dc=com
100.          cn=foobar1083182874,ou=users,dc=example,dc=com
101.          cn=foobar397077941,ou=users,dc=example,dc=com
102.          cn=foobar982743351,ou=users,dc=example,dc=com
103.          cn=foobar1305199694,ou=users,dc=example,dc=com
104.
105.Group: cn=ldap-confluence-users1970486739,ou=groups,dc=example,dc=com
106.Members:
107.          cn=foobar547480036,ou=users,dc=example,dc=com
108.          cn=foobar2092828899,ou=users,dc=example,dc=com
109.          cn=foobar1150791207,ou=users,dc=example,dc=com
110.          cn=foobar1648997840,ou=users,dc=example,dc=com
111.          cn=foobar632840294,ou=users,dc=example,dc=com
112.          cn=foobar1688785372,ou=users,dc=example,dc=com
113.          cn=foobar1402211544,ou=users,dc=example,dc=com
114.          cn=foobar1360676914,ou=users,dc=example,dc=com
115.          cn=foobar1900838417,ou=users,dc=example,dc=com
116.          cn=foobar174731773,ou=users,dc=example,dc=com
117.
118.Group: cn=ldap-confluence-users80556787,ou=groups,dc=example,dc=com
119.Members:
120.          cn=foobar763847686,ou=users,dc=example,dc=com
121.          cn=foobar532641707,ou=users,dc=example,dc=com
122.          cn=foobar1478107215,ou=users,dc=example,dc=com
123.          cn=foobar1457615713,ou=users,dc=example,dc=com
```

```
124.          cn=foobar1740817394,ou=users,dc=example,dc=com
125.          cn=foobar477128147,ou=users,dc=example,dc=com
126.          cn=foobar1686683086,ou=users,dc=example,dc=com
127.          cn=foobar625635825,ou=users,dc=example,dc=com
128.          cn=foobar122705808,ou=users,dc=example,dc=com
129.          cn=foobar827130393,ou=users,dc=example,dc=com
130.
131.Group: cn=ldap-confluence-users93075839,ou=groups,dc=example,dc=com
132.Members:
133.          cn=foobar1197670610,ou=users,dc=example,dc=com
134.          cn=foobar1478107215,ou=users,dc=example,dc=com
135.          cn=foobar477128147,ou=users,dc=example,dc=com
136.          cn=foobar122705808,ou=users,dc=example,dc=com
137.          cn=foobar847443321,ou=users,dc=example,dc=com
138.          cn=foobar39774714,ou=users,dc=example,dc=com
139.          cn=foobar797719296,ou=users,dc=example,dc=com
140.          cn=foobar1695434944,ou=users,dc=example,dc=com
141.          cn=foobar1251231239,ou=users,dc=example,dc=com
142.          cn=foobar2524660,ou=users,dc=example,dc=com
143.
144.Group: cn=ldap-confluence-users91103896,ou=groups,dc=example,dc=com
145.Members:
146.          cn=foobar991592891,ou=users,dc=example,dc=com
147.          cn=foobar1688694241,ou=users,dc=example,dc=com
148.          cn=foobar1924966176,ou=users,dc=example,dc=com
149.          cn=foobar1070601909,ou=users,dc=example,dc=com
150.          cn=foobar95317734,ou=users,dc=example,dc=com
151.          cn=foobar1999491022,ou=users,dc=example,dc=com
152.          cn=foobar62701314,ou=users,dc=example,dc=com
153.          cn=foobar14075397,ou=users,dc=example,dc=com
154.          cn=foobar1541805698,ou=users,dc=example,dc=com
155.          cn=foobar369641998,ou=users,dc=example,dc=com
156.
157.Group: cn=ldap-confluence-users78102241,ou=groups,dc=example,dc=com
158.Members:
159.          cn=foobar997638232,ou=users,dc=example,dc=com
160.          cn=foobar1588244344,ou=users,dc=example,dc=com
161.          cn=foobar1783755629,ou=users,dc=example,dc=com
162.          cn=foobar1371091481,ou=users,dc=example,dc=com
163.          cn=foobar904159761,ou=users,dc=example,dc=com
164.          cn=foobar794627990,ou=users,dc=example,dc=com
165.          cn=foobar104918275,ou=users,dc=example,dc=com
166.          cn=foobar839977082,ou=users,dc=example,dc=com
167.          cn=foobar725629285,ou=users,dc=example,dc=com
168.          cn=foobar1958046554,ou=users,dc=example,dc=com
169.
170.Group: cn=ldap-confluence-users236420517,ou=groups,dc=example,dc=com
171.Members:
172.          cn=foobar799557307,ou=users,dc=example,dc=com
173.          cn=foobar1258451414,ou=users,dc=example,dc=com
174.          cn=foobar1041699029,ou=users,dc=example,dc=com
175.          cn=foobar2092828899,ou=users,dc=example,dc=com
176.          cn=foobar547480036,ou=users,dc=example,dc=com
177.          cn=foobar2111497410,ou=users,dc=example,dc=com
178.          cn=foobar1648997840,ou=users,dc=example,dc=com
179.          cn=foobar1776604828,ou=users,dc=example,dc=com
180.          cn=foobar1121496402,ou=users,dc=example,dc=com
181.          cn=foobar1954769270,ou=users,dc=example,dc=com
182.
183.Group: cn=ldap-confluence-users283478949,ou=groups,dc=example,dc=com
184.Members:
185.          cn=foobar1588906497,ou=users,dc=example,dc=com
186.          cn=foobar111198746,ou=users,dc=example,dc=com
187.          cn=foobar357608135,ou=users,dc=example,dc=com
188.          cn=foobar1413209068,ou=users,dc=example,dc=com
189.          cn=foobar1263319669,ou=users,dc=example,dc=com
190.          cn=foobar1650270608,ou=users,dc=example,dc=com
191.          cn=foobar397077941,ou=users,dc=example,dc=com
192.          cn=foobar2116500618,ou=users,dc=example,dc=com
193.          cn=foobar1668222621,ou=users,dc=example,dc=com
194.          cn=foobar1032804410,ou=users,dc=example,dc=com
```

Output showing Failure to Connect

This is an example of the output when the LDAP server is not available:

```
01.
################################################################################################################
02.
03.LDAP Support Tool version 4.0
04.
05.
################################################################################################################
06.
07.Failed to connect to LDAP server: Failed to connect to LDAP server with provider URL ldap:
//ldap.example.com:389: ldap.example.com:389
08.com.atlassian.paddle.connection.ConnectionException: Failed to connect to LDAP server with
provider URL ldap://ldap.example.com:389: ldap.example.com:389
09.          at
com.atlassian.paddle.connection.DefaultConnectionFactory.createFriendlyLdapException(DefaultConnectio
55)
10.          at
com.atlassian.paddle.connection.DefaultConnectionFactory.createConnection(DefaultConnectionFactory.ja
35)
11.          at com.atlassian.paddle.task.DefaultTaskRunner.runTask(DefaultTaskRunner.java:33)
12.          at com.atlassian.paddle.Paddle.testConnectivity(Paddle.java:98)
13..
14..
15..
16.<<<snip>>>
```

Output showing No LDAP Configuration

This is an example of the output when the `atlassian-user.xml` file is not correctly configured for LDAP user management:

```
01.
################################################################################################################
02.
03.LDAP Support Tool version 4.0
04.
05.
################################################################################################################
06.
07.Error reading atlassian-user.xml file: No LDAP settings found in XML configuration.
08.com.atlassian.paddle.configuration.ConfigurationException: No LDAP settings found in XML
configuration.
09.          at
com.atlassian.paddle.configuration.AtlassianUserConfiguration.setAtlassianUserXml(AtlassianUserConfig
44)
10.          at
com.atlassian.paddle.configuration.AtlassianUserConfiguration.(AtlassianUserConfiguration.java:
21)
11.          at com.atlassian.paddle.Paddle.testConnectivity(Paddle.java:79)
12.          at
com.atlassian.confluence.admin.actions.ExternalUserConnectivityAction.execute(ExternalUserConnectivit
23)
13..
14..
15..
16.<<<snip>>>
```

RELATED TOPICS

Requesting External User Management Support
LDAP User Management
User Management

Troubleshooting the "Not Permitted" Screen under LDAP Integration

If you have set up Confluence with AtlassianUser integration and attempted a login with an LDAP/AD account and got a page titled **"Not Permitted"** here are the steps to troubleshoot this:

- Have you assiged USE permission to the relevant LDAP groups from the Administration > Global Permissons page?

Before an LDAP user can login and use Confluence, the LDAP group they belong to must be given USE permission directly. (Please note: nested groups is not supported in Confluence as yet, so you cannot specify the parent group. If you want this feature, please vote for it here.)

- Does an account exist on your LDAP/AD server that has the same name as your local admin account? (for example, is there an acccount on LDAP called 'admin'?)

If so, then you will not be able to login with your local admin account once you enable LDAP integration. To rectify this, you need to either rename your LDAP admin account or rollback your LDAP integration and create another Confluence admin account.

- You have assigned USE permission to the relevant LDAP groups, but LDAP users in those groups still get "Not Permitted"?

Here, you need to check if Confluence is actually aware that your LDAP users belong to those LDAP groups. To work this out, here is what we need from you:

1. Login as the local admin account you created when you first set up Confluence.
2. Enable profiling by appending `?profile=on` to the end of a Confluence URL (say the URL of the dashboard, if you happen to be on that page) and hit enter to refresh the page. This setting will now be enabled and cause additional information to be written to your log files to help us diagnose the problem.
3. Now browse to the `Administration > Manage Users` screen and do a search for any LDAP user.
4. Now click on the user to view their details (including the groups they belong to).
5. Now submit a support ticket at http://support.atlassian.com and attach:
    - A screenshot of the user details page.
    - Your server logs files.

If you are feeling brave, you can attempt to decipher the logs yourself. Here's how:

1. Your logs should display something similar to this:

```
1.[0ms] -
com.atlassian.user.impl.ldap.adaptor.LDAPStaticGroupAdaptor_search((&(objectClass=groupOfNames
2.      [0ms] - com.atlassian.user.impl.ldap.repository.DefaultLDAPRepository_getLDAPContext
3.      [0ms] -
com.atlassian.user.impl.ldap.adaptor.LDAPStaticGroupAdaptor_search_JNDI_RAW_((&(objectClass=gr
```

   This means that Confluence is using this LDAP search filter **(&(objectClass=groupOfNames)(member=cn=confadmin,ou=users,ou=people,ou=functest,dc=atlassian,dc=com))** to find the groups that the user 'confadmin' belongs to. Obviously, the objectClass and member attributes may differ in your install, but the filter should be similar.

2. Connect to your LDAP/AD server with JXplorer or an LDAP tool of your choice, and issue the above filter and check that you get the results you expect. This should help you to identify if and why the filter is incorrect and what it should be. Please add what you find in this step to the support ticket if you are unable to resolve it from here.

Migrating users from Confluence to JIRA

There is currently no way to delegate user management from JIRA to Confluence. So, if you are in a situation where your users are defined in Confluence and would like to take advantage of Confluence's ability to use JIRA user management, you will need to transfer all of your existing Confluence users into JIRA. You can do this manually, or if you have a large number of users, you can use the attached XML-RPC script.

You should backup your JIRA installation before running this script. This is an experimental script that has not gone through the same extensive testing as the Confluence and JIRA products.

Getting the migration tool

- Download the attached rpc-tools-0.9.zip.
- Extract the ZIP file to a temporary directory on your computer.

Running the migration tool

- Back up your JIRA database.
- If you do not have an 'admin' username with password 'admin' in both JIRA and Confluence, create it now.
- Ensure JIRA and Confluence have remote API access enabled. In both applications, it is configured in the *General Configuration* screen in Administration.
- Edit `connection.properties` in the extracted ZIP file to the set the XML-RPC URLs for JIRA and Confluence. The XML-RPC URLs in the default file correspond to the following application base URLs:
    - JIRA - http://localhost:8080
    - Confluence - http://localhost:8080/confluence
- Run `java -jar tools.jar` in the extracted ZIP file directory.

A successful run will generate output like the following:

```
$ java -jar tools.jar
- Transferring group: confluence-administrators
- Transferring group: confluence-users
- Transferring user: testuser1
- Transferring user: testuser2
```

Things to note:

- This script requires that both the Confluence and JIRA remote APIs are available and accessible to a username 'admin', password 'admin'. You can temporarily add this user to both system to run the script, then remove it afterwards.
- A random password will be assigned to each user that is transferred because it is not possible to access password information via the XML-RPC API. Therefore they will have retrieve a new password via the password reminder.

- JIRA does not allow users to sign up if they do not have a valid email address. Therefore we will assign them an email address of type username@example.com to any Confluence users that do not have a email address. This will allow you to find the users and help them to create a password or change the email address.

Trouble running the script?

- Ensure you have created a user 'admin' with password 'admin' in both Confluence and JIRA.
- Ensure both applications have remote API access enabled (see above).
- Ensure you have patched the Confluence RPC plugin if running Confluence 2.0.x.

If you're still having trouble, please raise a support request, and include a copy of the error you're getting.

Requesting External User Management Support

- LDAP Troubleshooting Resources
- Problems During Initial Setup
- Complex Authentication or Performance Problems
  - Confluence server
  - Confluence configuration files
  - User management system
  - Using Active Directory for LDAP?
  - Diagnostics

LDAP Troubleshooting Resources

- Add LDAP Integration
- Troubleshooting User Management and Login Issues

If the above resources don't help, continue below.

Problems During Initial Setup

Open a Support Ticket and include:

- Refer to Troubleshooting LDAP User Management. Run the tests and paste the output in your support ticket.
- Download an LDAP browser to make sure you've got the right values. Atlassian recommends LDAP Studio. Include screenshots of your user and group DN's.
- Attach your `atlassian-user.xml` file.

Complex Authentication or Performance Problems

Open a Support Ticket and include:

Confluence server

- Take a screenshot of Confluence's `Administration  System Information` (or save the page as HTML)
- Take a screenshot of Confluence's `Administration  Global Permissions`, if you are having problems with logging in
- Take a screenshot of the Space permissions page, if you are having problems with space or page permissions.

Confluence configuration files

- Attach a copy of `atlassian-user.xml`, found in confluence/WEB-INF/classes
- Attach a copy of `osuser.xml`, found in confluence/WEB-INF/classes.
- If you have implemented a custom authenticator or in any way modified `seraph-config.xml` or `seraph-paths.xml`, please provide the modified files as well.

User management system

- What is the name and version of your LDAP server?
- Does your LDAP server use dynamic or static groups?

Using Active Directory for LDAP?

Please include LDAP Studio Entry Editor snaphots with the information specified on this page

Diagnostics

- Enable profiling (as described here)
- Enable detailed user management logging by editing `confluence/WEB-INF/classes/log4j.properties`:

Change this section:

```
01.###
02.# Atlassian User
03.###
04.#log4j.logger.com.atlassian.user=DEBUG, confluencelog
05.#log4j.additivity.com.atlassian.user=false
06.#log4j.logger.com.atlassian.confluence.user=DEBUG,console
07.#log4j.additivity.com.atlassian.confluence.user=false
08.#log4j.logger.bucket.user=DEBUG,console
09.#log4j.additivity.bucket.user=false
```

To this:

```
01.###
02.# Atlassian User
03.###
04.log4j.logger.com.atlassian.user=DEBUG, confluencelog
05.log4j.additivity.com.atlassian.user=false
06.log4j.logger.com.atlassian.confluence.user=DEBUG,confluencelog
07.log4j.additivity.com.atlassian.confluence.user=false
08.log4j.logger.bucket.user=DEBUG,confluencelog
09.log4j.additivity.bucket.user=false
10.log4j.logger.com.atlassian.seraph=DEBUG, confluencelog
11.log4j.additivity.com.atlassian.seraph=false
12.log4j.logger.com.opensymphony.user=DEBUG, confluencelog
13.log4j.additivity.com.opensymphony.user=false
```

- After enabling both the above, please attempt a Confluence LDAP account login and attach a copy of the log files that are produced when the problem occurs. To do this, locate your install directory or exploded WAR directory, then zip the full /logs subdirectory into a single file for us to examine. An example location might be confluence-2.2.2-std/logs.



If you are using LDAP, run the External User Test tool and attach a copy of the output to the support ticket.

Paddle

From Confluence version 2.8 the '**External User Test**' tool is integrated into the Administration Console, and you can also run it outside of Confluence. In previous versions of Confluence, this tool was available only as a separate utility called '**Paddle**'.

- For **Confluence 2.8 and later**, please refer to Troubleshooting LDAP User Management.
- If you are running an earlier version of Confluence, please refer to the Confluence 2.7 documentation or choose the relevant Confluence version from the list of previous versions on the documentation home page.
- Sometimes for troubleshooting purposes it may still be useful to use the stand-alone Paddle so you don't have to start and stop Confluence each time you test. Place your atlassian-user.xml file in the same directory as the jar file and run: `java -jar paddle-2.0.jar`.



Be aware that:
The Atlassian-User optional setting `useUnqualifiedUsernameForMembershipComparison` is not supported by Paddle, so the tests may not be correct if this option is being used. In this case, the Connectivity test must can be done using Confluence.

Understanding User Management in Confluence



Looking for help with your user management configuration? See Requesting External User Management Support.



Try Atlassian Crowd for powerful user administration
Atlassian's Crowd is a web-based single sign-on (SSO) tool that simplifies application provisioning and identity management, including LDAP integration. For more information, please see the Crowd documentation on Integrating Crowd with Confluence.

On this page:

- Components of User Management
- Authentication
  - Seraph
  - XML-RPC and SOAP Authentication
  - Password Authentication and User Management
- Confluence User Management Frameworks
  - AtlassianUser
  - OSUser

- Related pages

Components of User Management

1. **Authentication** - determining what user identity is making a request to Confluence.
2. **User management** - storing and retrieving core information about users.
3. **Group membership** - storing and retrieving groups, and group membership.
4. **Profile information** - providing metadata associated with users.

It's important to understand that these are separate components of the user management system. The term *LDAP integration* is not really meaningful, because you could use LDAP repository for any or all of the above tasks.

For example, in OSUser authentication can be performed against different repository to that used for group membership queries. In AtlassianUser, authentication and group membership can be retrieved from LDAP, but profile information is still stored in the Confluence database.

Authentication

Seraph

Almost all authentication in Confluence (and JIRA) is performed through Seraph, Atlasian'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 performs no user management itself. It merely checks the credentials of the incoming request, and delegates any user-management functions (looking up a user, checking a user's password is correct) to Confluence's user-management subsystem.

If you were looking to integrate Confluence with a Single Sign-On (SSO) infrastructure, you would do so by writing a custom Seraph authenticator (and in fact, many customers have done so).

XML-RPC and SOAP Authentication

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

However, if the provided 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 and User Management

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.

Confluence User Management Frameworks

The rest of this document covers the user management frameworks used by Confluence: **AtlassianUser** and **OSUser**.

AtlassianUser

AtlassianUser is a new user and group management framework developed by Atlassian, and is the core framework used in Confluence since version 2.1. AtlassianUser was developed with the following goals in mind:

- Support LDAP as a fully functional repository for authentication, group management and profile information (profile information not yet implemented).
- Compatibility with both JIRA and Confluence (JIRA support not yet implemented).
- Be simple to configure.

AtlassianUser provides user, group and profile management services to Confluence. It does so by delegating tasks to configurable *repositories*. Multiple repositories can be configured, so for example Confluence can draw user information from both the database *and* an LDAP server.

Default Configuration

Configuration of AtlassianUser is done through the `<<confluence-install>>\confluence\WEB-INF\classes\atlassian-user.xml` file. See the `atlassian-user.xml` reference page. (In Confluence 2.1, configuration of AtlassianUser is done through the `atlassianUserContext.xml` file.)

For Confluence 2.7.0 and later:

- All user management is performed by AtlassianUser's native providers.
- OSUser delegation is still supported for customers who rely on the OSUser/JIRA bridge or the old-style OSUser LDAP support.
- Refer to the Confluence 2.7 Upgrade Guide for details of the automatic migration which may occur during the upgrade process, for

customers who are using the standard user management framework.

For Confluence 2.6.x and earlier:

- Confluence's AtlassianUser configuration delegates all user, group, profile and password authentication to OSUser.

Database (Hibernate) Support in AtlassianUser

AtlassianUser can store user, group and profile data directly in Confluence's database. This is the default behaviour for Confluence 2.7.0 and later.

LDAP Support in AtlassianUser

AtlassianUser currently supports *password authentication*, *user management* and *group management* with an LDAP server. Follow the instructions on configuring AtlassianUser LDAP integration.

At this point, only read-only access to LDAP is planned. Java's JNDI-LDAP interface does not support updating an LDAP repository, and the administration tools that come with LDAP servers such as Microsoft Active Directory are generally comprehensive and already available in enterprise IT departments.

JIRA Integration via AtlassianUser

AtlassianUser will not support delegating Confluence user management to JIRA. Instead, our goal is to implement AtlassianUser as the JIRA user management framework as well. Once this is done, both Confluence and JIRA can use the same LDAP server for their authentication and group management.

OSUser

OpenSymphony User was Confluence's core user management framework until it was replaced by AtlassianUser in version 2.1. OSUser is still supported through AtlassianUser's OSUser repositories. OSUser is also built around the model of pluggable providers, but its LDAP support is limited.

OSUser Database (Hibernate) Providers

In its default configuration, Confluence's OSUser providers store a list of users and groups together with profile information in tables in the Confluence database:

- os_user (authentication)
- os_group (group membership)
- os_user_group (group membership)
- os_propertyentry (profile information)

The hashed password in the os_user table is used to authenticate the user unless LDAP support is enabled. The os_user_group table is queried for group membership information.

OSUser configuration is controlled through the `<<confluence-install>>\confluence\WEB-INF\classes\osuser.xml` file.

LDAP Support in OSUser

OSUser only supports *authentication* against an LDAP server. That is, you can check user passwords against LDAP, but all other user information must be shadowed in the Confluence database. Follow the instructions on configuring OSUser LDAP authentication.

If you need support for LDAP user information or group membership as well, you should use AtlassianUser instead (see above).

Delegating User Management to JIRA via OSUser

Confluence can use OSUser to retrieve information for *authentication*, *group membership* and *profile information* from JIRA.

If you look at the discussion of OSUser's implementation above, you can see how this can works pretty easily. A data source to JIRA database is configured in Confluence which lets Confluence read directly from JIRA's os_* tables. For example, when a user is created in JIRA, the username and password goes in the os_user table in the JIRA database. Confluence looks at the same table in the JIRA database to authenticate the user.

Access to the JIRA database is read-only. For this reason, Confluence maintains a subset of the user's *profile information* locally in the Confluence database (things like last login time and user preferences that Confluence needs to be able to modify) .

Follow the instructions on configuring OSUser delegation to JIRA.

Related pages

HTTP authentication with Seraph
Single Sign-on Integration with JIRA and Confluence
Add LDAP Integration
LDAP Authentication with OSUser
Delegate user management to use JIRA logins

Migrating to new User Management
Confluence 2.7 Upgrade Guide
atlassian-user.xml reference

User Management Frequently Asked Questions

This page has been split into the LDAP FAQ and JIRA Integration FAQ.

Confluence Installation and Upgrade Guide

The pages listed below contain information on installing and upgrading Confluence:

- System Requirements
  - Server Hardware Requirements Guide
  - Example Size & Hardware Specifications From Customer Survey
- Supported Platforms
  - End of Support Announcements for Confluence
  - Supported Platforms FAQ
- Confluence Installation Guide
  - Installing Confluence Standalone
  - Installing the Confluence EAR-WAR Edition
  - Confluence Cluster Installation
  - Confluence Cluster Installation with Existing Data
  - Installing Sun JDK for Confluence
  - Confluence Unix and X11 Dependencies
  - Get A Confluence Licence
  - Running Confluence in a Virtualised Environment
  - Uninstalling Confluence Standalone
- Confluence Setup Guide
  - External Database
  - Load Content for the Site
  - Restoring from Backup During Setup
- Upgrading Confluence
  - Upgrading Confluence Standalone Distribution
  - Upgrading Confluence EAR-WAR Distribution
  - Upgrading Beyond Current Licensed Period
- Confluence Release Cycle
  - Release Notes
  - Development Releases
  - Coherence license changes SEPT 2009 - new Standard and Clustered Confluence Editions
  - Upgrade Notes Overview

System Requirements

Confluence is a 'web application', meaning it runs centrally on a server and users interact with it through web browsers from any computer. Hence, Confluence requires hardware and some additional software in order to operate. The software platforms currently supported for Confluence are listed on the Supported Platforms page. However, more detailed information regarding Confluence's software and hardware requirements can be found on this page. This page also covers information about:

- Platforms which Atlassian does not support, but should work with Confluence
- Platforms and other software that are incompatible with Confluence or have known problems running in conjunction with Confluence.

On this page:

- Introduction
  - Atlassian's Hosted Solutions are an Alternative
- Confluence Software Requirements
  - Client Software Requirements
    - Web Browsers
  - Server Software Requirements
    - Java
    - Operating Systems
    - Application Servers
    - Databases
    - Antivirus Software Configuration
- Confluence Hardware Requirements
  - Server Load
  - Disk Space

Introduction

Confluence works with a broad range of operating systems, database systems and application servers. Provided you have the technical knowledge, it is very likely that you will be able to run Confluence with an 8-year-old database or even on some 8-year-old hardware. Realistically, it is not technically feasible for us to provide our legendary support service on all environments available. There can only be a

finite number of platforms and release versions of those that we support.

Our rule of thumb when releasing a new version of Confluence is that we will officially support platforms that have been released within the last one to two years (or the latest version of that platform if no new version of it was released in that period). This does not necessarily mean that you will need to upgrade your database or application server every time you upgrade Confluence. However, if you do run into problems with an unsupported version of a database or application server, we may have to ask you to upgrade to something newer.

For example, you are running Confluence 2.7.3 with PostgreSQL 8.0 and everything works fine. You decide to upgrade to Confluence 2.8, which only supports version 8.1 of PostgreSQL. The chances are that you can run Confluence 2.8 with PostgreSQL 8.0 with no problems whatsoever. However, if you do run into problems, we may ask you to upgrade to PostgreSQL 8.1 or 8.2 before we dive deeper into the problem and provide further assistance.

Please refer to our Supported Platforms topic for details on platforms that we currently support in this version of Confluence and our Supported Platforms FAQ topic for details on our support handling procedures.

Atlassian's Hosted Solutions are an Alternative

If you do not have the resources to set up and maintain a Confluence installation locally, how about using one of our hosted solutions instead? We can run and maintain your own installation of Confluence on one of our servers and we will handle all the testing, monitoring and upgrading processes for you. For more information, please refer to our Confluence Hosted and our integrated JIRA Studio solutions on our website.

Confluence Software Requirements

Please read the Supported Platforms page for Confluence. This page contains important information about all client and server software requirements, which are supported for **Confluence 3.2.x**.

Client Software Requirements

Please read the following additional information regarding client software requirements for Confluence.

Web Browsers



Internet Explorer 6, Mozilla Firefox 2 and Safari 2 will no longer be supported in future versions of Confluence. Please read our End of Support Announcements for Confluence.

Server Software Requirements

Please read the following additional information regarding server software requirements for Confluence.

Java

You will need to install a Java Development Kit (JDK) on your operating system before before proceeding with a Confluence installation. For instructions on installing the Sun JDK for Windows and Linux/UNIX, please refer to Installing Sun JDK for Confluence.



Important notes about installing a JDK for Confluence

- Confluence requires the full installation of a JDK. It is not enough to run Confluence on a Java Runtime Environment (JRE) alone.
- While JDK 1.5 is fine, JDK 6 (1.6) is the preferred platform as it is faster and more reliable. Please refer to our End of Support Announcements for Confluence.

We do not provide support for non-Sun JDKs. However, if you wish to install a non-Sun JDK and you want to use SSL, you will also need to install the Sun JSSE package.

OpenJDK is currently not supported. A JIRA issue to request support for this JDK has been created (CONF-16431).

Operating Systems

If you would like to run Confluence on VMware, please read our Running Confluence in a Virtualised Environment document first.



Confluence on Virtualised Environments

Atlassian officially supports non-clustered installations of Confluence 3.0 and later on VMware. Although possible, we do not recommend (nor support) running versions of Confluence prior to 3.0 on VMware, since Confluence 3.0 resolved many performance issues that were present in earlier versions. Be aware that we also do not support clustered installations of Confluence on VMware.

Application Servers

Atlassian supports the application servers listed on the Supported Platforms page, provided they are running on Windows, Linux, a UNIX-based operating system (such as NetBSD, FreeBSD, OpenBSD and Solaris) or Mac OS X.

If you have no preference for a particular application server and wish to set up Confluence for production purposes, we highly recommend installing Confluence Standalone, which includes the Apache Tomcat application server.

We assume that **Confluence 3.2.x** works fine with the application server versions listed below. However, we do not test these versions regularly and we may ask you to migrate to one of the supported application servers before we can provide you with further support.

- **Apache Tomcat** — 5.5
- **BEA Weblogic** — 9.1, 10

The following application servers possess problems or incompatibilities with Confluence and consequently are not supported:

- **Unsupported J2EE-based Application Servers:**
    - Oracle OC4J / Oracle Application Server — see reported issue
    - Sun Application Server — see reported issue
    - GlassFish — see reported issue
    - SAP-J2EE — see reported issue
    - Macromedia JRun
- **Unsupported non-J2EE-based Application Servers:**
    - Microsoft IIS

Databases

Atlassian supports the databases listed on the Supported Platforms page, provided they are running on Windows, Linux, a UNIX-based operating system (such as NetBSD, FreeBSD, OpenBSD and Solaris) or Mac OS X.

If you have no preference for a particular database and wish to set up Confluence for production purposes, we highly recommend using PostgreSQL. This is a scalable, robust and free database server that is also easy to set up. For database setup information, please refer to Database Setup For Any External Database.

We assume that **Confluence 3.2.x** works fine with the database versions listed below. However, we do not test these versions regularly and we may ask you to migrate to one of the supported databases before we can provide you with further support.

- **PostgreSQL** — 8.0, 8.3
- **MySQL** — 5.0 - 5.0.27 (using the InnoDB storage engine, not MyISAM)
- **Oracle** — 11

Antivirus Software Configuration

The presence of antivirus software on your operating system running Confluence greatly decreases the performance of Confluence. Antivirus software that intercepts access to the hard disk is particularly detrimental and may even cause errors in Confluence.

You should configure your antivirus software to ignore the following directories:

- Confluence home directory
- Confluence's index directory
- All database-related directories

 This recommendation above is particularly important if you are running Confluence on Windows. No matter how fast your hardware is, antivirus software will almost always have a negative impact on Confluence's performance and may render Confluence impossible to use.

Confluence Hardware Requirements

Please be aware that while some of our customers run Confluence on SPARC-based hardware, Atlassian only officially supports Confluence running on x86 hardware and 64-bit derivatives of x86 hardware.

Server Load

Server load depends primarily on the number of users online at once and their usage of Confluence.

Under 25 concurrent users:

- 2GHz+ CPU or equivalent
- 512MB RAM

Over 25 concurrent users:

- Dual 2GHz+ CPU Xeon or equivalent
- 2GB RAM

Over 100 concurrent users:

- Quad 2GHz+ CPU Xeon or equivalent
- 4GB of RAM

See Server Hardware Requirements Guide for details.

Refer also to the tips on reducing out of memory errors, in particular the section on Permanent Generation Size.

Disk Space

Confluence Install Directory - 250MB

- Install files
- Nightly site backups
- Temporary files

Confluence Home Directory or **External Database** - 250MB minimum

- Text content
- File attachments

RELATED TOPICS

End of Support Announcements for Confluence
Confluence Installation Guide
Confluence Setup Guide
Installing Confluence Standalone Using the Windows Evaluation Installer
Installing the Confluence EAR-WAR Distribution
Confluence Cluster Installation
Example Size & Hardware Specifications From Customer Survey
Adding Confluence to a JIRA Standalone installation
Confluence Documentation Home
Server Hardware Requirements Guide
Supported Platforms FAQ

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 the Confluence server. Visitors are counted from their first page request until the connection is closed and if public access is enabled, this includes internet visitors as well as logged in users. Storage requirements will vary depending on how many pages and attachments you wish to store inside Confluence.

Minimum Hardware Requirements

On small instances, server load is primarily driven by peak visitors.

5 Concurrent Users

- 2GHz+ CPU
- 512MB RAM
- 5GB database space

25 Concurrent Users

- Quad 2GHz+ CPU
- 2GB+ RAM
- 10GB database space

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.

Example Hardware Specifications

These are example hardware specifications for non-clustered Confluence instances. It not recorded whether the RAM refers to either total server memory or memory allocated to the JVM, while blank settings indicate that the information was not provided.

Accounts	Spaces	Pages	CPUs	CPU (GHz)	RAM (Meg)	Notes
150	30	1,000	1	2.6	1,024	
350	100	15,000	2	2.8	700	
5,000	500		4	3	2,024	
10,000	350	16,000	2	3.8	2,024	
10,000	60	3,500	2	3.6	4,048	
21,000	950		2	3.6	4,048	
85,000	100	12,500	4	2.6	4,048	3 machines total: application server, database server, Apache HTTPD + LDAP tunnel server. See Accenture's slides and video for full details

Server Load & Scalability

When planning server hardware requirements for your Confluence deployment, you will need to estimate the server scalability based on peak visitors, the editor to viewer ratio and total content.

- The editor to viewer ratio is how many visitors are performing updates versus those only viewing content
- Total content is best estimated by a count of total spaces

Confluence scales best with a steady flow of visitors rather than defined peak visitor times, few editors and few spaces. Users should also take into account:

- Total pages is not a major consideration for performance. For example, instances hosting 80K of pages can consume under 512 meg of memory
- Always use an external database, and check out the performance tuning guides.

As mentioned on the documentation for Operating Large or Mission-Critical Confluence Installations, some important steps are loadtesting your usecase and monitoring the system continuously to find out where your system could do better and what might need to improve in order to scale further.

Maximum Reported Usages

These values are largest customer instances reported to Atlassian or used for performance testing. Clustering for load balancing, database tuning and other performance tuning is recommended for instances exceeding these values.

Most Spaces	1700
Most Internal Users	15K
Most LDAP Users	100K
Most Pages	80K

Hard Disk Requirements

All wiki content is stored in the database, while attachments use either the database or filesystem. For example, the wiki instance you are reading now uses approximately 1 GB of database space and 9.4 GB of disk space.

Here is a breakdown of the disk usage requirements for this wiki, as at December 2008:

Database size	1003 MB
Home directory size	9.4 GB

Size of selected database tables

Data	Rows	Size

Content bodies (incl. all versions of blogs, pages and comments)	170462	145 MB
Content metadata (incl. title, author)	188697	48 MB
Content and user properties	193652	42 MB
Users	20679	5.8 MB
Attachment metadata	25718	5.0 MB
Labels	43235	4.5 MB

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)	27484	5.9 GB
Usage index (now disabled)	240	2.6 GB
Search index	10	236 MB
Office Connector cache	44	222 MB
Temporary files	7269	201 MB
Plugin files	1508	139 MB
Thumbnails	10154	84 M
Did-you-mean search index	3	9.9 MB

Note: not all files are shown, and average file size may vary between instances.

Private & Online Comparison

Private instances manage their users either internally or through a user repository such as LDAP, while online instances have public signup enabled and must handle the additional load of anonymous internet visitors. Please keep in mind that these are examples only, not recommendations:

Use Case	Spaces	User Accounts	Editors	Editor To Viewer Ratio	Pages	Page Revisions	Attachments	Comments	Total Data Size (GB)	Notes
Online Documentation	140	11,500	1,000	9%	8,800	65,000	7,300	11,500	10.4	
Private Intranet	130	180	140	78%	8,000	84,000	3,800	500	4.5	
Company-Wide Collaboration	100	85,000	1,000+	1%+	12,500	120,000	15,000			Accenture - see slides and video for full details

Professional Assistance

For large instances, it may be worthwhile contacting an Atlassian partner for expertise on hardware sizing, testing and performance tuning. Simply contact a local partner directly or email our partner manager for a recommendation.

Related Pages

- Powered By Confluence
- Managing Application Server Memory Settings
- Example Size & Hardware Specifications From Customer Survey
- Confluence Installation Guide
- Confluence Clustering Overview
- Operating Large or Mission-Critical Confluence Installations

Performance Testing Scripts

Example Size & Hardware Specifications From Customer Survey

Below are the results of a survey conducted by Atlassian in July 2007, showing some capacity statistics for Confluence users. The figures are broken down by industry and number of users.

Num Users	Length of time in production	Database	Application Server	Num CPUs/Cores	Physical Memory/RAM	Operating System	Satisfactio with Confluenc Performar
Banking/Finance							
26 - 50	3-6 Months Ago	Microsoft SQL Server	Standalone/Apache Tomcat	2	2G	Windows	Neutral
26 - 50	2 Years Ago	Sybase ASE	Weblogic	>8	>16G	Unix	Satisfied
51 - 250	3-6 Months Ago	Oracle	Standalone/Apache Tomcat	2	4G	Unix	Neutral
501 - 1,000	3-6 Months Ago	Microsoft SQL Server	Websphere	2	2G	AIX	Satisfied
1,001 - 5,000	3-6 Months Ago	Oracle	Standalone/Apache Tomcat	2	4G	Windows	Satisfied
1,001 - 5,000	2 Years Ago	Oracle	Websphere	4	>16G	Solaris	Extremely Satisfied
5,001 - 10,000	10-12 Months Ago	Microsoft SQL Server	Standalone/Apache Tomcat	4	16G	Linux	Satisfied
Education							
1-25	2 Years Ago	DB2	Standalone/Apache Tomcat	2	2G	Linux	Satisfied
26 - 50	10-12 Months Ago	MySQL	Standalone/Apache Tomcat	2	2G	Linux	Extremely Satisfied
51 - 250	<3 Months Ago	Oracle	Standalone/Apache Tomcat		1G	Windows	Unsatisfiec
51 - 250	10-12 Months Ago	Oracle	Standalone/Apache Tomcat	1	2G	Unix	Extremely Satisfied
Engineering/Aerospace							
251 - 500	7-9 Months Ago	Oracle	Standalone/Apache Tomcat	1	1G	Mac OS X	Satisfied
1,001 - 5,000	7-9 Months Ago	Microsoft SQL Server	JBoss	2	4G	Linux	Satisfied
Entertainment							
1,001 - 5,000	10-12 Months Ago	PostgreSQL	Standalone/Apache Tomcat	2	8G	Linux	Extremely Satisfied
Government							
51 - 250	2 Years Ago	MySQL	Standalone/Apache Tomcat	2	2G	Mac OS X	Extremely Satisfied
Technology							
501 - 1,000	7-9 Months Ago	MySQL	Standalone/Apache Tomcat	1	2G	Linux	Satisfied
Telecommunications & Media							

1-25	3-6 Months Ago	Standalone/HSQL	Standalone/Apache Tomcat	1		Linux	Satisfied
1-25	7-9 Months Ago	MySQL	Standalone/Apache Tomcat	1	2G	Linux	Satisfied
26 - 50	10-12 Months Ago	MySQL	Standalone/Apache Tomcat	2	2G	Linux	Satisfied

Supported Platforms

This page describes the supported platforms for Confluence. Please review them before installing Confluence.

> **End of Support Announcements**
> Please read End of Support Announcements for Confluence for important information regarding the end of support for various platforms and browsers when used with Confluence.

For any further information about these supported platforms and for information on hardware requirements, please refer to our System Requirements document.

On this page:

Key: ✓ = Supported. ✗ = Not Supported

Java Version	
Sun JDK	✓ 1.5, 1.6 (Note: Java 6 update 18 64-bit has known issues) ✗ 1.4
Operating Systems	
Microsoft Windows (including 64-bit) (1)	✓
Linux / Solaris (1, 2)	✓
Apple Mac OS X (1)	✓
Application Servers	
Apache Tomcat	✓ 5.5.20 - 6.0
BEA Weblogic	✓ 9.2
IBM Websphere	✓ 6.1
Caucho Resin	✓ 3.0, 3.1.6, 3.1.7
Databases	
PostgreSQL	✓ 8.1, 8.2
MySQL (3)	✓ 5.0.28+
Oracle	✓ 10.1, 10.2
Microsoft SQL Server	✓ 2005
DB2	✓ 8.2, 9.7
HSQLDB (4)	✓ (for evaluation purposes only)

Web Browsers	
Microsoft Internet Explorer (Windows)	6, 7 and 8
Mozilla Firefox (all platforms)	3, 3.5, 3.6
Safari	3, 3.1, 4

1. Confluence is a pure Java application and should run on this platform provided all other JDK requirements are satisfied.
2. While some of our customers run Confluence on SPARC-based hardware, Atlassian only officially supports Confluence running on x86 hardware and 64-bit derivatives of x86 hardware.
3. Ensure that you configure your Confluence MySQL database to use the InnoDB storage engine as the MyISAM storage engine could lead to data corruption.
4. HSQLDB: Confluence ships with a built-in HSQL database. While this database is fine for evaluation purposes, it is somewhat susceptible to data loss during system crashes. Hence, for production environments, we recommend that you configure Confluence to use an external database.

RELATED TOPICS

Confluence Installation Guide
Confluence Setup Guide
Installing Confluence Standalone Using the Windows Evaluation Installer
Installing the Confluence EAR-WAR Distribution
Confluence Cluster Installation
Example Size & Hardware Specifications From Customer Survey
Adding Confluence to a JIRA Standalone installation
Confluence Documentation Home
Server Hardware Requirements Guide
Supported Platforms FAQ

End of Support Announcements for Confluence

This page contains announcements of the end of support for various platforms and browsers when used with Confluence. Please see the sections below for details.

On this page (most recent announcements first):

- 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 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.

Why is Atlassian doing this?

Atlassian is committed to delivering improvements and bug fixes as fast as possible. We are also committed to providing world class support for all the platforms our customers run our software on. However, as the complexity of our applications grows, the cost of supporting multiple platforms increases exponentially. Each new feature has to be tested on several combinations of databases, with setup and ongoing maintenance of automated tests. Moving forward, we want to reduce the time spent there to increase Confluence development speed significantly.

Deprecated Application Servers for Confluence (27 January 2010)

This section announces the end of Atlassian support for certain application servers for Confluence. End of support means that Atlassian will not fix bugs related to certain application servers past the support end date.

We will **stop supporting the following application servers**:

- From Confluence 3.2, due late Q1 2010, Confluence will no longer support JBoss application servers.
- From Confluence 3.3, due in Q3 2010, Confluence will no longer support Oracle Weblogic, IBM Websphere or Caucho Resin.

We are reducing our application server platform support to reduce the amount of testing time and help us speed up our ability to deliver market-driven features. We are committed to helping our customers understand this decision and assist them in migrating to Tomcat, our supported application server.

The details are below. Please refer to the Supported Platforms for more details regarding platform support for Confluence. If you have questions or concerns regarding this announcement, please email `eol-announcement at atlassian dot com`.

End of Life Announcement for Application Server Support

Application Servers	Support End Date
JBoss 4.2.2	When Confluence 3.2 releases, due late Q1 2010
Oracle Weblogic 9.2	When Confluence 3.3 releases, due Q3 2010
IBM Websphere 6.1	When Confluence 3.3 releases, due Q3 2010
Caucho Resin 3.0, 3.1.6, 3.1.7	When Confluence 3.3 releases, due Q3 2010

- **JBoss End of Support Notes:**
  - 'Support End Date' means that Confluence 3.1 and previous released versions will continue to work with stated application servers. However, we will not fix bugs affecting JBoss application servers.
  - Confluence 3.2 will not support JBoss application servers.

- **Websphere, Weblogic, 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?

Atlassian is committed to delivering improvements and bug fixes as fast as possible. We are also committed to providing world class support for all the platforms our customers run our software on. However, as the complexity of our applications grows, the cost of supporting multiple platforms increases exponentially. Each new feature has to be tested on several combinations of application servers, with setup and ongoing maintenance of automated tests. At times, 30% of the development team is busy coding solutions for edge cases in various application servers. Moving forward, we want to reduce the time spent there to increase Confluence development speed significantly.

We have chosen to standardise on Tomcat, because it is the most widely used application server in our user population. It is fast, robust, secure, well-documented, easy to operate, open source, and has a huge community driving improvements. It is the de facto industry standard, with several companies available that specialise in providing enterprise grade support contracts for it, ranging from customisations to 24/7 support.

Deprecated Java Platforms for Confluence (27 January 2010)

This section announces the end of Atlassian support for certain Java Platforms for Confluence.

We will **stop supporting the following Java Platforms**:

- From Confluence 3.3, due Q3 2010, support for Java Platform 5 (JDK/JRE 1.5) will end.

We are ending support for Java Platform 5, in line with Sun's Java SE Support Road Map (i.e. "End of Service Life" for Java Platform 5 dated October 30, 2009). We are committed to helping our customers understand this decision and assist them in updating to Java Platform 6, our supported Java Platform.

The details are below. Please refer to the Supported Platforms for more details regarding platform support for Confluence. If you have questions or concerns regarding this announcement, please email `eol-announcement at atlassian dot com`.

End of Life Announcement for Java Platform Support

Java Platform	Support End Date
Java Platform 5 (JDK/JRE 1.5)	When Confluence 3.3 releases, due Q3 2010

- **Java Platform 5 End of Support Notes:**
  - Atlassian intends to end support for Java Platform 5 in Q3 2010.
  - 'Support End Date' means that Confluence 3.2.x and previous released versions will continue to work with Java Platform 5 (JDK/JRE 1.5), however we will not fix bugs related to Java Platform 5 past the support end date.
  - Confluence 3.3 will only be tested with and support Java Platform 6 (JDK/JRE 1.6).
  - If you have concerns with this end of support announcement, please email `eol-announcement at atlassian dot com`.

Deprecated Web Browsers for Confluence (14 December 2009)

This section announces the end of Atlassian support for certain web browsers for Confluence.

We will **stop supporting older versions of web browsers** as follows:

- From Confluence 3.2, due late Q1 2010, support for Firefox 2 and Safari 2 will end.
- From 13 July 2010, in line with Microsoft's Support Lifecycle policy, support for IE6 will end.

The details are below. Please refer to the Supported Platforms for more details regarding platform support for Confluence. If you have questions or concerns regarding this announcement, please email `eol-announcement at atlassian dot com`.

End of Life Announcement for Web Browser Support

Web Browsers	Support End Date
Firefox 2	When Confluence 3.2 releases, late Q1 2010
Safari 2	When Confluence 3.2 releases, late Q1 2010
Internet Explorer 6	When Confluence 3.3 releases (target Q3 2010) or 13 July 2010, whichever is sooner

- **Firefox 2 and Safari 2 Notes:**
  - Confluence 3.1 is the last version to officially support Firefox 2 and Safari 2.
  - You may be able to use these older browser for the most common use cases like viewing and editing content, but official support for these browsers will end once you upgrade to Confluence 3.2.
  - Confluence 3.2 is currently targeted to release late Q1 2010 and will not be tested with Firefox 2 and Safari 2. After the Confluence 3.2 release, Atlassian will not provide fixes in older versions of Confluence for bugs affecting Firefox 2 and Safari 2.
- **Internet Explorer 6 Notes:**
  - Confluence 3.2 (due late Q1 2010) will be the last version to officially support Internet Explorer 6.
  - Confluence 3.3 is currently targeted to release Q3 2010 and will **not** support IE6.
  - Atlassian will support IE6 in Confluence until the 13th of July 2010, in line with Microsoft's Support Lifecycle policy. Beyond that date, released versions of Confluence will continue working with IE6 just as they did before, but we will not fix bugs affecting Internet Explorer 6.
  - You may be able to use Internet Explorer 6 for the most common use cases like viewing and editing content, but official support for this browser will end once you upgrade to Confluence 3.3.

Java 1.4 Support Timeline



This notice was first published on January 4th, 2008, and later updated with the release of Confluence 2.9. As from Confluence 2.9, Java 1.4 is no longer supported. You will need Java 5 or later.

What is happening?

As part of the ongoing development of Confluence, we have raised our minimum supported version of the Java platform.

- Confluence version 2.8 was the last major version to support Java 1.4.
- Confluence 2.9 and later require at least Java 5.

What does this mean to me?

I use Confluence

Users of Confluence websites should see absolutely no change.

I administer a Confluence Server

If you are running Confluence 2.8 or one of the 2.8.x patch releases, your current version of Confluence will continue to run in your current environment without change.

If you choose to upgrade to Confluence 2.9, you will need to ensure your environment is running at least Java 5.

You can check your current Java version in Confluence:

1. Go to the Confluence '**Administration Console**'. To do this:
    - Open the '**Browse**' menu and select '**Confluence Admin**'. The 'Administration Console' view will open.
2. Select '**System Information**' from the '**Administration**' section in the left-hand panel.
3. Refer to '**Java Version**'.
    - If the version is 1.5 or higher, you do not need to do anything.
    - If the version is 1.4, you need to upgrade your JDK before you can upgrade to Confluence 2.9.

If you are running the Confluence EAR-WAR edition against your own application server, you will need to check with your application server vendor about which JDK versions are supported.

I am a Confluence Plugin/Extension Developer

Plugin developers who want their plugins to work on Confluence 2.8 and earlier should continue to compile their plugins with the Java 1.4 compiler. Plugin developers specifically targeting Confluence 2.9 and later should use the Java 5 compiler and Java 5 language features.

Why Now?

Our normal policy for JDK support is to follow Sun's Java Technology End-of-Life policy, where only the most recent three major versions of Java are supported. On Sun's original timeline for the release of Java 7, Java 1.4 would have been scheduled for EOL in (Northern Hemisphere) Spring 2008. Sun's release roadmap for Java 7 has since been pushed back to 2009, but we feel that it is in the best interests of Confluence to stick to the original schedule.

Given Java 1.4's near-obsolescence, saved only by the slipping schedule of Java 7, IT departments should already be planning to transition away from Java 1.4. Our surveys of customers suggest that most are already running Java 5, and those that don't are running application servers that can easily support the new version. As such, the cost of continuing to support the old version, both in developer and support resources, cannot really be justified.

Progress on this issue can be tracked here: CONF-10365

Supported Platforms FAQ

Q: How does Atlassian choose which JDK 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 JDK versions to support based on the recommended environments for these servers.

Q: What is a supported platform?

A supported platform is one that:

- Confluence is regularly tested on during the development cycle
- One that is available within Atlassian for support technicians and developers to reproduce problems
- Bugs raised against it will be given a high priority

Supporting a platform means we know how to get Confluence running in that environment and can troubleshoot Confluence issues within it. It does not mean we have any particular expertise beyond that. As such, we may not be able to provide assistance with customising or tuning that application server or database. (Atlassian support is not a substitute for a good database administrator.)

Q: Can I get assistance with running Confluence on a platform that is *not* supported?

If you are running Confluence on an unsupported platform, then we can not guarantee providing any support for it. Furthermore, we will recommend that you switch to a platform which is supported.

Q: If you write your application to standards like J2EE, JDBC and SQL, doesn't that mean it should run on any compliant server?

Confluence is a complicated application and we commonly encounter interesting edge-cases where different servers have interpreted the specifications differently. Then again, each server has its own different collection of bugs.

Q: How can I get Atlassian to support Confluence on a new platform?

Supporting a new platform involves a significant investment of time by Atlassian, both up-front costs to set up new testing environments and fix any issues we might encounter and the ongoing costs involved in maintaining the application against this new environment in the future. As such, supporting a new platform is not something we will do unless we know there is significant demand for it.

Please be aware that your interest alone will not be enough for us to add support for your application server or database. We would need to see a significant number of votes on the issue raised in our public JIRA site or a significant level of interest in our forums, before considering supporting that platform.

Q: My organisation has standardised on an operating environment that Confluence does not support. What can I do?

In this situation, you have the following two options:

1. Run Confluence in the unsupported environment, with the caveats mentioned above.
2. Make an exception to your standardised operating environment and set up Confluence based on its supported platforms.

Confluence Installation Guide

Welcome to the Confluence Installation Guide!

> **What is Confluence?**
>
> Confluence is an enterprise wiki that makes it easy for your team to collaborate and share knowledge. A wiki is a web application that lets you edit web pages easily and immediately. No waiting, just click, type, and click again.

Want one? Let's go!

Before you Start

Please check the following points:

1. Make sure that your system meets the **minimum requirements** to run Confluence:
   - If you are installing Confluence for **evaluation** purposes, it should be pretty easy. You can use the Confluence Installer for Windows or Mac, or the zip archives. You will need a web browser — we recommend Firefox or Internet Explorer.
   - For production installations, use the zip archives. Please read the detailed system requirements.
2. Please verify that this version of the Confluence documentation matches that of the Confluence version you are installing. The Confluence documentation version you are currently viewing is indicated toward the top of the page tree on the left or in the 'breadcrumb trail' in the top banner of this page. If you need to access a different version of the Confluence documentation, use the control at the top of the page tree on the left or you can access it from the documentation home page.

Installing and Setting Up Confluence

> **Two phases: Installation and Setup**
>
> There are two phases to the Confluence installation and setup procedure. The instructions on this and following pages will lead you through both phases. Overview:
>
> - **Install Confluence onto your computer**. We provide detailed instructions for installing the Standalone distribution and the EAR/WAR distribution.
> - **Provide initial setup information.** The Setup Wizard will prompt you for the information needed to get you up and running.

Choose the type of Confluence installation you'd like from the table below, and follow the link(s) to the remaining installation instructions. When you have finished the installation phase, you will be prompted to start the setup phase.

Installation Type	Description
Evaluation installer for Windows or Evaluation Installer for Mac OS X	These are the simplest options for installing and evaluating Confluence. Choose these options if you wish to **evaluate** Confluence on a Windows- or a Mac OS X-based system. ⚠ These installers are specialised versions of the Standalone distribution (below) that should **only be used to evaluate Confluence**.

Standalone distribution	The 'Standalone' distribution provides the simplest method of installing Confluence for long term or production use. Choose this option if any of the following is true: • You are **evaluating** Confluence and may likely use this installation for **long term** or **production** purposes. • You want to run the Confluence application on the Apache Tomcat application server bundled with this distribution. • You are not sure **what** you want, except to get Confluence up and running with minimum fuss.
EAR/WAR distribution	This distribution allows you to deploy Confluence onto your own existing application server, instead of the Apache Tomcat server bundled with the Standalone distribution.
Confluence Clusters	Please read the Confluence Clustering Overview and the Cluster Checklist before you consider installing Confluence in a cluster.
Upgrading Confluence	Choose this option if you want to upgrade an existing Confluence installation to a new release of Confluence.

RELATED TOPICS

Upgrading Confluence
System Requirements

Installing Confluence Standalone

Choose the type of Confluence Standalone installer you'd like from the table below and follow the link to the installation instructions. When you have finished the installation phase, you will be prompted to start the setup phase.

If you have not already done so, please verify that this version of the Confluence documentation matches that of the Confluence version you are installing. The Confluence documentation version you are currently viewing is indicated toward the top of the page tree on the left or in the 'breadcrumb trail' in the top banner of this page. If you need to access a different version of the Confluence documentation, use the control at the top of the page tree on the left or you can access it from the documentation home page.

Confluence Standalone installer	Description
Automatic installer for Windows or Evaluation Installer for Mac OS X	These are the simplest options for installing and evaluating Confluence. Choose these options if you wish to **evaluate** Confluence on a Windows- or a Mac OS X-based system. ⚠ These installers are specialised versions of the Standalone distribution that should **only be used to evaluate Confluence**.
Windows zip file	Instructions for installing Windows-based Confluence production systems using the Confluence Standalone Windows zip file installer.
Mac OS X gzipped tar file	Instructions for installing Mac OS X-based Confluence production systems using the Confluence Standalone gzipped tar file installer.
Linux/Unix gzipped tar file	Instructions for installing Linux- or Unix-based Confluence production systems using the Confluence Standalone gzipped tar file installer.

The Standalone distribution includes Apache Tomcat as the standalone application server. If you want to install an EAR/WAR distribution for deployment on your own existing application server, please refer to Installing the Confluence EAR-WAR Distribution.

Installation Videos

Check out our installation videos for Windows or Mac.

Take me back to the Confluence Installation Guide.

Change listen port for Confluence Standalone

Problem

This page tells you what to do if you get errors like the following when starting Confluence Standalone, and can't access Confluence on port **8080**.

```
java.net.BindException: Address already in use: JVM_Bind:8080
```

This means you are running other software on Tomcat's default port **8080**. This may either be another Tomcat or some other process. 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 **8080**.

Solution: Change the Ports which Confluence Listens On

To change the ports for Confluence Standalone, 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

```
1.<Server port="8000" shutdown="SHUTDOWN" debug="0">
2.          <Service name="Tomcat-Standalone">
3.                    <Connector className="org.apache.coyote.tomcat4.CoyoteConnector" port="8080"
minProcessors="5" maxProcessors="75"
4.                              enableLookups="true" redirectPort="8443" acceptCount="10" debug="0"
connectionTimeout="20000" useURIValidationHack="false"/>
5.                    ...
```

You need to modify both the **server** port (default is 8000) and the **connector** port (default is 8080) to ports that are free on your machine.

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 '8015' and '8090':

Modified conf/server.xml using ports 8015 and 8090

```
1.<Server port="8015" shutdown="SHUTDOWN" debug="0">
2.          <Service name="Tomcat-Standalone">
3.                    <Connector className="org.apache.coyote.tomcat4.CoyoteConnector" port="8090"
minProcessors="5" maxProcessors="75"
4.                              enableLookups="true" redirectPort="8443" acceptCount="10" debug="0"
connectionTimeout="20000" useURIValidationHack="false"/>
5.                    ...
```

To access Confluence in this configuration, point your web browser to http://localhost:8090/.

NOTES

[1] For more information on netstat, see using netstat on Windows, or netstat man page (Linux).

[2] JIRA Standalone also runs on port **8080** by default. If you're looking to change the port of JIRA Standalone, see Changing JIRA Standalone's port.

RELATED PAGES

Installing Confluence Standalone Using the Windows Evaluation Installer
Running Tomcat on a Different Port
Documentation Home

Installing Confluence Standalone on Mac OS X from Zip File

These instructions apply to:

- The **Standalone distribution** of Confluence. The Standalone distribution includes Apache Tomcat as the standalone application server. If you want to install an EAR/WAR distribution for deployment on your own existing application server, please refer to the Confluence Installation Guide.
- **Mac OS X systems**. For other operating systems please refer to the Confluence Installation Guide.
- Installation via a **zipped download file**. For a simpler installation procedure for evaluation purposes, please consider using the Confluence Installer instead.
- Preparation for a production instance of Confluence.

Also, please check the version of Confluence which you are installing. Refer to the documentation home page to verify the latest Confluence version and to find documentation for older versions.

On this page:

- 1. Before you Start
- 2. Check for the JDK (Java Development Kit)
- 3. Download and Unzip the Confluence Installation File
- 4. Define your Confluence Home Directory
- 5. Check the Ports
- 6. Select an External Database
- 7. Start Confluence
- 8. Next Step is the Confluence Setup Wizard

1. Before you Start

Please check the following points:

1. Ensure that your system meets the minimum requirements to run Confluence. For more information, please read the detailed system requirements.
2. Have your Confluence license key ready. You can obtain a trial, free or commercial license now, or retrieve your existing license key.
3. You must be able to use a command prompt and install Java to continue. If not, please contact your system administrator to assist you or consider the Confluence Hosted evaluation option.
4. We recommend that you run Software Update to ensure that your operating system is up to date.

2. Check for the JDK (Java Development Kit)

The jdk is installed on OS X by default. From Applications > Utilities, open Terminal. Run `java -version` to see if its installed.

3. Download and Unzip the Confluence Installation File

1. If you have not downloaded Confluence already, download the Standalone tar.gz file for Macs.
2. Create a folder in your user's home folder called 'confluence'. The Confluence application and data will live in two separate folder underneath this folder once you are done.
    - Ensure that there are no spaces in the folder name, or in the names of any of its parent folders.
3. Find the downloaded tar.gz file in the Finder. Drag it to the confluence directory in your home folder. Double-click on the file to extract it to a directory called `confluence-2.7.0-std` or something similar.
    - This folder that includes the Confluence application is called the Confluence Installation directory.
4. Set up an environment variable which tells Confluence where your Java Virtual Machine is installed.
    - Inside the Confluence Installation directory, find the `bin/setenv.sh` file.
    - Open the `setenv.sh` file with TextEdit.
    - Add the following line as the first line of the file:

```
1.export JAVA_HOME=/Library/Java/Home
```

    - Save the file.
5. Create a new folder named `data` inside the 'confluence' folder in your user's home folder. This new folder is called the Confluence Home directory.
    - Do not put your Confluence Home directory inside the Confluence Installation folder ('confluence-2.7.0-std') or your installation will have problems with upgrades.

4. Define your Confluence Home Directory

Now you need to define the Confluence Home directory. This is where Confluence will store its configuration information, indexes and attachments.

Tip: Another term for 'Home directory' would be 'data directory'.

We suggest using different paths for your installation and home directories. This will facilitate upgrades.



Examples of Installation and Home Directories
Installation directory: `/Users/example/confluence/confluence-2.7.0-std`
Home directory: `/Users/example/confluence/data`

1. Open your Confluence Installation directory (created when you unzipped Confluence — see above).
2. Under the Installation directory, find this file: `/confluence/WEB-INF/classes/confluence-init.properties`
3. Open the `confluence-init.properties` file in a text editor such as Textedit.
4. Scroll to the bottom and find this line:

```
1.# confluence.home=c:/confluence/data
```

5. Remove the '#' and the space at the beginning of this line, so that Confluence no longer regards the line as a comment. The line should now begin with `confluence.home`
6. Change the Confluence Home directory to the Confluence Home directory you have just created:
   - Place the cursor after the '=' sign.
   - Find the the Confluence Home directory which you created earlier and drag it to TextEdit. This will insert the full path of the folder into the file.
   - Make sure that you did delete the leading '#' character, and that the path after the '=' sign starts with a '/' character. The line should now look something like this:

```
1.confluence.home=/Users/example/confluence/data
```

5. Check the Ports

If you have another application running on your machine which is using the same ports that Confluence uses by default, you may need to change the port which Confluence will use. For example, if you have a Standalone installation of JIRA running on this machine, JIRA might be already using the port which Confluence requests by default.

By default, Confluence listens on port '8080'. If this port is already in use in your installation, follow these instructions to change the ports:

- To change the ports for Confluence Standalone, 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="8080"
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 8080) to ports that are free on your machine.

   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 '8015' and '8090':

Modified conf/server.xml using ports 8015 and 8090

```
<Server port="8015" 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"/>
        ...
```

  To access Confluence in this configuration, point your web browser to http://localhost:8090/.

You will find more information on this page.

6. Select an External Database

This step is optional for users evaluating Confluence. However, if you are installing Confluence for production purposes, this step is mandatory. Please refer to the database requirements listed on our System Requirements topic for help in choosing an external database.

External databases are those listed on our Supported Platforms topic, excluding HSQLDB, which is bundled with Confluence and should not be used in production.

When you have chosen your external database, follow the database setup guide for setting up your database to work with Confluence.

You can learn more about migration from an existing installation or use of the evaluation database here. You will continue to use the Database Setup Guide during the Confluence Setup Wizard. (See step 8 below.)

7. Start Confluence

1. Go to your Confluence Installation directory (created when you unzipped Confluence — see above).
2. Open the `bin` folder and run **OS X - Run Confluence In Background**.
3. Once Confluence is running, open a web browser and visit http://localhost:8080/.
   Hint: If you changed the port earlier, use the port you specified in step 5 above.

8. Next Step is the Confluence Setup Wizard

The Confluence Setup Wizard should appear in your web browser, prompting you to enter your license key. Follow the instructions on the screens, and read more guidelines on the Confluence Setup Wizard.

If the web browser shows an error instead of the Setup Wizard, check the Installation FAQ.

RELATED TOPICS

Change listen port for Confluence Standalone
Adding SSL for Secure Logins and Page Security
Confluence Setup Guide
Confluence Configuration Guide
Documentation Home

Installing Confluence Standalone on Unix or Linux

These instructions apply to:

- The **Standalone Distribution** of Confluence. The Standalone distribution includes Apache Tomcat as the standalone application server. If you want to install an EAR/WAR distribution for deployment on your own existing application server, please refer to the Confluence Installation Guide.
- **Unix, Linux or Solaris** systems. If you are installing Confluence on a Windows or Mac OS X system, please refer to Installing Confluence Standalone Using the Windows Evaluation Installer.

Also, please check the version of Confluence which you are installing. Refer to the documentation home page to verify the latest Confluence version and to find documentation for older versions.

Hint: If you are evaluating Confluence on Unix or you are unsure which version to install, this is the one for you. Just follow the instructions below.

On this page:

- 1. Before you Start
- 2. Install the JDK (Java Development Kit)
- 3. Install X11 Dependencies
- 4. Download and Extract the Confluence Installation File
- 5. Define your Confluence Home Directory
- 6. Check the Ports
- 7. Select an External Database
- 8. Start Confluence
- 9. Next Step is the Confluence Setup Wizard

1. Before you Start

Please check the following points:

1. Ensure that your system meets the minimum requirements to run Confluence. For more information, please read the detailed system requirements.
2. Have your Confluence license key ready. You can obtain a trial, free or commercial license now, or retrieve your existing license key.
3. You must be able to use a command prompt and install Java to continue. If not, please contact your system administrator to assist you or consider the Confluence Hosted evaluation option.
4. Make sure that you use a Gnu version of zip application - Sun/Solaris and AIX are known to have problems with zip, because they use their own (old) versions instead of the Gnu version.

2. Install the JDK (Java Development Kit)

Confluence requires Java 5 (JDK 1.5) or later
Confluence needs JDK 1.5 or newer to be installed on your computer.

- A JRE (Java Runtime Environment) is not enough.
- JDK 6 is the preferred platform, because it is faster and more reliable.
- JDK 1.5 is fine.

OpenJDK is currently **not** supported. A JIRA issue to request support for this JDK has been created.

1. If you are not sure whether you have JDK installed correctly, please confirm by doing the following:
   a. Open a shell console.
   b. Type `echo $JAVA_HOME` in the shell console and then press Enter
   c. View the result:
      - If a line is displayed such as `/opt/jdk1.6.0_12` or `/usr/lib/jvm/java-6-sun`. If you see something like this, then your JDK is installed and properly configured.
      - If nothing is displayed, then you either need to install your JDK or set the `$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 (in step 4 below) by defining this path in your Confluence installation's `setenv.sh` file, usually located in the Confluence `bin` directory.
      - If you have installed a non-Sun JDK and you want to use SSL then you need to install the Sun JSSE package.
2. If you need to install the JDK, follow these instructions:
   - Go to the Java Sun download page.
   - Download the version entitled 'JDK 6 Update XX', where 'XX' stands for some number. (Sun will provide the latest version on that page.)
   - When the download has finished, run the Java installer. Detailed installation instructions are provided on Sun's website.
     At one point, you will be asked to choose an installation directory. Make a note of this directory for use later.

3. Install X11 Dependencies

On Unix-based operating systems, the Java runtime makes use of certain parts of the platform's native X11 graphics libraries. The X Server does not have to be running, but the libraries must be available on the server. Confluence will run on a server that does not have Xlib installed, but parts of the application that manipulate graphics: PDF exports, image thumbnailing, the image gallery macro, CAPTCHA, and the resizing of profile pictures, will fail.

Mac OS X
You do not need to install X11 on Mac OS X, as it has its own graphics libraries.

If X11 is not present, you may see any of the following errors

- "This Confluence installation can not generate thumbnails: no image support in Java runtime"
- "Exception in thread "main" java.lang.UnsatisfiedLinkError: /usr/local/j2sdk1.4.2_09/jre/lib/i386/libawt.so: libXp.so.6: cannot open shared object file: No such file or directory" when exporting a PDF
- "NoClassDefFoundError" when uploading a profile picture

If This Doesn't Help
If you have X11 installed and thumbnailing still does not work, please ensure that you are running Java in headless mode — see the FAQ entitled Confluence doesn't generate thumbnails.

Specific Installation Instructions

Fedora Core

On Fedora Core, you will need to install the `xorg-x11-deprecated-libs` package. (https://bugzilla.redhat.com/bugzilla/show_bug.cgi?id=130239)

Fedora Core 6, RHEL 5

- libXp
- libXp-devel (if you wish to compile against this library)

Debian Linux

On Debian, you will need to instll the following packages (CONF-6411):

```
apt-get install libx11-6 libx11-dev libxt6 libxt6-dbg libxext6 libxtst-dev libxtst6 xlibs-dbg xlibs-dev
```



You'll only need the xlibs-dbg package if you're running an older version of Debian (3.0). It's a dummy package for smoothing the transition to a new set of graphics libraries, so if you can't locate it, you most likely don't need it.

Gentoo Linux

```
emerge libICE libSM libX11 libXext libXp libXt libXtst
```

Solaris 10

Please refer to the following forum for more information.

Ubuntu

Execute the following:

```
> apt-get install libice-dev libsm-dev libx11-dev libxext-dev libxp-dev libxt-dev libxtst-dev
```

Note: '**sudo**' enables you to be superuser for one operation. You will need to supply your user password.

4. Download and Extract the Confluence Installation File

1. If you have not downloaded Confluence already, download the Standalone TAR.GZ file.
2. Use your unzip program to unzip the installation file to a directory such as `/home/jsmith/confluence-2.7.0-std/`.



Most Linux/UNIX users can use any unzip program (such as GNU Tar) to extract the Confluence installer. However, Solaris users should not use the Solaris Tar program due to a known issue associated with its use in extracting Confluence. Use another application such as GNU Tar instead.

For example, change directory to your home directory in Linux and enter the following commands in the shell console:

- `gunzip confluence-<version>-std.tar.gz`
- `tar -xf confluence-<version>-std.tar`
  (where `<version>` refers to the Confluence version you downloaded.)



As usual on Linux/UNIX-based operating systems, avoid using spaces in your directory path.
The directory into which you unzipped the Confluence installation is called the Confluence Installation directory. Next you will define the Confluence Home directory.

5. Define your Confluence Home Directory

Now you need to define the Confluence Home directory. This is where Confluence will store its configuration information, indexes and attachments.



Tip: Another term for 'Home directory' would be 'data directory'.

We suggest using different paths for your installation and home directories. This will facilitate upgrades. Examples of Installation and Home Directories:

- Installation directory: `/usr/local/confluence/`

  If you wish to install or maintain multiple versions of Confluence, you can add a version number to the Confluence installation directory name like `/usr/local/confluence-3.1-std/` and optionally, create the symbolic link `/usr/local/confluence/` that points to `/usr/local/confluence-3.1-std/`

- Home directory: `/usr/local/confluence-data/`

1. Open your Confluence Installation directory (created when you unzipped Confluence — see above).
2. Under the Installation directory, find this file: `confluence/WEB-INF/classes/confluence-init.properties`
3. Open the `confluence-init.properties` file in a text editor.
4. Scroll to the bottom and find this line:

```
1.# confluence.home=c:/confluence/data
```

5. Remove the '#' and the space at the beginning of this line, so that Confluence no longer regards the line as a comment. The line should now begin with `confluence.home`
6. If you decide to change the Confluence Home directory from the default, use an absolute path rather than a symbolic link to specify the path and file name. For example:

```
1.confluence.home=/home/jsmith/confluence-data/
```

6. Check the Ports

If you have another application running on your machine which is using the same ports that Confluence uses by default, you may need to change the port which Confluence will use. For example, if you have a Standalone installation of JIRA running on this machine, JIRA might be already using the port which Confluence requests by default.

By default, Confluence listens on port '8080'. If this port is already in use in your installation, follow these instructions to change the ports:

- To change the ports for Confluence Standalone, 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="8080"
  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 8080) to ports that are free on your machine.

  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 '8015' and '8090':

  Modified conf/server.xml using ports 8015 and 8090

  ```
  <Server port="8015" 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"/>
          ...
  ```

  To access Confluence in this configuration, point your web browser to http://localhost:8090/.

You will find more information on this page.

7. Select an External Database

This step is optional for users evaluating Confluence. However, if you are installing Confluence for production purposes, this step is mandatory. Please refer to the database requirements listed on our System Requirements topic for help in choosing an external database.

External databases are those listed on our Supported Platforms topic, excluding HSQLDB, which is bundled with Confluence and should not be used in production.

When you have chosen your external database, follow the database setup guide for setting up your database to work with Confluence.

You can learn more about migration from an existing installation or use of the evaluation database here. You will continue to use the Database Setup Guide during the Confluence Setup Wizard. (See step 8 below.)

8. Start Confluence

1. Go to your Confluence Installation directory (created when you unzipped Confluence — see above).
2. Under your Confluence Installation directory, open the `bin` directory and run the startup script: `startup.sh`.
3. Once Confluence is running, open a web browser and visit http://localhost:8080/.

   Hint: If you changed the port earlier, use the port you specified in step 6 above.

9. Next Step is the Confluence Setup Wizard

The Confluence Setup Wizard should appear in your web browser, prompting you to enter your license key. Follow the instructions on the screens, and read more guidelines on the Confluence Setup Wizard.

If the web browser shows an error instead of the Setup Wizard, check the Installation FAQ.

RELATED TOPICS

Change listen port for Confluence Standalone
Adding SSL for Secure Logins and Page Security
Confluence Setup Guide
Confluence Configuration Guide
Documentation Home

Installing Confluence Standalone on Windows from Zip File

These instructions apply to:

- The **Standalone distribution** of Confluence. The Standalone distribution includes Apache Tomcat as the standalone application server. If you want to install an EAR/WAR distribution for deployment on your own existing application server, please refer to the Confluence Installation Guide.
- **Windows** systems. For other operating systems please refer to the Confluence Installation Guide.
- Installation via a **zipped download file**. For a simpler installation procedure for evaluation, please consider using the Confluence Evaluation Installer instead.

Also, please check the version of Confluence which you are installing. Refer to the documentation home page to verify the latest Confluence version and to find documentation for older versions.

On this page:

- 1. Before you Start
- 2. Install a Java Development Kit (JDK)
- 3. Download the Confluence Installation File
- 4. Define your Confluence Home Directory
- 5. Check the Ports
- 6. Select an External Database
- 7. Start Confluence
- 8. Next Step is the Confluence Setup Wizard
- 9. Start Confluence automatically on Windows as a Service

1. Before you Start

Please check the following points:

1. Ensure that your system meets the minimum requirements to run Confluence. For more information, please refer to our Supported Platforms topic and for further details, our System Requirements topic.
2. Have your Confluence license key ready. You can obtain a trial, free or commercial license now, or retrieve your existing license key.

2. Install a Java Development Kit (JDK)

Please refer to the Installing Sun JDK for Confluence topic for details on installing a JDK for Confluence. If you are certain that this has already been installed and that the `JAVA_HOME` environment variable has been correctly configured, then proceed to the next step.

3. Download the Confluence Installation File

1. If you have not downloaded Confluence already, download the Standalone zip file.
2. Please check your unzip program before extracting the downloaded zip file. Some archive-extract programs cause errors when unzipping the Confluence zip file. You should use a third-party unzip program like 7Zip or Winzip. If you do not have one, please download and install one before continuing:
    - 7Zip — Recommended. If in doubt, download the '32-bit.exe' version
    - Winzip
3. Use your unzip program to unzip the installation file to a directory such as `c:\confluence`.
    - Do not use spaces in your directory path.

The directory into which you unzipped the Confluence installation is called the Confluence Installation directory. Next you will define the Confluence Home directory.

4. Define your Confluence Home Directory

Now you need to define the Confluence Home directory. This is where Confluence will store its configuration information, indexes and attachments.

Tip: Another term for 'Home directory' would be 'data directory'.

We suggest using different paths for your installation and home directories. This will facilitate upgrades.



Examples of Installation and Home Directories
Installation directory: `c:\confluence\confluence-2.7.0-std`
Home directory: `c:\confluence\data`

1. Open your Confluence Installation directory (created when you unzipped Confluence — see above).
2. Under the Installation directory, find this file: `\confluence\WEB-INF\classes\confluence-init.properties`
3. Open the `confluence-init.properties` file in a text editor such as Notepad.
4. Scroll to the bottom and find this line:

```
1.# confluence.home=c:/confluence/data
```

5. Remove the '#' and the space at the beginning of this line, so that Confluence no longer regards the line as a comment. The line should now begin with `confluence.home`
6. If you decide to change the Confluence Home directory from the default, please note the following:
   - Avoid spaces in the directory path or file name.
   - Use forward slashes '/' to define the path.
     For example:

```
1.confluence.home=c:/data/confluence-home
```

5. Check the Ports

If you have another application running on your machine which is using the same ports that Confluence uses by default, you may need to change the port which Confluence will use. For example, if you have a Standalone installation of JIRA running on this machine, JIRA might be already using the port which Confluence requests by default.

By default, Confluence listens on port '8080'. If this port is already in use in your installation, follow these instructions to change the ports:

- To change the ports for Confluence Standalone, 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="8080"
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 8080) to ports that are free on your machine.



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 '8015' and '8090':

Modified conf/server.xml using ports 8015 and 8090

```
<Server port="8015" 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"/>
        ...
```

To access Confluence in this configuration, point your web browser to http://localhost:8090/.

You will find more information on this page.

6. Select an External Database

This step is optional for users evaluating Confluence. However, if you are installing Confluence for production purposes, this step is mandatory. Please refer to the database requirements listed on our System Requirements topic for help in choosing an external database.



External databases are those listed on our Supported Platforms topic, excluding HSQLDB, which is bundled with Confluence and should

not be used in production.

When you have chosen your external database, follow the database setup guide for setting up your database to work with Confluence.

You can learn more about migration from an existing installation or use of the evaluation database here. You will continue to use the Database Setup Guide during the Confluence Setup Wizard. (See step 8 below.)

7. Start Confluence

1. Go to your Confluence Installation directory (created when you unzipped Confluence — see above).
2. Under your Confluence Installation directory, open the `bin` directory and run the startup script: `startup.bat`. A command prompt window should appear.
   Please do not close this command prompt window. If you do so, Confluence will stop running.

   Troubleshooting
   If the window closes immediately when started, this means that an error is preventing Confluence from starting. To view this error:
   a. Open a command prompt: Click on your 'Start' menu, then click 'Run'. In the Run box, type `cmd` and click 'OK'.
   b. From the command prompt, go to your Confluence Installation directory.
   c. Go into the `bin` subdirectory.
   d. Run `catalina.bat run`.
      You should not run `startup.bat` at this point, because that would still produce a popup window that would close straight away.
   e. Read the error message.
   f. Find the solution to that error in the Installation FAQ.
3. Once Confluence is running, open a web browser and visit http://localhost:8080/.
   If you changed the port earlier, use the port you specified in step 5 above.
   If your web browser window shows an error, try waiting for 30 seconds or so and then refresh the browser page.

8. Next Step is the Confluence Setup Wizard

The Confluence Setup Wizard should appear in your web browser, prompting you to enter your license key. Follow the instructions on the screens, and read more guidelines on the Confluence Setup Wizard.

If the web browser shows an error instead of the Setup Wizard, check the Installation FAQ.

9. Start Confluence automatically on Windows as a Service

Confluence should be run as a service.

RELATED TOPICS

Change listen port for Confluence Standalone
Adding SSL for Secure Logins and Page Security
Confluence Setup Guide
Confluence Configuration Guide
Confluence Documentation Home

Installing Confluence Standalone Using the Windows Evaluation Installer

This page contains instructions for using the **Windows Evaluation Installer** for **Confluence Standalone**.

Important Notice:

- These instructions apply to the **Windows evaluation installer**, which installs an evaluation version of the **Standalone distribution** of Confluence – this package is **NOT RECOMMENDED** for production use. It is strictly for temporary, evaluation purposes only.
- For information about installing other Confluence distributions, such as those for production use or other operating systems, see the instructions for other distributions below.

On this page:

- 1. Before you Begin
- 2. Install a Java Development Kit (JDK)
- 3. Download and Run the Confluence Windows Evaluation Installer
- 4. Wait while the Confluence Evaluation Application Sets Up Confluence
- 5. Run and Configure Confluence
- 6. Follow the Confluence Setup Wizard

- Instructions for Other Distributions
- Support

1. Before you Begin

Please read the following before you begin:

- This evaluation installer should work on most recent computer hardware running Windows. However, for more information, please refer to our Supported Platforms topic and for further details, our System Requirements topic.
- Have your Confluence license key ready. You can obtain a trial, free or commercial license now, or retrieve your existing license key.
- Note that the Windows Evaluation Installer version has only been tested for evaluation purposes. If you wish to set up a production (permanent) instance of Confluence, please use one of the other Confluence packages.

With those conditions satisfied, please continue on to Step 2.

2. Install a Java Development Kit (JDK)

Please refer to the Installing Sun JDK for Confluence topic for details on installing a JDK for Confluence. If you are certain that this has already been installed and that the `JAVA_HOME` environment variable has been correctly configured, then proceed to the next step.

3. Download and Run the Confluence Windows Evaluation Installer

Follow these steps to download and run the Windows evaluation installer and application for Confluence:

1. Download the Windows evaluation installer if you have not already done so. Select the file labelled '**Confluence x.x.x - One-Click Evaluation Installer (EXE)**' with the following filename:

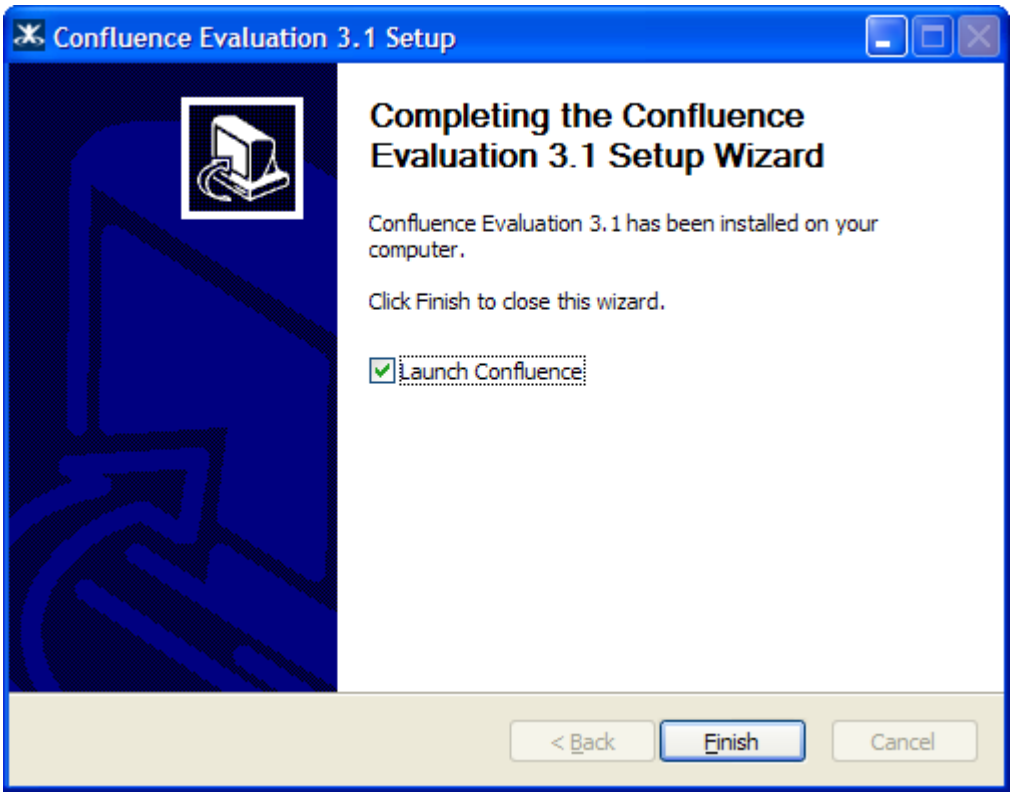
   ```
   confluence-x.x.x-evaluation.exe
   ```

   '**x.x.x**' represents the version of Confluence that will be installed. (You may only see two version numbers separated by a single decimal point.)
   If you are on the Download page and cannot see the installer, click '**Show All**', then the '**Windows**' tab.

2. Save the downloaded file onto your computer.
3. When the download has finished, start the Confluence Windows Evaluation Installer by opening the file you have just downloaded.
   If you see an '**Open File - Security Warning**' message, proceed with running the installer (for example, by clicking '**Run**'), which starts the 'Setup Wizard'.
4. Proceed through each step of the 'Setup Wizard', using the default options at each step (unless you wish to modify the installation location and Windows Start Menu folder options).
5. At the last step of the 'Setup Wizard', leave the 'Launch Confluence' option selected and click the '**Finish**' button.

You can also start Confluence via the Windows '**Start**' -> '**(All) Programs**' folder, by selecting the '**Confluence Evaluation x.x.x**' folder and choosing the '**Confluence x.x.x Evaluation**' menu item.

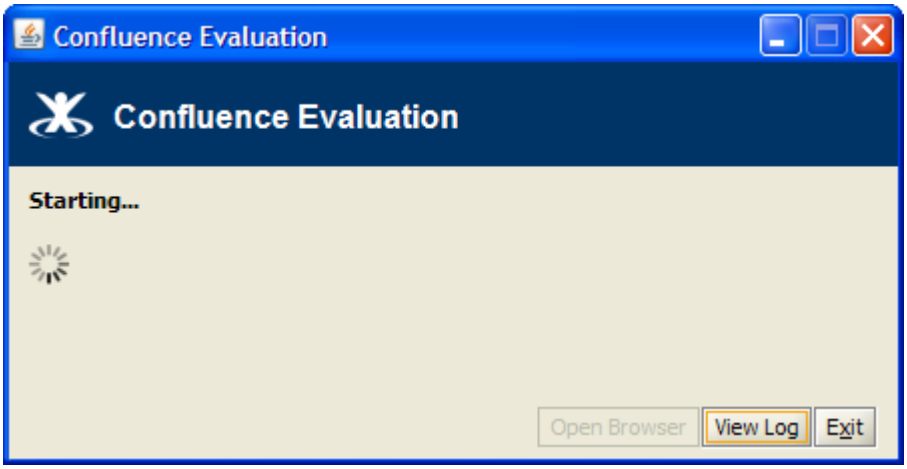
Screenshot: Setup Wizard - Last Step



4. Wait while the Confluence Evaluation Application Sets Up Confluence

Upon starting the Confluence Evaluation application, the 'Confluence Evaluation' dialog box opens and indicates 'Starting'. During this time, the application installs and sets up Confluence for its first run.

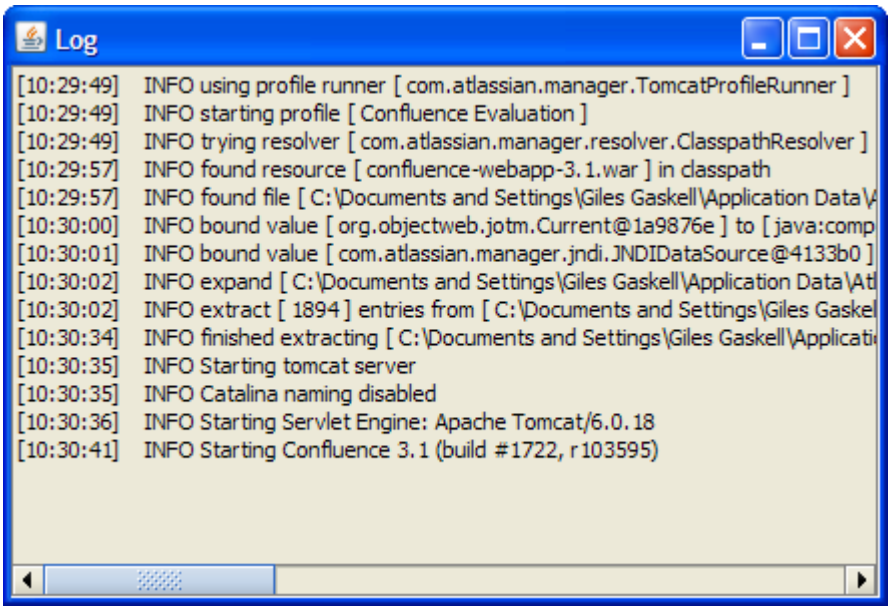
Screenshot: Confluence Evaluation Application – Initial Setup



You will have a short wait while the Confluence Evaluation application installs its files (a process that should take no more than a few minutes to complete).

You can view the Confluence log at any time by clicking the '**View Log**' button. This opens the contents of the Confluence log in a separate window.

Screenshot: Confluence Evaluation Application – Confluence Log Window

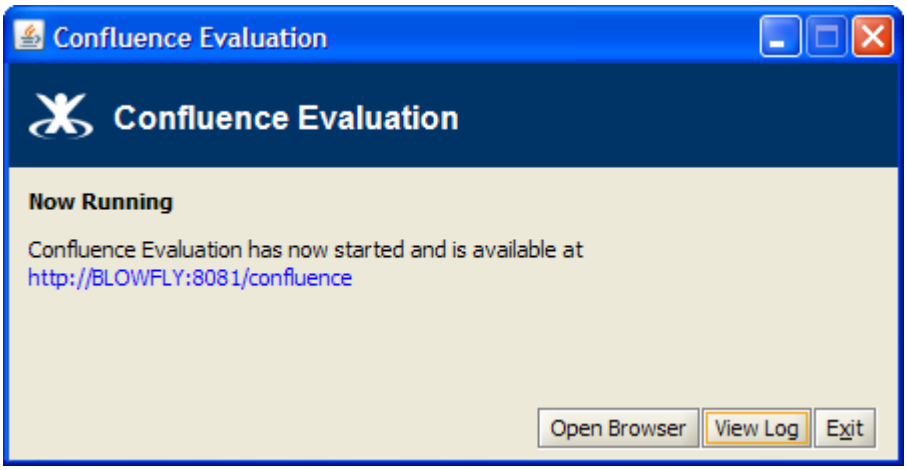


Viewing the Confluence Evaluation application's log during the installation phase shows the progress of the installation process. It also shows the usual activities logged while Confluence is being used.

5. Run and Configure Confluence

Once the installation is complete, the 'Confluence Evaluation' dialog box indicates 'Now Running'.

Screenshot: Confluence Evaluation Application – Ready to Run Confluence



The dialog box indicates the web address of your Confluence Evaluation installation. Clicking on this link or on the **'Open Browser'** button opens Confluence (at this address) directly in your web browser. If your web browser is already open, Confluence will be opened in a new tab.

By default, the Confluence Evaluation application sets itself up to launch from your local computer's network address. You can share the URL shown in the window with people on your local network to allow them to access your instance of Confluence.

You must leave the Confluence Evaluation application's dialog box open to keep Confluence running. You should only close this dialog box or click the **'Exit'** button when you have completely finished using Confluence.

Confluence installation is now complete! You can now move on to configuration, covered below in Step 6.

If you restart your computer (or close the 'Confluence Evaluation' dialog box) and later want to run Confluence again, simply open the **Confluence x.x.x Evaluation** menu item from the Windows Start menu (end of section 2 above) and wait until the 'Confluence Evaluation' dialog box indicates 'Now Running'. The wait will be shorter than the first time as the Confluence Evaluation application's files will have already been installed and it is only the start-up phase that needs to execute.

6. Follow the Confluence Setup Wizard

Now that you have started Confluence, the Confluence Setup Wizard should appear in your web browser, prompting you to enter your license key. Follow the instructions on the setup screens. If you like, you can also read the guide to the Confluence Setup Wizard.

If you encounter any errors, check the installation FAQ.

Instructions for Other Distributions

- If you want to install an EAR/WAR distribution for deployment on your own existing application server, please refer to the Installing the Confluence EAR-WAR Edition.
    - If you want to install Confluence Standalone on a Windows-based system using the zip file for evaluation or production purposes, please refer to Installing Confluence Standalone on Windows from Zip File.
- If you wanted to install Confluence Standalone on a Unix-based system, please refer to Installing Confluence Standalone on Unix or Linux.
- If you wanted to install Confluence Standalone on a Mac OS X-based system:
    - Using the simplified installer for evaluation purposes, please refer to Installing Confluence Standalone Using the Mac OS X Evaluation Installer.
    - Using the zip file for evaluation or production purposes, please refer to Installing Confluence Standalone on Mac OS X from Zip File.

Support

If you have trouble using the Windows Evaluation Installer, please raise an issue in our online support system under the Confluence project.

Installing Confluence Standalone Using the Mac OS X Evaluation Installer

This page contains instructions for using the **Mac OS X Evaluation Installer** for **Confluence Standalone**.



Important Notice:

- These instructions apply to the **Mac OS X evaluation installer**, which installs an evaluation version of the **Standalone distribution** of Confluence – this package is **NOT RECOMMENDED** for production use. It is strictly for temporary, evaluation purposes only.
- For information about installing other Confluence distributions, such as those for production use or other operating systems, see the instructions for other distributions below.

On this page:

- 1. Before you Begin
- 2. Download and Run the Confluence Mac OS X Evaluation Installer
- 3. Wait while the Confluence Evaluation Application Sets Up Confluence
- 4. Run and Configure Confluence
- 5. Follow the Confluence Setup Wizard
- Instructions for Other Distributions
- Support

1. Before you Begin

Please read the following before you begin:

- This evaluation installer should work on all Mac OS X operating systems running on x86 hardware. However, for more information, please refer to our Supported Platforms topic and for further details, our System Requirements topic.
- Have your Confluence license key ready. You can obtain a trial, free or commercial license now, or retrieve your existing license key.
- Note that the Mac OS X Evaluation Installer version has only been tested for evaluation purposes. If you wish to set up a production (permanent) instance of Confluence, please use one of the other Confluence packages.

With those conditions satisfied, please continue on to Step 2.

2. Download and Run the Confluence Mac OS X Evaluation Installer

Follow these steps to download and run the Mac OS X evaluation installer and application for Confluence:

1. Download the Mac OS X evaluation installer if you have not already done so. From the Download page, click '**Show All**', then the '**Mac OS X**' tab. Select the file labelled '**Confluence x.x.x - One-Click Evaluation Installer (DMG)**' with the following filename:

    ```
    Confluence x.x.x Evaluation.dmg
    ```

     '**x.x.x**' represents the version of Confluence that will be installed.

2. Save the downloaded file onto your computer.
3. When the download has finished, start the Confluence Mac OS X Evaluation Installer by opening the file you have just downloaded. This 'mounts' the .dmg file in Mac OS X Finder and creates a Confluence Evaluation 'application' with the name `Confluence x.x.x Evaluation`, where `x.x.x` represents the version of Confluence that will be run.

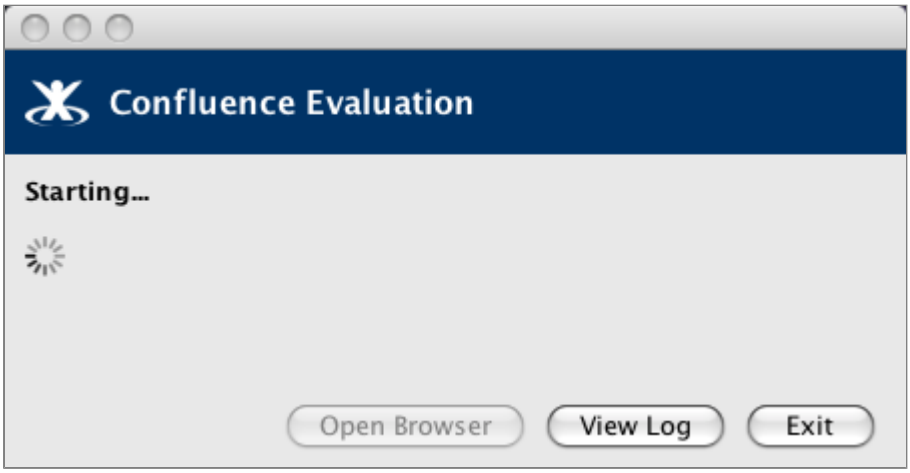
     We recommend that you drag the `Confluence x.x.x Evaluation` file into the `Applications` folder, so that you'll be able to easily launch it next time you start the computer.

4. Start Confluence by opening the newly created **Confluence x.x.x Evaluation** application.

3. Wait while the Confluence Evaluation Application Sets Up Confluence

Upon starting the Confluence Evaluation application, the 'Confluence Evaluation' dialog box opens and indicates 'Starting'. During this time, the application installs and sets up Confluence for its first run.

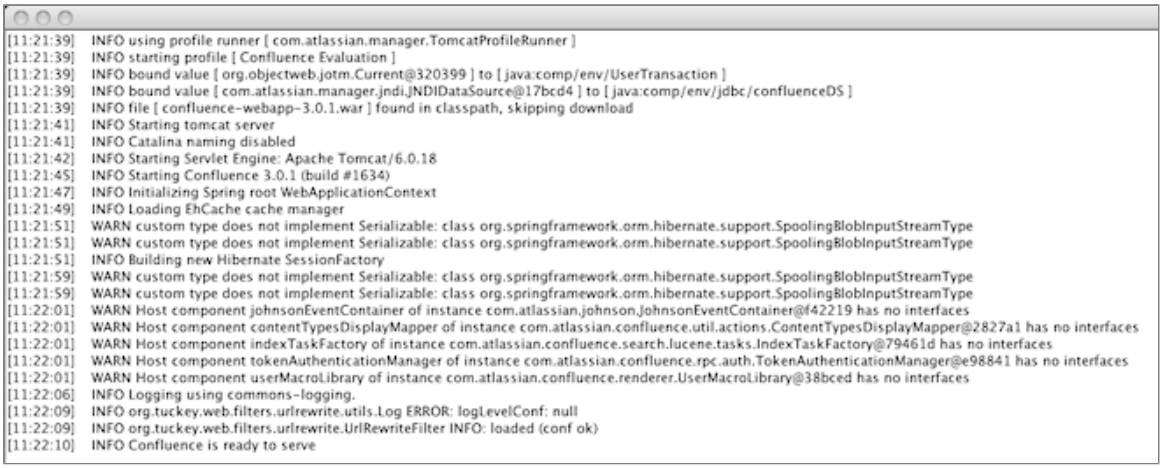
Screenshot: Confluence Evaluation Application – Initial Setup



You will have a short wait while the Confluence Evaluation application installs its files (a process that should take no more than a few minutes to complete).

You can view the Confluence log at any time by clicking the '**View Log**' button. This opens the contents of the Confluence log in a separate window.

Screenshot: Confluence Evaluation Application – Confluence Log Window

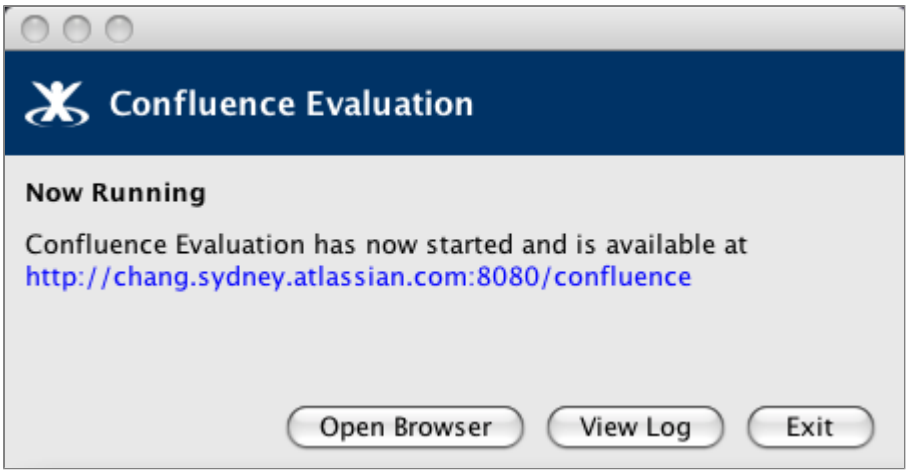


Viewing the Confluence Evaluation application's log during the installation phase shows the progress of the installation process. It also shows the usual activities logged while Confluence is being used.

4. Run and Configure Confluence

Once the installation is complete, the 'Confluence Evaluation' dialog box indicates 'Now Running'.

Screenshot: Confluence Evaluation Application – Ready to Run Confluence



The dialog box indicates the web address of your Confluence Evaluation installation. Clicking on this link or on the '**Open Browser**' button opens Confluence (at this address) directly in your web browser. If your web browser is already open, Confluence will be opened in a new tab.

By default, the Confluence Evaluation application sets itself up to launch from your local computer's network address. You can share the URL shown in the window with people on your local network to allow them to access your instance of Confluence.

You must leave the Confluence Evaluation application's dialog box open to keep Confluence running. You should only close this dialog box or click the '**Exit**' button when you have completely finished using Confluence.

Confluence installation is now complete! You can now move on to configuration, covered below in Step 5.

If you restart your computer (or close the 'Confluence Evaluation' dialog box) and later want to run Confluence again, simply open the **Confluence x.x.x Evaluation** file from the Applications folder (end of section 2 above) and wait until the 'Confluence Evaluation' dialog box indicates 'Now Running'. The wait will be shorter than the first time as the Confluence Evaluation application's files will have already been installed and it is only the start-up phase that needs to execute.

5. Follow the Confluence Setup Wizard

Now that you have started Confluence, the Confluence Setup Wizard should appear in your web browser, prompting you to enter your license key. Follow the instructions on the screens that follow, reading the instructions for the Confluence Setup Wizard.

If you encounter any errors, check the Installation FAQ page.

Instructions for Other Distributions

- If you want to install an EAR/WAR distribution for deployment on your own existing application server, please refer to the Installing the Confluence EAR-WAR Edition.
    - If you want to install Confluence Standalone on a Mac OS X-based system using the tar.gz (zip) file for evaluation or production purposes, please refer to Installing Confluence Standalone on Mac OS X from Zip File.
- If you wanted to install Confluence Standalone on a Unix-based system, please refer to Installing Confluence Standalone on Unix or Linux.
- If you wanted to install Confluence Standalone on a Windows-based system:
    - Using the simplified installer for evaluation purposes, please refer to Installing Confluence Standalone Using the Windows Evaluation Installer.
    - Using the zip file for evaluation or production purposes, please refer to Installing Confluence Standalone on Windows from Zip File.

Support

If you have trouble using the Mac OS X Evaluation Installer, please raise an issue in our online support system under the Confluence project.

Installing the Confluence EAR-WAR Edition

The Confluence EAR-WAR distribution is intended for deployment into an existing J2EE application server.

> To use this method of installation, you or your application server administrator must already know how to deploy a web application on the application server of choice. If not, please use the Confluence Standalone distribution instead.

On this page:

- System Requirements for Confluence EAR-WAR Distribution
- Follow the Application-Specific Instructions

System Requirements for Confluence EAR-WAR Distribution

1. Please check the Confluence system requirements.
2. In addition to the above requirements, the EAR-WAR distribution requires an application server. For more information on Confluence's supported application servers, please refer to our Supported Platforms page.
3. If deploying as an unexploded WAR, Ant 1.3 or later is required. This is bundled with the war download.
4. Confluence, the database and application server must use the same character encoding. UTF-8 is recommended.
5. If you are installing Confluence EAR-WAR on Unix, Linux or Solaris, please ensure that the required X11 libraries are installed.

> **Support**
> If you deploy Confluence on an unsupported server, server-related issues cannot be covered by Atlassian technical support. You can try the user forums for assistance instead.

Follow the Application-Specific Instructions

- Installing Confluence EAR-WAR on Tomcat
- Installing Confluence EAR-WAR on Weblogic
- Installing Confluence EAR-WAR on Websphere
- Installing Confluence EAR-WAR on Resin

Installing Confluence EAR-WAR on Tomcat

On this page

- Step 1. Check the Known Issues for Tomcat
- Step 2. Download and Extract EAR-WAR Installation File
- Step 3. Check for Patches
- Step 4. Check for Application Server Issues
- Step 5. Review Application Server Memory Allocation
- Step 6. Configure `confluence-init.properties`
- Step 7. Edit Tomcat Context Descriptors
- Step 8. Add UTF-8 Encoding
- Step 9. Optional: Run Tomcat on a Different Port
- Step 10. Optional: Running Confluence as a Windows Service
- Proceed Through the Confluence Setup Wizard
  - TROUBLESHOOTING
  - RELATED TOPICS

Step 1. Check the Known Issues for Tomcat

Read through the Known Issues for Apache Tomcat.

Step 2. Download and Extract EAR-WAR Installation File

This section gives detailed instructions for installing Confluence EAR-WAR edition on an **Apache Tomcat 5.5, or 6** server.

1. Download the Confluence EAR/WAR zip file. (You need to click the '**Show all**' link to see the EAR/WAR zip file.)
2. Please check your unzip program before extracting the downloaded zip file. Some archive-extract programs cause errors when unzipping the Confluence zip file:
   - Windows users must avoid the Windows built-in unzip utility, as it doesn't extract all the files. Use a third-party unzip program like 7Zip or Winzip.
   - Solaris users will need to use GNU tar to handle the long file names.
3. Extract the downloaded zip file.

> Tomcat users, take care not to unzip the Confluence installation into your Tomcat `webapps` folder, as this may cause Confluence to be deployed more than once. It may cause a Cluster Panic error.

4. You have now unzipped your *Confluence installation directory*, which should contain the version number e.g. `confluence-2.10.0` or `confluence-2.10.2`. This directory will be later referred to as the *Confluence installation directory*. Inside is a `confluence` subdirectory, referred to later as the *(Exploded) Confluence WAR directory*. Record the absolute path to the *Confluence WAR directory*.

Step 3. Check for Patches

Review the Release Notes for your Confluence version and apply any patches listed.

Step 4. Check for Application Server Issues

Note that these are optionals and you do not need to go through them to make Confluence EAR/WAR run for the first time. They are listed here only to assist you in getting the best matching configuration for your production requirement.

Step 5. Review Application Server Memory Allocation

Confluence requires a maximum heap allocation (Xmx) of at least 256 MB for normal operation. See Increasing Application Server Memory.

Do not configure a heap allocation so large that it does not allow enough remaining physical memory for your operating system and other applications on the server. The heap allocation should be large enough for Confluence, but not so large that the memory would be paged to disk during normal operation.

Step 6. Configure `confluence-init.properties`

1. Inside the *Confluence installation directory*, edit `...confluence/WEB-INF/classes/confluence-init.properties` in a text editor.
2. Now define your Confluence Home directory, by setting the `confluence.home` property to a directory of your choosing.
   We suggest using different paths for your installation and home directories. This will facilitate upgrades. This is the directory that will contain all of Confluence's configuration, backup and attachment files.
   Tip: Another term for 'Home directory' would be 'data directory'.

Step 7. Edit Tomcat Context Descriptors

1. Create a file called `confluence.xml` in your Tomcat installation's `conf/Catalina/localhost` directory. (If you have set up a different hostname for your tomcat instance, please specify that instead of *localhost*.) For Tomcat 6, you must create the Catalina and localhost directories.
2. Open your new `confluence.xml` and add these lines:

```
1.<Context path="/confluence" docBase="c:/applications/confluence-2.9.2/confluence" debug="0"
reloadable="true">
2.</Context>
```

   To run Confluence *without* a context path of "/confluence", change the path in the Context tag to an empty string (""). If this is Tomcat, name the file ROOT.xml rather than confluence.xml. If you wish to change the context path to a different name, change both the context path and the name of the xml file.

3. For *docBase*, specify the value you noted down earlier.
4. Restart Tomcat, and Confluence should be accessible under `/confluence/` on your Tomcat server.
5. Follow the link below to proceed with the setup wizard.

Step 8. Add UTF-8 Encoding

1. Edit `conf/server.xml` and find the line where the `Coyote HTTP Connector` is defined. It will look something like this, possibly with more parameters:

```
1.<Connector port="8080"/>
```

2. Add a `URIEncoding="UTF-8"` property to the connector:

```
1.<Connector port="8080" URIEncoding="UTF-8"/>
```

Step 9. Optional: Run Tomcat on a Different Port

See Running Tomcat on a Different Port.

Step 10. Optional: Running Confluence as a Windows Service

Confluence can be run as a service.

Proceed Through the Confluence Setup Wizard

When you visit the Confluence application in your browser, the Confluence Setup Wizard should appear. The first screen is a prompt for you to enter your license key. Follow the instructions on the screens, and read more guidelines on the Confluence Setup Wizard.

TROUBLESHOOTING

Installation FAQ

RELATED TOPICS

Confluence Setup Guide
Start Confluence automatically on Windows as a Service
Confluence Configuration Guide
Confluence Home Directory
Confluence Cluster Installation
Fedora or RHEL/Centos Install Guide — *Contributed by a Confluence user*
Known Issues with Enterprise or Webhosting environments
Pull down RSS Feeds or use the Repository plugin through a web proxy
Setting the JAVA_HOME Variable in Windows
Setup a mail session in standalone version

Confluence Documentation Home
Confluence FAQ
Installing Confluence EAR-WAR on Weblogic
Configuring a MySQL Datasource in Apache Tomcat

Known Issues for Apache Tomcat

On this page:

- Supported Application Servers
- Tomcat Documentation
- Known Issues
    - RELATED TOPICS

Supported Application Servers

Check the list of supported application servers on the Supported Platforms topic.

Tomcat Documentation

An excellent resource for Tomcat configuration is the Apache documentation.

Known Issues

 Setup Fails Creating MySQL Schema due to Tomcat Incompatibility

 Slow Page Rendering of Large Pages due to HTTP POST Limitations

 Unable to start Tomcat after Confluence user management delegation to JIRA

 Fix 'Not supported by BasicDataSource' Setup or Startup Error

 Confluence Does Not Start due to NullPointerException in FelixOsgiContainerManager

 NotSerializableException on shutdown

 Confluence Menus do not work, or confluence fails to startup, when running in the same Tomcat as Jira 4.0, 4.0.1 or Crowd 2.0.x

 Confluence Can't Start and Doesn't Create Logfiles due to CATALINA_HOME Being Set

 Confluence Startup Referencing a Different Tomcat

 Troubleshooting Application Servers

RELATED TOPICS

Running Confluence behind Apache
Configuring a MySQL Datasource in Apache Tomcat

Installing Confluence EAR-WAR on Weblogic

This documentation provides a guide to install Confluence on Weblogic 9.2 with basic settings and an external database. Examples on the page are given for Windows. Instructions for Linux should be similar to this guide.

On this page:

- Step 1. Check the Known Issues
- Step 2. Create a Weblogic Domain
- Step 3. Configuring Confluence
- Step 4. Install Confluence with Weblogic and an External Database via Direct JDBC
- Step 5. Starting Confluence

Step 1. Check the Known Issues

Peruse the Known Issues for WebLogic.

Step 2. Create a Weblogic Domain

Assuming that your Weblogic has been installed in `C:\bea`, follow these steps:

1. Open Weblogic Configuration Wizard:

```
1.C:\bea\weblogic92\common\bin\consolew.exe /NOWINDOW /c
"C:\bea\weblogic92\common\bin\config.cmd
```

OR

```
1.Start >> Programs >> BEA Products >> Tools >> Configuration Wizard
```

2. Choose `Create New Weblogic Domain`
3. Click next all the way.
4. You will be prompted to enter your Weblogic's admin account, enter and click next.
5. Sun's JDK is recommended.
6. Click next until prompted for a domain name. Enter `confluence`.

More information can be found in BEA documentation.

Step 3. Configuring Confluence

1. Download the Confluence EAR/WAR zip file. (You need to click the '**Show all**' link to see the EAR/WAR zip file.)
2. Check your unzip program before extracting the downloaded zip file. Some archive-extract programs cause errors when unzipping the Confluence zip file:
   - Windows users must avoid the Windows built-in unzip utility, as it doesn't extract all the files. Use a third-party unzip program like 7Zip or Winzip.
   - Solaris users will need to use GNU tar to handle the long file names.
3. Extract the downloaded zip file.
4. You have now unzipped your *Confluence installation directory*, which should contain the version number e.g. `confluence-3.0`. This directory will be later referred to as the *Confluence installation directory*. Inside is a `confluence` subdirectory, referred to later as the *(Exploded) Confluence WAR directory*. Record the absolute path to the *Confluence WAR directory*.
5. Modify its **confluence-init.properties** and set Confluence Home.
   a. Inside the *Confluence installation directory*, edit `...confluence/WEB-INF/classes/confluence-init.properties` in a text editor.
   b. Now define your Confluence Home directory, by setting the `confluence.home` property to a directory of your choosing. We suggest using different paths for your installation and home directories. This will facilitate upgrades. This is the directory that will contain all of Confluence's configuration, backup and attachment files.

      Tip: Another term for 'Home directory' would be 'data directory'.
6. Replace the `confluence/WEB-INF/web.xml` with `confluence/WEB-INF/web.servlet2-4.xml`. See Known Issues for WebLogic.
7. Create a file called `weblogic.xml` in `<confluence install directory>\confluence\WEB-INF` with the following content:

```
01.<?xml version="1.0" encoding="UTF-8"?>
02.<weblogic-web-app
03.   xmlns="http://www.bea.com/ns/weblogic/90"
04.   xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
05.   xsi:schemaLocation="http://www.bea.com/ns/weblogic/90
http://www.bea.com/ns/weblogic/90/weblogic-web-app.xsd">
06.          <jsp-descriptor>
07.                   <precompile>false</precompile>
08.          </jsp-descriptor>
09.          <container-descriptor>
10.                           <servlet-reload-check-secs>-1</servlet-reload-check-secs>
11.                           <prefer-web-inf-classes>true</prefer-web-inf-classes>
12.          </container-descriptor>
13.          <context-root>yourContextPathName</context-root>
14.</weblogic-web-app>
```

From Confluence 3.1 onwards it is essential to specify `prefer-web-inf-classes` due to CONF-17827:

```
1.<container-descriptor>
2.       <prefer-web-inf-classes>true</prefer-web-inf-classes>
3.</container-descriptor>
```

This parameter allows Weblogic to prefer the libraries from Confluence and in many cases helps to solve many libraries conflict. However, depending on the library version you have in your installation this parameter can also cause Confluence fail to start. If that happens you need to figure out which package is causing the conflict. A related bug reports for this issue has been filed in CONF-15163 and in CONF-15332, however these bugs have been fixed for version 3.1 of Confluence.

If you discover that Weblogic still uses its own library despite of setting the `prefer-web-inf-classes` field to true, you may need to tell it explicitly to load the correct jar version. Please refer to **Forcing Weblogic to load certain Jar files** in Weblogic Configuration Tips.

The full `weblogic.xml` syntax is described in BEA's documentation.



Due to a bug in WebLogic 9.2's implementation of the J2EE ServletRequest certain deployment paths can lead to failed resource requests. You must avoid deploying Confluence to any of the following contexts:

- `/confluence`
- `/resources`
- `/download`

If you do not set the `context-root` in `weblogic.xml` Weblogic will define the context root automatically to `/confluence`, which is not what we want. See CONF-13712 for further details.
This problem has been fixed in Weblogic version 9.2 MP2.



Confluence must be deployed as an exploded war file to WebLogic.

Step 4. Install Confluence with Weblogic and an External Database via Direct JDBC

1. Create a Weblogic Domain eg. `your_domain_name`
2. Go to `C:\bea\user_projects\domains\your_domain_name\lib` and place the JDBC driver here.
3. Start Weblogic Admin Server.
4. Access Weblogic Admin console: http://localhost:7001/console
5. On the left hand panel, click **Lock & Edit**
6. On the left hand panel, look up **Domain Structure > your_domain_name > Deployments**
7. Click on the **Install** button and look up for your Confluence installation folder. Eg. `c:\atlassian\confluence\war\confluence-3\confluence`.
8. Once you found the folder, click on the radio button and click Next all the way.
9. Name the deployment as **confluence**.
10. Click Finish.
11. On the left hand panel, click **Activate Changes**.

Step 5. Starting Confluence

1. Start your Weblogic Admin Server
2. Go to **Deployments** and locate your installed Confluence
3. Tick Confluence
4. Click on Start button and choose **Servicing all requests**
5. Click Yes

To access Confluence go to http://localhost:7001/yourContextPathName and you can proceed to Confluence Setup guide Step number 2.

RELATED TOPICS

Known Issues for WebLogic

Known Issues for WebLogic

On this page:

- Installation Instructions
- Supported Application Servers
- Important Configuration Information
- WebLogic Configuration tips
  - Forcing Weblogic to load certain Jar files
  - Disabling JSP reload checks
  - Disabling servlet reload checks
  - Avoiding JSP recompiles on redeploy
  - Configuring Character Encoding
  - Configuring HTTP BASIC Auth
- Knowledge Base Articles

Installation Instructions



Check out the step by step guide for help installing Weblogic.

Supported Application Servers

Check the list of supported application servers on the Supported Platforms topic.

Important Configuration Information

Configuring weblogic.xml

If you are seeing the following error -- or other NoSuchMethodErrors -- then this tip is for you!

```
1.java.lang.NoSuchMethodError:
org.objectweb.asm.ClassVisitor.visit(ILjava/lang/String;Ljava/lang/String;[Ljava/lang/String;Ljava/la
```

It appears that WLS 9 ships with an incompatible version of org.objectweb.asm.* packages which clash with the ones required by Hibernate. You can fix this by adding to the weblogic.xml deployment descriptor the settings.

```
1.<container-descriptor>
2.     <prefer-web-inf-classes>true</prefer-web-inf-classes>
3.</container-descriptor>
```

Configuring web.xml

In order to run Confluence under Weblogic 9.2, you need to use a 2.4 servlet compatible version of `web.xml`.

This has been bundled in the main Confluence download as:

```
1.confluence/WEB-INF/web.servlet2-4.xml
```

To enable it, simply remove the default `confluence/WEB-INF/web.xml` and then rename `web.servlet2-4.xml` to `web.xml`. Now restart.

If you do not use the 2.4 servlet compatible version, Confluence will not be able to display its pages in a nice URL format and will render blank instead.

You will also see this error in the log file:

```
[confluence.util.velocity.VelocityUtils] getRenderedTemplate Error occurred rendering template:
/decorators/main.vmd
org.apache.velocity.exception.MethodInvocationException: Invocation of method 'requireResource' in
class $Proxy566 threw exception java.lang.NullPointerException @ /decorators/main.vmd[1,25]
 at org.apache.velocity.runtime.parser.node.ASTMethod.execute(ASTMethod.java:286)
```

Confluence Deployment Context

Due to a bug in WebLogic's implementation of the J2EE `ServletRequest`, certain deployment paths can lead to failed resource requests. Hence, you should avoid deploying Confluence to any of the following contexts:

- /confluence
- /resources
- /download

See the JIRA issue CONF-13712 for further details.

Exploded War

Confluence must be deployed as an *exploded* war file to WebLogic. Otherwise you may experience errors as such:

```
java.lang.IllegalStateException: Spring Application context has not been set
  at bucket.container.SpringContainerContext.getComponent(SpringContainerContext.java:98)
```

WebLogic Configuration tips



Enabling logging
Confluence sends its log output to standard out, so by default Weblogic does not record it. To redirect Confluence's log output to a file follow these instructions Redirecting System.out and System.err to a File

Forcing Weblogic to load certain Jar files

If you discover that despite using `prefer-web-inf-classes` parameter Weblogic still loads its own library classes, you can use tell Weblogic to load the required jar file. Copy the required jar file from **`<confluence install>/confluence/WEB-INF/lib`** to **`C:\bea\user_projects\domains\<confluence_domain>\lib`** (or Linux equivalent).

Then, edit the `setDomainEnv.cmd/sh` file in the **C:\bea\user_projects\domains\<confluence_domain>\bin** folder (or Linux equivalent).

Add the following:

```
1.set
PRE_CLASSPATH=%PRE_CLASSPATH%;C:\bea\user_projects\domains\confluence_domain>\lib\the_required_jar_fi
```

And restart your deployment.

Disabling JSP reload checks

Performance can be considerably affected by Weblogic's tendency to scan JSPs for changes on every page load.
This can be disabled by setting **pageCheckSeconds to -1**.

Disabling servlet reload checks

Similar to the JSP reload checks, Weblogic will scan servlets for modifications every second (by default), hurting performance.
This can be prevented by setting the **servlet-reload-check-secs element to -1**.

Avoiding JSP recompiles on redeploy

If you Weblogic instance is often restarted, performance can be improved by explicitly setting a directory to save compile JSPs to (the workingDir parameter), and setting **precompile to 'false'**.

Configuring Character Encoding

A user has reported that the following configuration can ensure that Chinese character sets are displayed correctly. To correct this, add the following entry to the weblogic.xml file.

```
<weblogic-web-app>
                              <jsp-descriptor>
                                   <jsp-param>
                                        <param-name>encoding</param-name>
                                        <param-value>UTF-8</param-value>
                                   </jsp-param>
                                   <jsp-param>
                                        <param-name>compilerSupportsEncoding</param-name>
                                        <param-value>false</param-value>
                                   </jsp-param>
                              </jsp-descriptor>
                         </weblogic-web-app>
```

Most of this page duplicates JIRA documentation - you might also like to read
http://www.atlassian.com/software/jira/docs/latest/servers/weblogic.html.

Configuring HTTP BASIC Auth

If you wish to use HTTP BASIC Auth for Confluence (this is currently required for authenticated RSS feeds), you will need to modify your config.xml
for your WebLogic domain as by default WebLogic will direct all BASIC Auth requests to its admin console.

Modify {WEBLOGIC_DOMAIN_LOCATION}\config\config.xml to include

```
<enforce-valid-basic-auth-credentials>false</enforce-valid-basic-auth-credentials>
```

to correct this.

Knowledge Base Articles

Site Import Fails on Weblogic when Using PostgreSQL

Unable to Deploy Evaluation Database on Weblogic

Troubleshooting Application Servers

RELATED TOPICS

Installing Confluence EAR-WAR on Weblogic

Installing Confluence EAR-WAR on Websphere

This documentation provides a guide to install Confluence on Websphere 6.1 with basic settings and an external database.

On this page:

- Step 1. Configuring the Confluence War File
- Step 2. Install Confluence in Websphere
- Step 3. Starting Confluence
- Step 4. Modifying Files in an Exploded War

After installing Websphere, follow the instructions below.

Step 1. Configuring the Confluence War File

1. Download the Confluence EAR/WAR zip file. (You need to click the '**Show all**' link to see the EAR/WAR zip file.)

2. Check your unzip program before extracting the downloaded zip file. Some archive-extract programs cause errors when unzipping the Confluence zip file:
    - Windows users must avoid the Windows built-in unzip utility, as it doesn't extract all the files. Use a third-party unzip program like 7Zip or Winzip.
    - Solaris users will need to use GNU tar to handle the long file names.
3. Extract the downloaded zip file.
4. You have now unzipped your *Confluence installation directory*, which should contain the version number e.g. `confluence-2.10.0` or `confluence-2.10.2`. This directory will be later referred to as the *Confluence installation directory*. Inside is a `confluence` subdirectory, referred to later as the *(Exploded) Confluence WAR directory*. Record the absolute path to the *Confluence WAR directory*.
5. Modify its **confluence-init.properties** and set Confluence Home.
    a. Inside the *Confluence installation directory*, edit `...confluence/WEB-INF/classes/confluence-init.properties` in a text editor.
    b. Now define your Confluence Home directory, by setting the `confluence.home` property to a directory of your choosing. We suggest using different paths for your installation and home directories. This will facilitate upgrades. This is the directory that will contain all of Confluence's configuration, backup and attachment files.

    Tip: Another term for 'Home directory' would be 'data directory'.
6. For Confluence versions 2.10.1, 2.10.2 and 2.10.3, delete the package-scanner-0.7.2.jar from the WEB-INF/lib directory, and place package-scanner-0.7.jar in its place. See Known Issues for Websphere.
7. Add your database driver jar file to the WEB-INF/lib directory. If your driver already exists in the main classpath and is used for other applications, you can skip this step. Check the Database Installation Instructions. If you're not already committed, we recommend Postgres.

    The built-in evaluation database may not work on a Websphere deployment.

8. Build the war file.

Windows
a. Open the command line prompt b. Navigate to the *Confluence installation directory* c. Run the command `build war`

Linux
a. Open a terminal window b. Navigate to the *Confluence installation directory* c. Add executable mode to `build.sh` — `chmod +x build.sh` d. Run `build.sh` e. The WAR file is now located in the `dist` subfolder in your *Confluence installation directory*

Step 2. Install Confluence in Websphere

1. Start Websphere Admin Server.
2. Access Websphere Admin console: https://localhost:9043/ibm/console
3. On the left hand panel, click **Servers > Application Servers.**
4. Select your server.
5. Choose **Allow** for **Access to internal server classes**.
6. Choose **Multiple** for **Classloader Policy**.
7. On the left hand panel, click **Applications > Enterprise Applications.**
8. Import your war. Note your context path and choose all the defaults.
9. Once the war file is imported, select it, then choose **Class loading and update detection**.
10. Select **Classes loaded with application class loader first** and **Single class loader for application**.
11. Restart Websphere to make sure that all settings are picked up.

Step 3. Starting Confluence

1. Start your Websphere Server.
2. Go to **Applications > Enterprise Applications** and locate your installed Confluence.
3. Tick the box next to the Confluence war file and click **Start**.
4. Access Confluence at http://localhost:9080/contextPath (chosen above)

To access Confluence go to http://localhost:9080/contextPath

Step 4. Modifying Files in an Exploded War

It's a bit inconvenient to modify files - add patches or configure velocity or xml files - if you have to rebuild and redeploy the war file. Websphere requires deploying Confluence as a built .war archive originally. After that, however, you can drop files in the exploded war directory and restart the server to pick up changes (just the application isn't enough, given classloading issues). For the case of Velocity (vm or vmd) files or jsp files, you don't even need to reload the application - just refresh your page!

A typical exploded war directory might be found in a path similar to:

```
/usr/lib/IBM/WebSphere/AppServer/profiles/AppSrv03/installedApps/serverName-laptopNode01Cell/confluen
```

Keep in mind that this isn't a technique for deploying war files, but rather for modifications *after* deploy.

RELATED TOPICS

Known Issues for Websphere
Websphere JVM Tuning

Known Issues for Websphere

On this page:

- Installation Instructions
- Supported Application Servers
- Classloader settings
    - Confluence 3.0
- Oracle and Data Sources
- Configuring Encoding
- Known Issues

Installation Instructions



See the Installing Confluence EAR-WAR on Websphere page for step-by-step instructions.

Supported Application Servers

Check the list of supported application servers on the Supported Platforms topic.

Classloader settings

Confluence 3.0

Confluence 3.0 uses the Plugins 2 framework. The default classloader settings for WebSphere interferes with Plugins 2 dependencies.

The solution is to ensure your WebSphere application is configured with the following settings:

1. Application classloader policy to SINGLE
2. Application classloader mode to child-first

Please refer to Running on WebSphere 6.1 for more details.

Oracle and Data Sources

Please use a direct JDBC connection instead. See CONF-3580 for more details.

Configuring Encoding

I am having trouble with the £ and € characters

Confluence allows you to specify which character encoding should be used for its requests. To use the £ and € characters, you will need to use the "UTF-8" option. However, in some cases, WebSphere will encode the characters using ISO-8859-1 before they get to Confluence. If you have set the Confluence character encoding to "UTF-8" and are having problems, have a look at http://www-306.ibm.com/software/globalization/j2ee/encoding.jsp for an explaination of WebSphere and Character encoding. In particular, the -Dclient.encoding.override=UTF-8 system property allows you to override WebSpheres defaults.

Known Issues

 NoClassDefFoundError due to Path Size Limit in Windows

 Personal Space Results in 500 Error due to Siteminder Configuration

 Redirected to Dashboard on Login Due to Page Titles with Spaces (Websphere)

 No Trusted Certificate Found when accessing plugin repository (Websphere)

 Unable to Install Confluence or Load Plugins on Websphere 6.1

 Formatting Buttons Missing on Toolbar of Rich Text Editor due to Application Security on Websphere

 Unable to Install Oracle on Websphere due to Classloader Preferences

 Unable to Configure Oracle Data Source in Websphere

 Troubleshooting Application Servers

MalformedInputException when rebuilding the Did You Mean search index

Problems Loading Plugins due to org.osgi.framework.BundleException

Installing Confluence EAR-WAR on Resin

This documentation provides a guide to install Confluence on Resin with basic settings and an external database.

On this page:

- Step 1. Check the Known Issues
- Step 2. Configuring the Confluence War File
- Step 3. Install Confluence in Resin
- Step 4. Starting Confluence

Step 1. Check the Known Issues

Check the Known Issues for Resin 3.x page before proceeding.

Step 2. Configuring the Confluence War File

1. Download the Confluence EAR/WAR zip file. (You need to click the '**Show all**' link to see the EAR/WAR zip file.)
2. Please check your unzip program before extracting the downloaded zip file. Some archive-extract programs cause errors when unzipping the Confluence zip file:
    - Windows users must avoid the Windows built-in unzip utility, as it doesn't extract all the files. Use a third-party unzip program like 7Zip or Winzip.
    - Solaris users will need to use GNU tar to handle the long file names.
3. Extract the downloaded zip file.
4. You have now unzipped your *Confluence installation directory*, which should contain the version number e.g. `confluence-2.10.0` or `confluence-2.10.2`. This directory will be later referred to as the *Confluence installation directory*. Inside is a `confluence` subdirectory, referred to later as the *(Exploded) Confluence WAR directory*. Record the absolute path to the *Confluence WAR directory*.
5. Modify **confluence-init.properties** and set Confluence Home.
    a. Inside the *Confluence installation directory*, edit `/WEB-INF/classes/confluence-init.properties` in a text editor.
    b. Now define your Confluence Home directory, by setting the `confluence.home` property to a directory of your choosing. We suggest using different paths for your installation and home directories. This will facilitate upgrades. This is the directory that will contain all of Confluence's configuration, backup and attachment files.

    Tip: Another term for 'Home directory' would be 'data directory'.
6. Add your database driver jar file to the WEB-INF/lib directory. If your driver already exists in the main classpath and is used for other applications, you can skip this step. Check the Database Installation Instructions. If you're not already committed, we recommend Postgres.
7. Rename /WEB-INF/web.resin3.xml to web.xml. Keep the prior web.xml as a backup.

Step 3. Install Confluence in Resin

1. Move the *Confluence war directory* to Resin's `webapps` folder. You are deploying the exploded war file which is supported in Resin.

    If you'd like to run Confluence without a context path, rename the folder to ROOT and remove the existing ROOT folder.

Step 4. Starting Confluence

1. Start Resin. Use the shell script in the bin folder or resin.jar in the lib folder.

To access Confluence go to http://localhost:8080/contextPath

Known Issues for Resin 3.x

On this page

- Installing Confluence on Resin
- Check that your version is supported
- Configure Resin to use a XSD-aware parser
- Resin and SOAP
- Configure web.xml
- Knowledge Base Articles

Installing Confluence on Resin

Check out the Resin step-by-step installation instructions

Check that your version is supported

Refer to the Supported Platforms document for a list of currently supported versions of Caucho Resin.

Configure Resin to use a XSD-aware parser

Confluence is bundled with Xerces, which is an XSD-aware parser. Instructions on how to configure resin to use Xerces can be found here. Please ensure the Xerces related directives are added directly within the **<resin>** tag and not embedded inside the **<server>** tag. For more information, refer to the Spring Framework documentation.

Resin and SOAP

If you want to use Confluence's SOAP interface under Resin, you will need to follow these instructions.

Configure web.xml

 Confluence will not start in Resin without proper configuration of the `web.xml` file.

Resin 3.0 does not correctly support the Servlet 2.3 specification. Confluence relies on this specification to function. If you use Confluence under Resin 3.0:

- You can't seem to stay logged in, and are continually asked to enter your username and password
- The Edit tab (or other links) only appear if you grant Anonymous users permission to perform those functions

See also:

- CONF-1000
- CONF-1509
- CONF-1511
- CONF-3397

The solution is to modify Confluence's `web.xml` file so that it is conformant with the Servlet 2.4 specification. The resin file is bundled with Confluence distributions - you can find it in the /WEB-INF folder. A big thanks to Matthew Schmidt for supplying us with this workaround.

Knowledge Base Articles

 Confluence and SOAP on Resin

 Troubleshooting Application Servers

Confluence Cluster Installation

Overview

There are two methods of installing Confluence in a cluster, depending on whether you have existing data. **This page describes a fresh installation with no existing data.**

See also Confluence Cluster Installation with Existing Data.

 **Oracle Coherence Licensing Change:**

- Due to a license agreement change between Atlassian and Oracle over the Coherence technology, from September 2009, Confluence will be made available in two editions:
  - **Standard Edition** — Confluence with Ehcache's caching technology (available to customers with non-clustered Confluence licenses).

     **If you are currently running a clustered installation of Confluence, please do not upgrade it with a standard edition of Confluence.**
  - **Clustered Edition** — Confluence with Oracle's Coherence clustering and distributed caching technology (available to customers with Confluence clustered licenses only).
- For more information about these changes, please refer to the Coherence License Changes document.
- If you have a Confluence clustered license, are running a clustered installation of Confluence and wish to upgrade to Confluence version 2.6 or later after late September 2009, please ensure that you download only a clustered edition of Confluence and please refer to the Confluence 3.0.1 Upgrade Notes for additional upgrade information.

Installation with no existing data

To get Confluence running in a two-node cluster, you must do the following:

1. Ensure you meet the clustering requirements, including obtaining a clustered license key from Atlassian for each node
2. Install Confluence on a single node, configuring an external database and a cluster name
3. Load test the single node installation, see whether clustering is required
4. Shut down the first node, copy the Confluence application and Confluence home directory to the second node
5. Start the first node, wait until it is running, then bring up the second node and it will automatically join the cluster
6. Test the cluster is working correctly
7. Configure a load balancer in front of the two clustered nodes.

Each of these steps will be described in detail below.

1. Clustering requirements

Your Confluence cluster installation must meet *all* the following criteria for clustering:

- you must be running Confluence 2.3 or later
- you must have a clustered license
- you must use an external database
- you must use a load balancer with session affinity in front of the cluster.

Clustered commercial licenses may be purchased through Confluence website. Clustered evaluation licenses may be obtained by emailing sales@atlassian.com.

A cluster can run using two copies of Confluence Standalone. However, cluster administrators must understand how to configure an application server and web server with load balancing, so we recommend you are comfortable installing Confluence as a EAR/WAR in your application server before proceeding with a clustered installation.

2. Installation on first node

Cluster administrators should already be comfortable with the normal installation method, so it won't be repeated here. There are two differences in the Confluence Setup Wizard from a normal installation:

- you must use an external database
- you must enter a cluster name.



Enter a cluster name to create a new cluster

> **Technical note**
> The cluster name will be converted into a unique multicast IP address and port for your Confluence cluster. UDP multicast traffic is used for Confluence to automatically discover other nodes in the cluster when they start up.

3. Load test the single node

Most Confluence installations do not need to be clustered. Ensure you have tested your single node installation with the number of users you expect to host before going ahead with the additional complexity of clustering.

Check out our performance tuning tips for ways to improve the performance of a single instance of Confluence.

You can upgrade your single node to a multi-node cluster at any time by resuming this guide from step 4 below.

4. Copy Confluence to second node

Confluence clusters must use the same JDK, application server and application. The easiest way to ensure this is to shut down Confluence on the first node, then copy its web application and home directory to the second node:

1. Shut down Confluence on node #1.
2. Shut down your application server on node #2, or stop it automatically loading web applications.
3. Copy the Confluence web application from node #1 to node #2.
4. Copy the Confluence home directory from node #1 to node #2.

Copying the web application ensures any modifications you have made to the application itself, custom LDAP settings (`atlassian-user.xml`), and any other advanced configuration are copied to node #2.

Copying the home directory ensures the Confluence search index (the `index/` directory), the database and cluster configuration (`confluence.cfg.xml`), and any other home directory settings are copied to node #2.

5. Start Confluence on the first node, wait, then start Confluence on second node

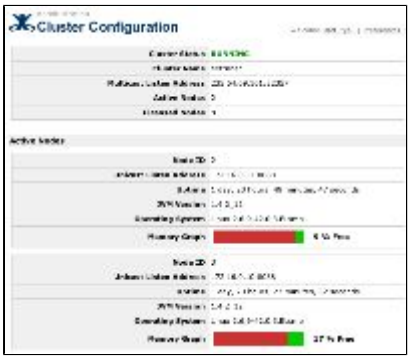
For the most stable start-up process, it is important to start Confluence one server at a time.

1. Start Confluence on node #1.
2. Wait for Confluence to become available on node #1.
3. Start Confluence on node #2.
4. Wait for Confluence to become available on node #2.

6. Test cluster connectivity

The Cluster Administration page (Administration, Cluster Configuration) includes information about the active cluster. When the cluster is running properly, this page displays:

- a correct count of the nodes in the cluster
- a status display for each node in the cluster
- an uptime for each node that is accurate.



Cluster Administration page

A simple process to ensure your cluster is working correctly is:

1. Create a new document on node #1
2. Ensure the new document is visible by accessing it directly on node #2
3. Wait one minute (Confluence does batch indexing once per minute)
4. Search for the new document on node #1, ensure it appears
5. Search for the new document on node #2, ensure it appears.

> **Technical note**
> If Confluence detects more than one instance accessing the database but not in a working cluster, it will shut itself down in a *cluster panic*. This can be fixed by troubleshooting the network connectivity of the cluster.

7. Configure load balancer

For the moment, configuring the load balancer is outside the scope of this document.

However, a simple Apache and Tomcat load-balancing configuration is available, which includes sample configuration for the Apache Tomcat and the Apache web server, using its load-balancing JK connector.

Troubleshooting

If you have problems with the above procedure, please see our Cluster Troubleshooting guide.

Upgrading a cluster

It is important that upgrades follow the procedure for Upgrading a Confluence Cluster.

Related documentation

Overview of Confluence Clusters
Clustering in Confluence
Confluence Cluster Installation with Existing Data
Confluence Installation Guide
Upgrading a Confluence Cluster
Cluster Administration page

Apache and Tomcat load balancing

Overview

The following is a description of how to setup a Confluence Cluster on a Windows machine using Apache and mod_jk to handle the loadbalancing.

The characteristics of this cluster are:

- Session affinity: sessions are associated with single servers.
- Failover: if a server dies, a connection will be directed to the nearest available server. (NOTE: sessions are not replicated)
- Failback: when a server comes back online, it will rejoin the cluster.
- Weighted load balancing: the load balancing can be controlled to take into account machine differences. (See the mod_jk documentation for details on this.)

What do you need?

1. Download and install one copy of Apache httpd. Do not install Apache as a service, but set it to listen on port 8080. (Tested with Apache httpd 2.0.55.)
2. Download the latest version of mod_jk. Copy this file into the Apache modules/ directory and rename it to `mod_jk.so`. (Tested with JK-1.2.19.)
3. Download and extract one copy of the ZIP distribution of Apache Tomcat. (Tested with Tomcat 5.5.)
4. Download JDBC drivers for the external database you will be using. Put the drivers in Tomcat's common/lib/ directory. (Tested with Postgresql 8.1, postgresql-8.1-404.jdbc3.jar).

Apache configuration

Edit the main Apache config file, conf/http.conf:

- add the following immediately after the other LoadModule directives:

```
1.LoadModule jk_module modules/mod_jk.so
```

- add the following just before the end of the file:

```
1.JkWorkersFile conf/workers.properties
2.JkLogFile logs/mod_jk.log
3.JkLogLevel info
4.
5.JkMount /confluence loadbalancer
6.JkMount /confluence/* loadbalancer
```

Create a workers.properties file in the Apache conf/ directory. This version of the workers.properties file is configured to use 2 Tomcat instances: *tomcat1* and *tomcat2*.

```
01.worker.list=loadbalancer
02.
03.worker.tomcat1.port=18081
04.worker.tomcat1.host=localhost
05.worker.tomcat1.type=ajp13
06.worker.tomcat1.lbfactor=1
07.
08.worker.tomcat2.port=28081
09.worker.tomcat2.host=localhost
10.worker.tomcat2.type=ajp13
11.worker.tomcat2.lbfactor=1
12.
13.worker.loadbalancer.type=lb
14.worker.loadbalancer.balanced_workers=tomcat1, tomcat2
15.worker.loadbalancer.method=Busyness
```

Tomcat configuration

The Tomcat configuration below will run multiple instances from the same binaries in the main Tomcat directory. For complete documentation of this configuration, see the `RUNNING.txt` file in the Tomcat distribution.

Create instance home directories

Create a directory for each instance of Tomcat, somewhere outside where you installed Tomcat. For example, if you extracted Tomcat to /opt/apache/tomcat-5.5, your instances could be in /var/tomcat-instances/tomcat1, /var/tomcat-instances/tomcat2. These folders will be referred to as the *instance home directories*.

Copy the following folders from the Tomcat installation directory into each instance home directory. Some of the folders may be empty, but copy them anyway.

- conf
- logs
- shared
- webapps

Configure server.xml in each instance

Edit conf/server.xml in the instance home directories to include the Confluence application and have distinct listen ports for Server, HTTP Connector and AJP13 Connector. All nodes can use the same Confluence webapp as long as you set confluence.home via a system property (see startup scripts below).

Attached are two sample configurations:

- tomcat1/conf/server.xml - listens on port 18080 (http) and 18081 (ajp13)
- tomcat2/conf/server.xml - listens on port 28080 (http) and 28081 (ajp13)

To use these sample config files, you will need to edit them to set the Confluence web-app location and the data source configuration.

If editing the configuration files yourself, the points to note are:

- 'Server' port must be distinct
- 'Connector' for HTTP must be uncommented and use a distinct port. Use this port for testing the node individually.
- 'Connector' for AJP13 must be uncommented and use a distinct port. This port must match the port of the worker in the Apache workers.properties.
- 'Engine' for localhost must have jvmRoute matching the name of the worker in Apache's workers.properties.
- 'Context' for Confluence must be added inside the 'Host' tag, and include a 'Resource' for the datasource, as per normal Confluence installation under Tomcat.

Create a startup script for each instance

The startup scripts for each instance must set the CATALINA_BASE environment variable and confluence.home system property. The variables in the sample scripts below should reference:

- CATALINA_HOME - Tomcat installation directorty
- CATALINA_BASE - Tomcat instance home directory (distinct for each node)
- JRE_HOME - Java runtime directory
- JAVA_OPTS - include a confluence-home system property (distinct for each node)

tomcat1/startup.bat:

```
1.set CATALINA_HOME=C:\home\mryall\opt\apache\apache-tomcat-5.5.16
2.set CATALINA_BASE=C:\home\mryall\var\tomcat-instances\tomcat1
3.set JRE_HOME=C:\Java\jre1.5.0_06
4.set JAVA_OPTS=-Dconfluence.home=C:\home\mryall\data\confluence\cluster\tomcat1 -Xmx512m
5.%CATALINA_HOME%\bin\startup.bat
```

tomcat2/startup.bat:

```
1.set CATALINA_HOME=C:\home\mryall\opt\apache\apache-tomcat-5.5.16
2.set CATALINA_BASE=C:\home\mryall\var\tomcat-instances\tomcat2
3.set JRE_HOME=C:\Java\jre1.5.0_06
4.set JAVA_OPTS=-Dconfluence.home=C:\home\mryall\data\confluence\cluster\tomcat2 -Xmx512m
5.%CATALINA_HOME%\bin\startup.bat
```

Continue setting up Confluence

Follow the Confluence Cluster Installation procedure with the steps following the app server setup.

Troubleshooting

General advice

The above tomcat configurations enable HTTP connectors on each Tomcat instance so that you can connect to the nodes individually. To check whether the load balancer (Apache & mod_jk) is causing the problem, try connecting to the individual Tomcat instances. Please note that you should not allow users to directly access individual nodes in production mode: You don't want people to bookmark nodes since the node details might change, or single nodes may be taken out of the cluster for maintenance while the cluster itself is still available.

Session-affinity doesn't seem to be working?

Ensure the name you use for your worker in workers.properties (e.g. tomcat1) matches the jvmRoute attribute of the engine tag in your Tomcat server.xml. For an example, search for 'Engine' in the attached sample config.

For troubleshooting your Confluence cluster, see Cluster Troubleshooting.

References

General

http://raibledesigns.com/tomcat/ http://httpd.apache.org/

Tomcat Clustering support

http://tomcat.apache.org/tomcat-5.0-doc/cluster-howto.html http://tomcat.apache.org/tomcat-5.0-doc/balancer-howto.html
http://tomcat.apache.org/tomcat-3.3-doc/mod_jk-howto.html

Clustering and Load Balancing in Tomcat 5, Part 1
Clustering and Load Balancing in Tomcat 5, Part 2

Confluence Cluster Installation with Existing Data

Overview

There are two methods of installing Confluence in a cluster, depending on whether you have existing data. **This page describes how to upgrade an existing Confluence instance into a cluster.**

See also Cluster installation without existing data.

Oracle Coherence Licensing Change:

- Due to a license agreement change between Atlassian and Oracle over the Coherence technology, from September 2009, Confluence will be made available in two editions:
  - **Standard Edition** — Confluence with Ehcache's caching technology (available to customers with non-clustered Confluence licenses).

    If you are currently running a clustered installation of Confluence, please do not upgrade it with a standard edition of Confluence.
  - **Clustered Edition** — Confluence with Oracle's Coherence clustering and distributed caching technology (available to customers with Confluence clustered licenses only).
- For more information about these changes, please refer to the Coherence License Changes document.
- If you have a Confluence clustered license, are running a clustered installation of Confluence and wish to upgrade to Confluence version 2.6 or later after late September 2009, please ensure that you download only a clustered edition of Confluence and please refer to the Confluence 3.0.1 Upgrade Notes for additional upgrade information.

Cluster installation from an existing copy of Confluence

To upgrade an existing copy of Confluence to run in a two-node cluster, you must do the following:

1. Ensure that your standalone version of Confluence has been upgraded to the version you want to run the Cluster on. **Do not upgrade your version of Confluence and switch to the clustered version at the same time**. First upgrade your system (e.g. from Confluence 2.5.8 Standalone to 2.7.1 Standalone) and make sure everything works fine (e.g. for a week) before switching (e.g. from Confluence 2.7.1 Standalone to 2.7.1 Clustered)
2. Ensure you meet the clustering requirements, including obtaining a clustered license key from Atlassian for each node
3. Due to CONF-8959, you need to perform attachment migration to the database before you change your license to a clustered license
4. Upgrade the existing Confluence instance to a clustered license
5. Enable clustering and select a cluster name
6. Shut down the first node, copy the Confluence application and Confluence home directory to the second node
7. Start the first node, wait until it is running, then bring up the second node and it will automatically join the cluster
8. Test the cluster is working correctly
9. Configure a load balancer in front of the two clustered nodes.

Each of these steps will be described in detail below.

1. Clustering requirements

Your Confluence cluster installation must meet *all* the following criteria for clustering:

- you must be running Confluence 2.3 or later
- you must have a clustered license
- you must use an external database
- you must use a load balancer with session affinity in front of the cluster.

Clustered commercial licenses may be purchased through Confluence website. Clustered evaluation licenses may be obtained by emailing sales@atlassian.com.

A cluster can run using two copies of Confluence Standalone. However, cluster administrators must understand how to configure an application server and web server with load balancing, so we recommend you are comfortable installing Confluence as a EAR/WAR in your application server before proceeding with a clustered installation.

You can follow the instructions to Migrate Confluence to an external database.

2. Upgrade existing instance to clustered license

Once you've obtained your clustered license from Atlassian, you can simply update the license in your running Confluence instance:

1. Go to 'Administration'.
2. Go to 'License Details', and paste in the new license.
3. Click 'Save'.

When you enter a clustered license, you will see a new line appear on this page: *Licensed Clustered Nodes*. This tells you how many nodes your Confluence license will allow.

Organisation	**Atlassian**
Date Purchased	**Aug 15, 2006**
License Type	**Confluence: Commercial Server**
Licensed Users	**Unlimited**
Licensed Clustered Nodes	**8 nodes** (2 nodes currently clustered).

License Details page shows the number of cluster nodes permitted

3. Enable clustering and select a cluster name

When you change Confluence to use a cluster license, a new menu option will be revealed on the left: *Cluster Configuration*. Open this menu item to configure your cluster.

On this page, you need to enter a name for your cluster, and possibly select which network interface on your server will be used for the cluster (only if you have more than one suitable interface).

Clicking 'Start cluster' will migrate your attachments to the database, then prevent any access to Confluence until it is restarted.

> **Technical note**
> The cluster name will be converted into a unique multicast IP address and port for your Confluence cluster. UDP multicast traffic is used for Confluence to automatically discover other nodes in the cluster when they start up.

4. Copy Confluence to second node

For the remaining steps in setting up a cluster with existing data, please continue from step 4 in the normal Confluence cluster installation guide.

5. Start Confluence on the first node, wait, then start Confluence on second node

See comment in step 4.

6. Test cluster connectivity

See comment in step 4.

7. Configure load balancer

See comment in step 4.

Troubleshooting

If you have problems with the above procedure, please see our Cluster Troubleshooting guide.

Upgrading a cluster

It is important that upgrades follow the procedure for Upgrading a Confluence Cluster.

Related documentation

Overview of Confluence Clusters
Confluence Cluster Installation
Confluence Installation Guide
Upgrading a Confluence Cluster
Confluence User Guide

Upgrading a Confluence Cluster

This page contains instructions for **upgrading an existing Confluence cluster** to a new version of Confluence. If you are not running a clustered instance of Confluence and wish to, see Confluence Cluster Installation with Existing Data.

Oracle Coherence Licensing Change:

- Due to a license agreement change between Atlassian and Oracle over the Coherence technology, from September 2009, Confluence will be made available in two editions:
  - **Standard Edition** — Confluence with Ehcache's caching technology (available to customers with non-clustered Confluence licenses).

    If you are currently running a clustered installation of Confluence, please do not upgrade it with a standard edition of Confluence.
  - **Clustered Edition** — Confluence with Oracle's Coherence clustering and distributed caching technology (available to customers with Confluence clustered licenses only).
- For more information about these changes, please refer to the Coherence License Changes document.
- If you have a Confluence clustered license, are running a clustered installation of Confluence and wish to upgrade to Confluence version 2.6 or later after late September 2009, please ensure that you download only a clustered edition of Confluence and please refer to the Confluence 3.0.1 Upgrade Notes for additional upgrade information.

You can download the latest version of Confluence from here.

Overview

The steps involved in upgrading a multi-node Confluence cluster are:

1. Backup your confluence instance.
2. Read the Release Notes for this version and check you have the required expertise to perform the upgrade.
3. Stop each node in the cluster.
4. Install the new version into the application server on the first node.
5. Install the new version into the application server onto the remaining nodes.

Step One: Backing up

We highly recommend that you backup your Confluence home and install directories and your database before proceeding.

For specific files to backup see Upgrading Confluence.

Step Two: Things you need to check ...

- Always check the release-notes for the version of Confluence you are installing for upgrade instructions specific to that version.
- To perform this upgrade you must be familiar with the usage of the application server running your Confluence Cluster, and the web server load balancing it.
- Check the Confluence Configuration Guide for your application server and database, to make sure there isn't anything extra you need to do to get Confluence running.
- Check that you know what configurations or customisations have been made to your Confluence instance. These may include specialised user management configurations and changes to Confluence's Java classes and Velocity templates.

Step Three: Stopping the cluster

It is vital that all nodes in the cluster are running the same version of Confluence. That's why the first step is to stop all the nodes.

Stop the Confluence application on each node using your application server.

Step Four: Upgrading the first node

We advise configuring your load balancing web server to redirect traffic away from Confluence until the upgrade is complete on multiple nodes.

Upgrading a cluster node uses the same process as Upgrading Confluence.

1. Unzip the new version.
2. Edit its `confluence-init.properties` to point to the existing home directory.
3. Port any immediately required customisations from the old version to the new one. Eg `atlassian-user.xml`.
4. Install the new version into the application server. Eg for Tomcat edit `confluence.xml` or `server.xml` to point to the new location, and restart Tomcat.
5. Wait for the Node to finish upgrading and confirm that you can log in and view pages before continuing to Step Five.
6. Port any additional customisations from the old version to the new version. Eg modifications to Java classes or Velocity templates.

Step Five: Upgrading other nodes

Copy the confluence installation, complete with customisations, to the next node.

1. Edit its confluence-init.properties to point to the existing home directory.
2. Install the new version into the application server. Eg for Tomcat 5 edit confluence.xml to point to the new location, and restart Tomcat.
3. Wait for the Node to finish upgrading and confirm that you can log in and view pages before continuing with the next node.

Troubleshooting

For suggested troubleshooting techniques, see our Cluster Troubleshooting page.

Related documentation

Overview of Confluence Clusters
Confluence Installation Guide
Cluster Troubleshooting
Confluence Cluster Installation
Confluence Cluster Installation with Existing Data
Confluence User Guide

Installing Sun JDK for Confluence

This page contains instructions for installing the Sun Java Development Kit (JDK) for Windows and Linux/UNIX.

Please refer to our Supported Platforms topic for details on the Sun JDK versions which are supported for Confluence.


Mac OS X users can ignore these instructions because this operating system already comes pre-installed with a JDK.

Installing the Sun JDK on Windows

1. If you are not sure whether you have Sun's JDK installed correctly, please confirm by doing the following:
    - Run a complete directory search (using the 'All File and Folders' option if available) on your drives for the occurrence of '`jdk` ' in the filename.
        - If your results retrieve a folder with the name '**jdk**' immediately followed by a series of version numbers (for example, `jdk1.5.0_06`), then your Sun JDK has been installed. You should double-check the contents of this folder to ensure that the installation files are intact. If you are unsure about this, proceed to step 2 to re-install the Sun JDK. Otherwise, proceed to step 3 to check that your `JAVA_HOME` environment variable has been set correctly.
        - If you do not retrieve a result like this, then it is highly likely that your Sun JDK has not been installed.
2. To install the Sun JDK, follow these instructions:
    - Go to the Java Sun download page.
    - Download the version entitled 'JDK 6 Update XX', where 'XX' stands for some number. (Sun will provide the latest version on that page.)
    - When the download has finished, run the Java installer. At one point, you will be asked to choose a directory to install to. Copy or write this directory down for use later.
3. Check that the `JAVA_HOME` environment variable has been set correctly.
    - Open the 'Start' menu, choose 'Run', type `cmd` in the 'Run' dialog box and click the '**OK**' button.
    - In the command prompt window, type `echo %JAVA_HOME%` and then press *Enter*.
    - View the result:
        - If a directory path is display that looks similar to one of the following examples, with the letters '**jdk**' immediately preceding a series of version numbers, and this path matches the location where you installed the Sun JDK in step 2, then your Sun JDK has been successfully installed and your `JAVA_HOME` environment variable has been set correctly.
          Examples of typical JAVA_HOME environment variable values:
            - `C:\Program Files\Java\jdk1.5.0_06`
            - `C:\Progra~1\Java\jdk1.5.0_06`
            - `C:\Java\jdk1.5.0_06`
            - `C:\jdk1.5.0_06`
        - If nothing is displayed or you do not see '**jdk**' immediately followed by a series of version numbers (like one of the examples above), then you need to set the `JAVA_HOME` environment variable. Please follow these instructions to set your `JAVA_HOME` environment variable to the directory you where you have just installed the JDK. By default, this directory is under `C:\Program Files\Java`.

Setting the JAVA_HOME Variable in Windows


This information is only relevant if you are installing Confluence on a Windows server.

After you have installed the Java Development Kit (JDK) in Windows, you must set the `JAVA_HOME` environment variable to the JDK installation directory.

Stage 1. Locate the JDK Installation Directory

If you already know the installation path for the Java or Software Development Kit, go to *Stage 2* below. Otherwise, find the installation path by following these instructions:

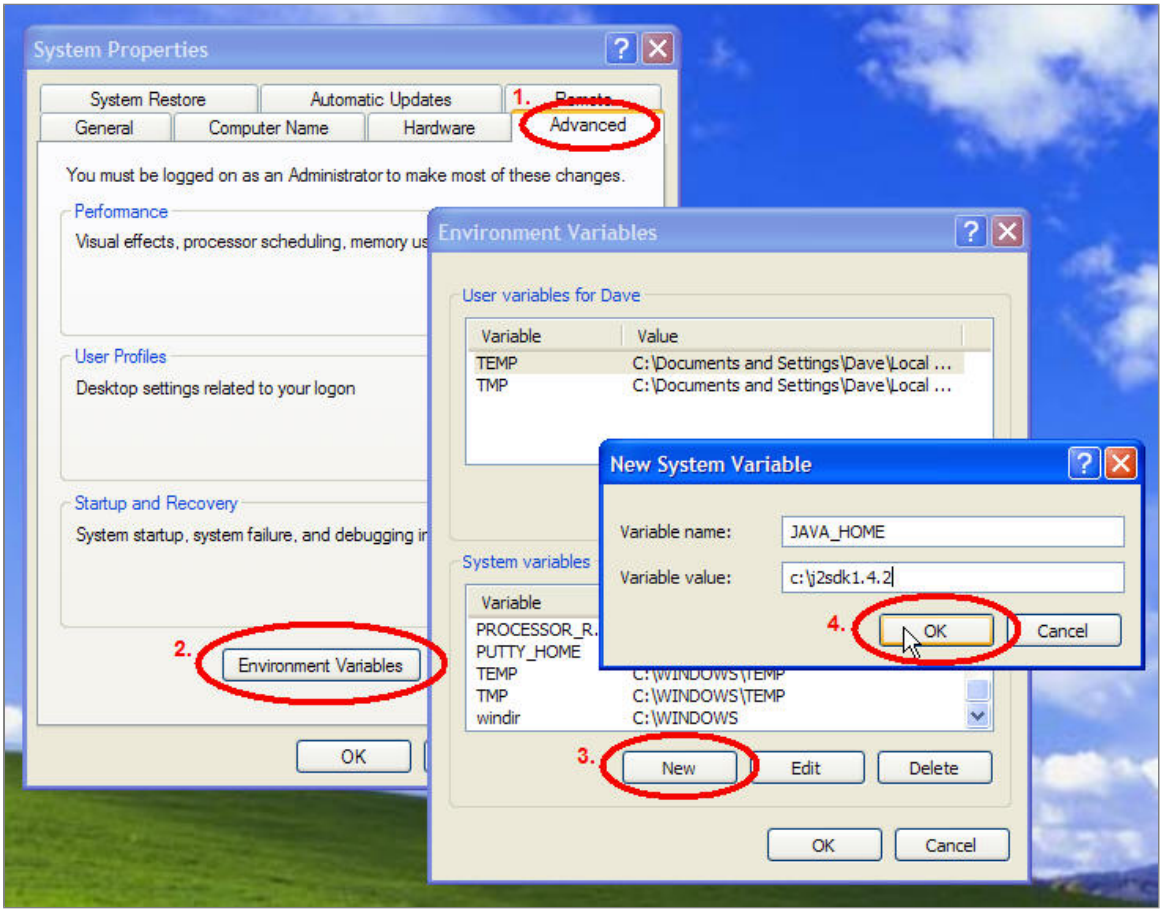
1. Unless you changed the installation path for the Java Development Kit during installation, it will be in a directory under `C:\Program Files\Java`. Using Explorer, open the directory `C:\Program Files\Java`.
2. Inside that path will be one or more subdirectories such as `jdk1.5.0_08`. If you have just installed the Java Development Kit, it will be installed to the newest directory, which you can find by sorting by date. For example, it may be installed in `C:\Program Files\Java\jdk1.5.0_08`. This is the installation path.

Stage 2. Set the JAVA_HOME Variable

Once you have identified the JDK installation path:

1. Right-click the **My Computer** icon on your desktop and select **'Properties'**.
2. Click the **'Advanced'** tab.
3. Click the **'Environment Variables'** button.
4. Under **'System Variables'**, click **'New'**.
5. Enter the variable name as `JAVA_HOME`.
6. Enter the variable value as the installation path for the Java Development Kit.
7. Click **'OK'**.
8. Click **'Apply Changes'**.
9. If you are running Confluence as a EAR/WAR rather than the Standalone, you may need to restart your application server.

This diagram shows setting the `JAVA_HOME` variable to an installation path of `c:\j2sdk1.4.2`:



If you came here from Installing Confluence Standalone Using the Windows Evaluation Installer, go back and begin *Stage 3*.

RELATED TOPICS

Installing Confluence Standalone Using the Windows Evaluation Installer
Starting Tomcat as a Windows Service
Installing Confluence in Linux

Confluence Unix and X11 Dependencies

Java X11 Dependencies

On Unix-based operating systems, the Java runtime makes use of certain parts of the platform's native X11 graphics libraries. The X Server does not have to be running, but the libraries must be available on the server. Confluence will run on a server that does not have Xlib installed, but parts of the application that manipulate graphics: PDF exports, image thumbnailing, the image gallery macro, CAPTCHA, and the resizing of profile pictures, will fail.



Mac OS X
You do not need to install X11 on Mac OS X, as it has its own graphics libraries.

If X11 is not present, you may see any of the following errors

- "This Confluence installation can not generate thumbnails: no image support in Java runtime"
- "Exception in thread "main" java.lang.UnsatisfiedLinkError: /usr/local/j2sdk1.4.2_09/jre/lib/i386/libawt.so: libXp.so.6: cannot open shared object file: No such file or directory" when exporting a PDF
- "NoClassDefFoundError" when uploading a profile picture



If This Doesn't Help
If you have X11 installed and thumbnailing still does not work, please ensure that you are running Java in headless mode — see the FAQ entitled Confluence doesn't generate thumbnails.

Specific Installation Instructions

Fedora Core

On Fedora Core, you will need to install the `xorg-x11-deprecated-libs` package. (https://bugzilla.redhat.com/bugzilla/show_bug.cgi?id=130239)

Fedora Core 6, RHEL 5

- libXp
- libXp-devel (if you wish to compile against this library)

Debian Linux

On Debian, you will need to instll the following packages (CONF-6411):

```
apt-get install libx11-6 libx11-dev libxt6 libxt6-dbg libxext6 libxtst-dev libxtst6 xlibs-dbg xlibs-dev
```



You'll only need the xlibs-dbg package if you're running an older version of Debian (3.0). It's a dummy package for smoothing the transition to a new set of graphics libraries, so if you can't locate it, you most likely don't need it.

Gentoo Linux

```
emerge libICE libSM libX11 libXext libXp libXt libXtst
```

Solaris 10

Please refer to the following forum for more information.

Ubuntu

Execute the following:

```
> apt-get install libice-dev libsm-dev libx11-dev libxext-dev libxp-dev libxt-dev libxtst-dev
```

Note: '**sudo**' enables you to be superuser for one operation. You will need to supply your user password.

Get A Confluence Licence

Need a Confluence licence or licence key?

- If you do not yet have a licence, you can get a free multi-user Evaluation licence or a 10-user Starter licence immediately.
- If you already have a Confluence licence, you can retrieve your key or generate a new key from the Licence Viewer.
- For enterprise, non-profit, open source and educational licences, see Confluence Licencing and Pricing.
- If you can't find your key or are having problems, contact sales@atlassian.com.

Running Confluence in a Virtualised Environment

This page provides some performance data and observations on running Confluence with VMware. The information on this page is intended to help you decide whether or not to run Confluence using a VMware product. It does not contain detailed instructions on how to set this up (please refer to the appropriate VMware product documentation instead).

On this page:

- Summary
- Recommendations
  - General
  - VMware ESX 3.5
  - VMware ESX 4i
- Performance Testing Setup
  - Server Configuration
  - Installed Software
  - Testing Tool
- Test Results
  - Result Descriptions
  - Low-to-Medium-load Confluence
  - Medium-to-High-load Confluence

Summary

Confluence is generally slower in a virtualised environment. As can be seen in the test results below, the amount by which Confluence slows down varies based on the workload.

Under low load there are several operations which are in fact faster under VMware. This is probably due to the 4CPU VM instance running on 8 real CPUs as opposed to there being only 4 real CPUs on the baseline machine.

Please note, no performance tuning was applied to VMware for these tests. It may be possible to improve Confluence's performance by tuning VMWare. However, this may cause other applications to run more slowly on the virtual environment. We recommend that you consult the VMware documentation before deciding whether to do this.

Recommendations

General

- If you are a running a **medium-to-high-load** instance, your biggest performance gain will be to run the application and database on a real machine and not on virtual infrastructure.
- Under **medium-to-high-load**, moving the database onto another machine will help.
- Always ensure that there are enough virtual CPUs and memory allocated to the virtual instance. This may not be possible under VMware ESX 3.5 due to limitations of 4 vCPUs per VM.
- Always ensure that there is enough CPU time and memory available on the physical host to service all VMs. Applications should not go into swap.
- Use modern CPUs with VT extensions — there is still a noticeable performance penalty for using a VM with these CPUs, but it will likely be much higher when using old CPUs.
- Carefully monitor your VMware hosts to ensure that there is no resource starvation.

VMware ESX 3.5

- If possible, upgrade to VMware ESX 4i.
- Under **low-to-medium-load**, using a non-virtualised database will generally result in better response times.

VMware ESX 4i

- Under **low-to-medium-load**, keep the database inside the virtual machine if there is enough CPU time for both the database and application.
- Using VMware EX 4i and virtual machine version 7, you will be able to allocate up to 8 vCPUs to an instance.

Performance Testing Setup

Server Configuration

All testing was performed on the following hardware. In the case of virtual machines, one VM per machine was configured.

Platform	CPU	Real Ram	Disk	Virtualisation Software	Virtual machine version	Virtual CPU's	Virtual Ram
Dell R610	2 x Intel 'Nehalem' Xeon E5520 (Quad Core) [1]	32Gb (8x 4Gb DDR3)	2 x 15K 146Gb SAS, Raid 1 [4,5]	VMware ESX 3.5	4	4	32Gb
Dell R610	2 x Intel 'Nehalem' Xeon E5520 (Quad Core) [1]	32Gb (8x 4Gb DDR3)	2 x 15K 146Gb SAS, Raid 1 [4,5]	VMware ESXi 4	7	4	32Gb
Dell R610	2 x Intel 'Nehalem' Xeon E5520 (Quad Core) [2,3]	32Gb (8x 4Gb DDR3)	2 x 15K 146Gb SAS, Raid 1 [5]	N/A	N/A	N/A	N/A

Notes:

1. *VT extensions were enabled in the BIOS on the machines running VMWare.*

2. *VT extensions were disabled in the BIOS on the machines not running VMWare, as per Dell best practices.*
3. *In order to limit the CPUs in the baseline test to match the number in VMWare, the kernel boot parameter `maxcpus=4` was added to the startup.*
4. *The full disk was allocated to VMware.*
5. *The filesystem used in all machines was EXT3.*

Installed Software

Each server was set up with identical software, as follows:

Atlassian Product	Confluence 3.0.1-rc2
Database	PostgreSQL 8.2.6
Application Server	Tomcat 6.0.14
Java	Java(TM) SE (build 1.6.0_07-b06), Java HotSpot(TM) 64-Bit Server VM (build 10.0-b23, mixed mode)

Operating System	Redhat Enterprise Linux 5.3 (Tikanga) 64bit (Kernel 2.6.18-128.2.1.el5). The file system used for all tests was EXT3 with the default options.
	The following tuning was applied to the operating system, in order to allow for more memory usage by the database server and better network throughput:
	`net.ipv4.ip_forward = 0`
	`net.ipv4.conf.default.rp_filter = 1`
	`net.ipv4.conf.default.accept_source_route = 0`
	`kernel.sysrq = 0`
	`kernel.core_uses_pid = 1`
	`net.ipv4.tcp_syncookies = 1`
	`kernel.msgmnb = 65536`
	`kernel.msgmax = 65536`
	`kernel.shmmax = 1310720000`
	`kernel.shmall = 4294967296`
	`net.core.rmem_max = 16777216`
	`net.core.wmem_max = 16777216`
	`net.ipv4.tcp_rmem = 4096 87380 16777216`
	`net.ipv4.tcp_wmem = 4098 65536 16777216`
	`net.ipv4.tcp_no_metrics_save = 1`
	`net.ipv4.tcp_moderate_rcvbuf = 1`
	`net.core.netdev_max_backlog = 2500`

Testing Tool

Performance tests were conducted with **Apache Jakarta JMeter 2.3.4** using the standard Confluence performance tests.

Test Results

The following tests were performed for each application. In each case, the test was performed with a database local to the host instance (i.e. in the same operating system image) and also with the database residing on a separate, non-virtualised physical server of the same specifications as above.

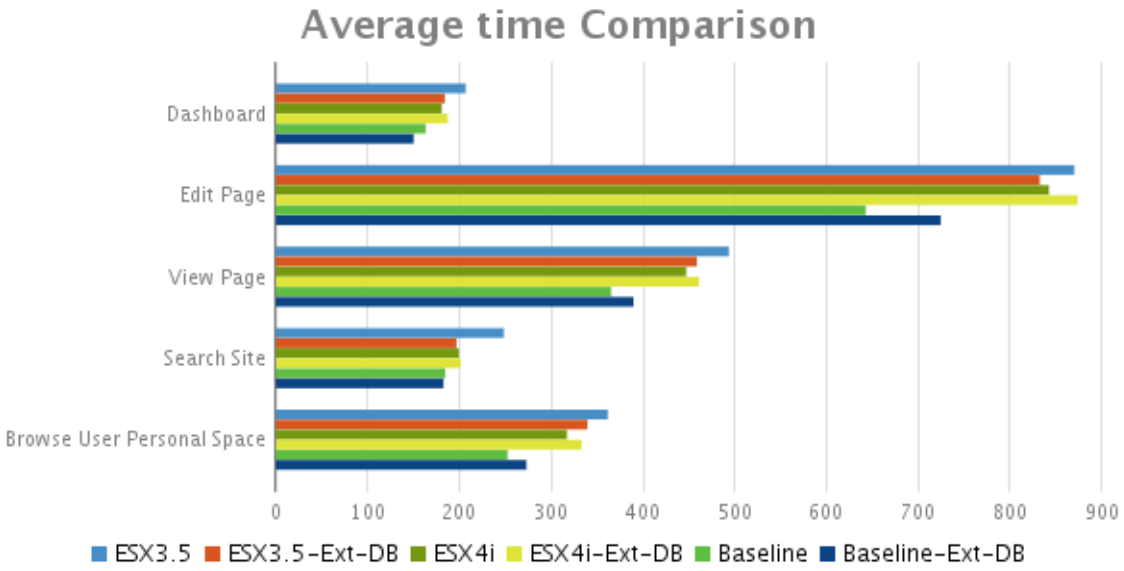
Result Descriptions

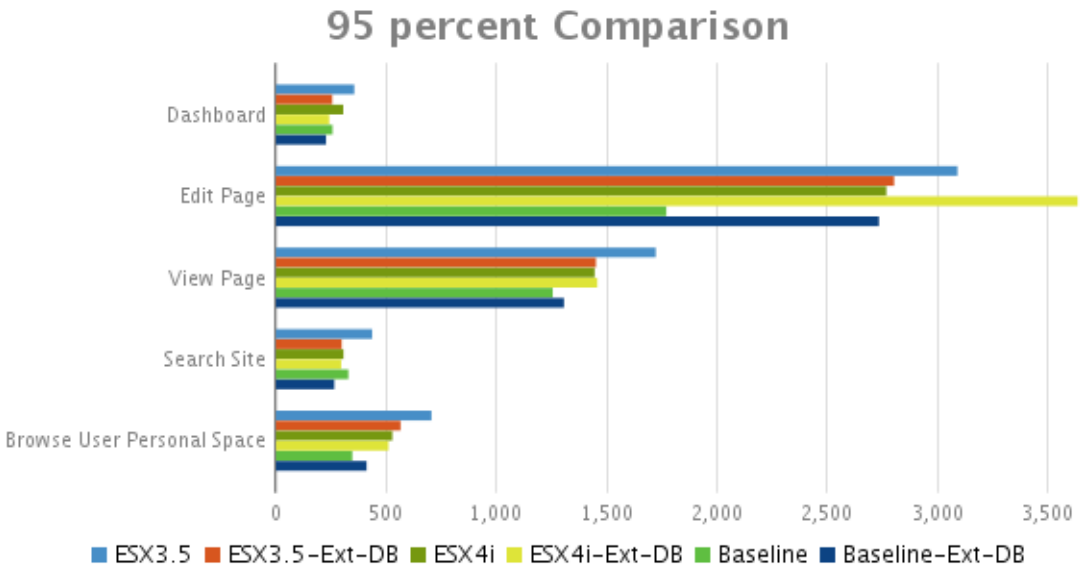
The following descriptions relate to the result graphs below.

- **Average time Comparison** — The average response time of the requests in the scenario - the lower the better.
- **95 percent Comparison** — The time (in milliseconds) by which 95% of all requests in the scenario have completed. This is not an average value – rather, you can think of it as a 'how long the slowest requests (except the very worst 5% cases) take to complete' scenario.
- **Scenarios:**
  - **Dashboard** — Simulates visiting the Confluence dashboard.
  - **Edit Page** — Simulates saving a page back to Confluence and notifying all people who are watching this page.
  - **View Page** — Simulates loading one out of hundreds of different Confluence pages. Some are short, others are long. Some have many images, others have many comments. Some have many macros, others do not. The pages are accessed through their full URL, as if someone had clicked a link within the application or a bookmark.
  - **Search Site** — Simulates a search across the whole system.
  - **Browse User Personal Space** — Simulates regular browsing of pages in a user's personal space.
- **Ext-DB** (In the legend of each graph) — Indicates scenarios in which the database resides on a separate, non-virtualised physical server.

Low-to-Medium-load Confluence

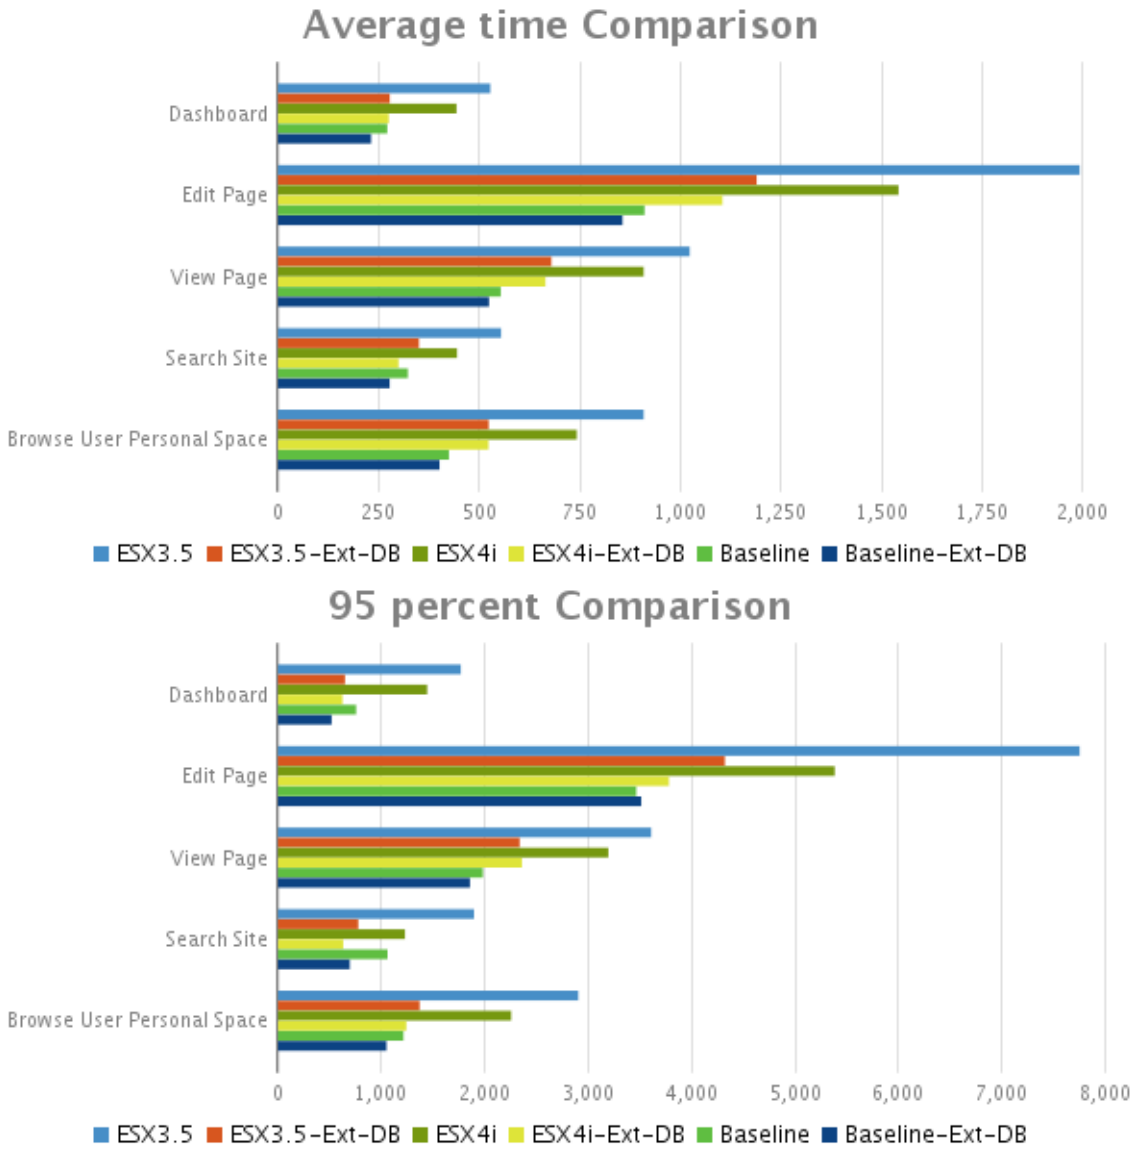
This test performs around **18 requests/second** on the Confluence instance. This is not enough to saturate the host CPU time and during the test there is around **50-75% idle time**. You could expect to see similar results if your Confluence instance has enough resources available to it.





Medium-to-High-load Confluence

This test tries to perform double the requests/second of the lower load test (i.e. approximately **36 requests/second**) on the Confluence instance. This is enough load to saturate the available CPU time on a 4 CPU machine. This test is designed to simulate an instance which does not have enough resources to serve the number of requests being made to it.



Uninstalling Confluence Standalone

If you installed Confluence Standalone using the Confluence Installer (Confluence 2.8 or later), you can **uninstall Confluence** using the uninstaller.

- On Windows, click '**Start**'...> '**Programs**'...> '**Atlassian Confluence Uninstaller**'.
- On the Mac, click '**Applications**'...> '**Atlassian Confluence Uninstaller**'.

Home Directory will survive
The Confluence Uninstaller will not clear your Home directory. This means that your Confluence database will be safe, even if you are using the embedded HSQLDB database. The uninstaller will clear the Installation only.

Confluence Setup Guide

Before running the Confluence Setup Wizard described below, please follow the instructions on 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. Enter your License Key
- 3. Choose your Installation Type
- 4. Custom Installation: Database Configuration
- 5. Custom Installation: External Database

- 6. Custom Installation: Load Content
- 7. Custom Installation: Restore Data from Backup
- 8. Enter Details of your Confluence System Administrator
- 9. Setup is Complete

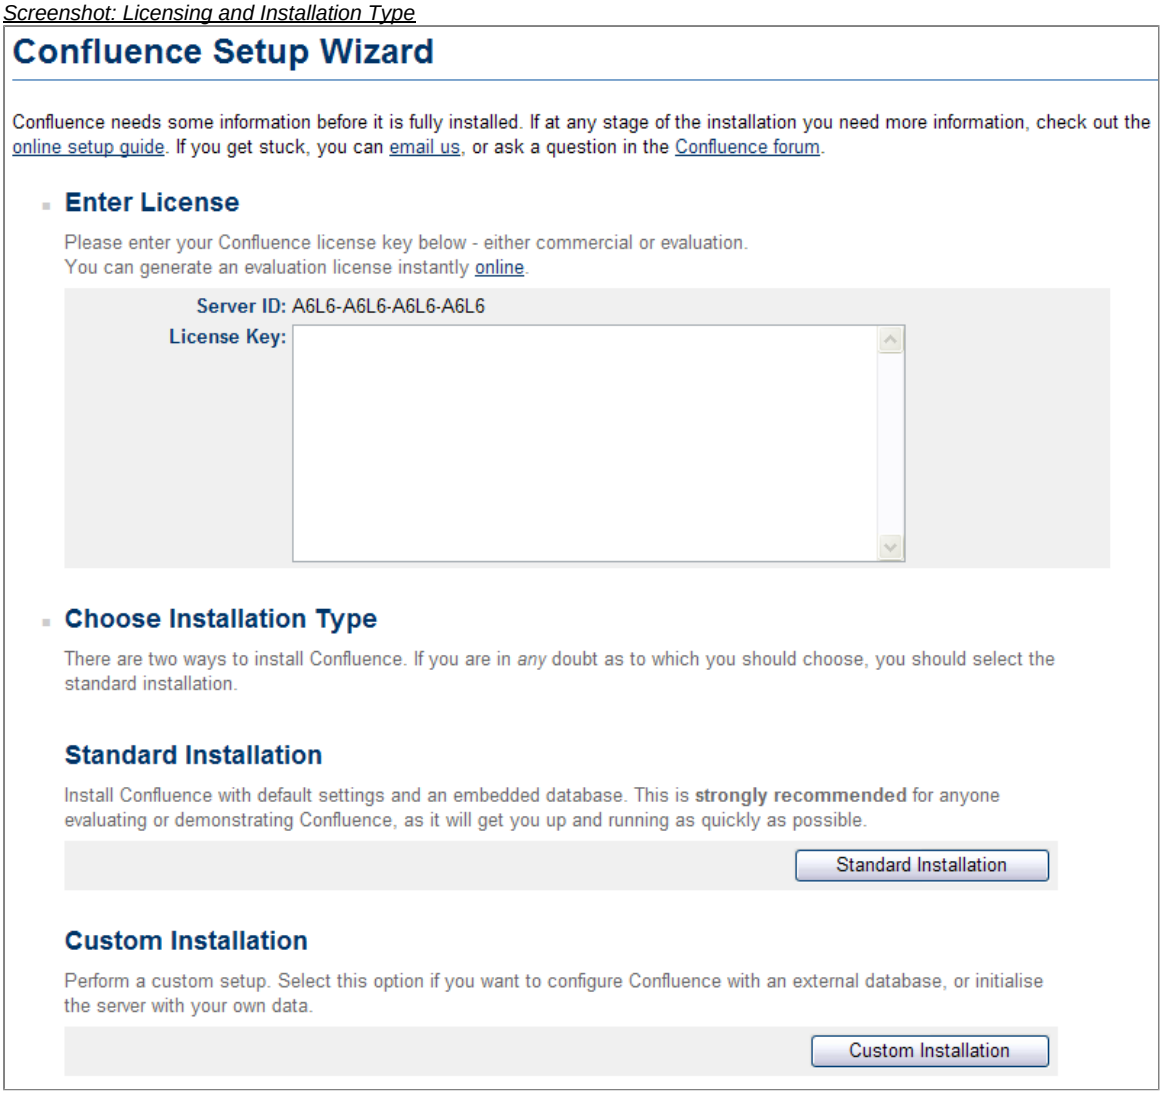
1. Start the Setup Wizard

1. If Confluence is not already running, i.e. if you did not configure it to start automatically during installation, you need to start it now:
    - If you are running Confluence Standalone on Windows, click '**Start**'..., '**Programs**'..., '**Atlassian Confluence**'..., '**Start in Console**'.
    - If you are running Confluence Standalone on a MAC, click '**Applications**'..., '**Atlassian Confluence**'.
    - Or run the start-up script found in the `bin` folder of your Installation directory:
        - `startup.bat` for Windows.
        - `startup.sh` for Unix-based systems.
        - Mac OS X users will be prompted to choose an application. Choose the **Terminal** application in the Utilities folder.
2. Go to the following web address in your web browser: http://localhost:8080

    The above web address uses port '8080'. If you chose a different port during installation, change '8080' to the number you chose.
    - You should see the Licensing screen described below.
    - If an error message appears, first check that you're using the port which you specified during installation. Then check the Installation FAQ.

2. Enter your License Key

Screenshot: Licensing and Installation Type



Hint: The above image and all the images on this page are screenshots. Clicking an image will not configure Confluence.

1. Find your Confluence license key:
    - If you don't have your license key handy, you can retrieve your existing key.
    - If you do not already have a Confluence license, you can obtain one now:
        - Copy your '**Server ID**' from the Setup Wizard's Licensing screen, shown on the screenshot above.
        - Obtain a trial, free or commercial licence.
2. Type or paste your license key into the '**License Key**' field, shown on the screenshot above.

3. Choose your Installation Type

Refer to the screenshot above. In this step, you will choose whether you want a standard or a custom installation.

Option 1: Standard Installation — Set up Confluence with the embedded HSQLDB database and default settings. This option will also install a 'Demonstration Space' with some example content to get you working with Confluence as quickly and easily as possible. You may upgrade to another type of database later on.

Hint: Who should choose this option?

- Choose the standard installation if you are evaluating Confluence or if you are new to Confluence.
- This option is not recommended for production instances of Confluence.

For production use, we strongly recommend that you connect to an external database rather than using the embedded database. The standard installation is therefore not suitable for production installations.

Next, you will be asked for details of your system administrator. Go to step 8 below. Yes, you really can skip all the steps between

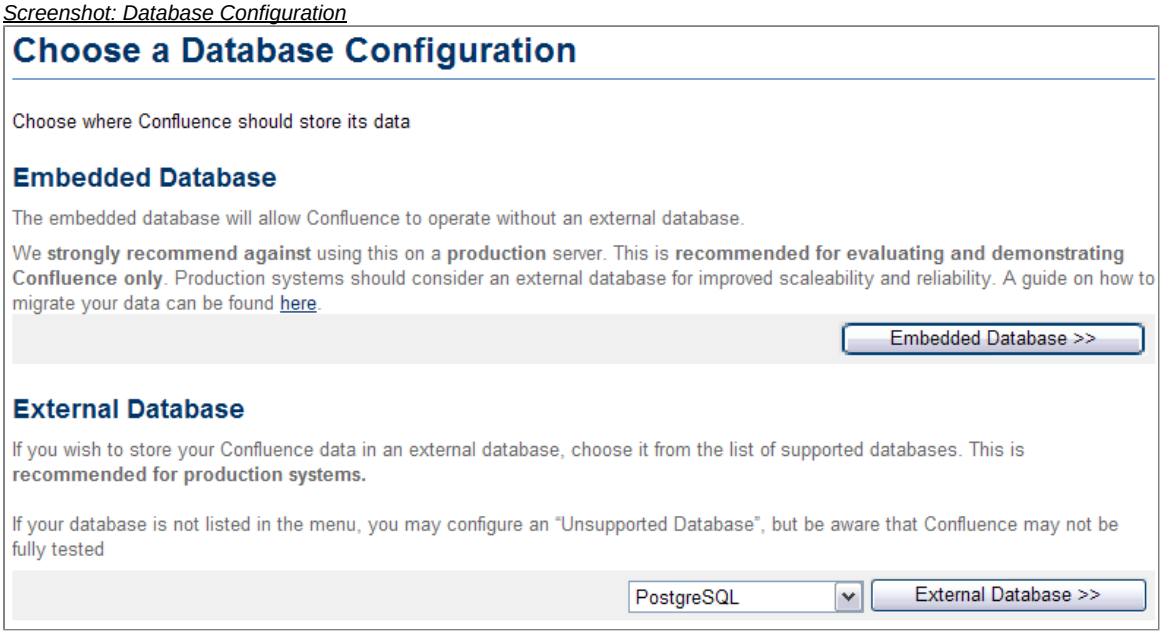
Option 2: Custom Installation — Customise your Confluence instance to use your own database and your own data.

Hint: What options does the custom installation offer?

- Connect Confluence to an external database. **Recommended for production instances of Confluence**.
- Restore data from an existing Confluence database.
- Install Confluence without the demonstration content.

4. Custom Installation: Database Configuration

Screenshot: Database Configuration



The above screen appears if you have chosen a custom installation of Confluence. You can choose to use the embedded database supplied with your Confluence installation, or to connect to an external database.

- **Option 1: Embedded Database** — If you select this option, Confluence will use an embedded HSQLDB database. You should only

select this option for the purposes of ***evaluating*** or ***demonstrating*** the use of Confluence only.

You can migrate to an external database later on if you wish.

- **Option 2: External Database** — If you wish Confluence to use an external database, select your database type from the database dropdown list and then click the '**External Database**' button.

  For production purposes, you ***should only use*** an external database to ensure your data is kept safe and consistent.

  Read the Supported Platforms topic for more information about which databases are supported and for details about choosing an external database, refer to the System Requirements topic.

5. Custom Installation: External Database

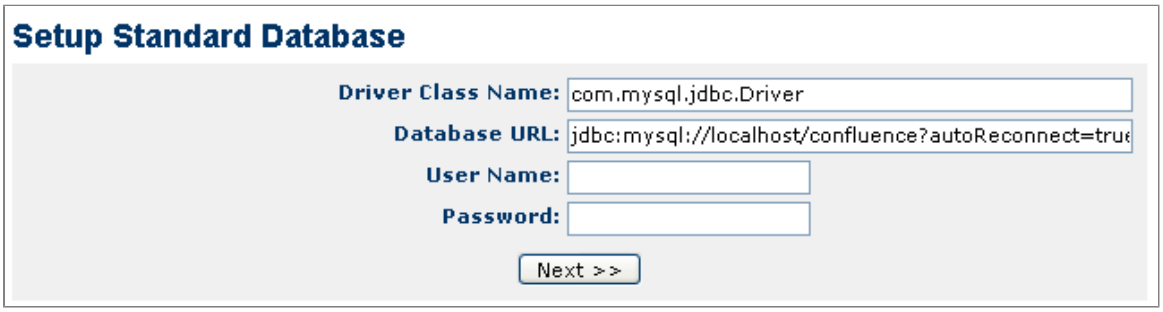
Before you Start

- **Character encoding:**
  - We strongly recommend that character encoding is consistent across your database, application server and web application, and that you use **UTF-8** encoding.
  - Before setting up your database, please read about configuring character encoding.
- **Database name:** When creating a new external database, give it the name '**confluence**'.

You can choose to configure your database via a standard JDBC connection or via a server-managed datasource connection. Choose one of the two options below.

Option 1: Standard Database Connection — This uses a standard JDBC database connection. Connection pooling is handled within Confluence.

Screenshot: Standard (JDBC) Connection



Supply the following information:

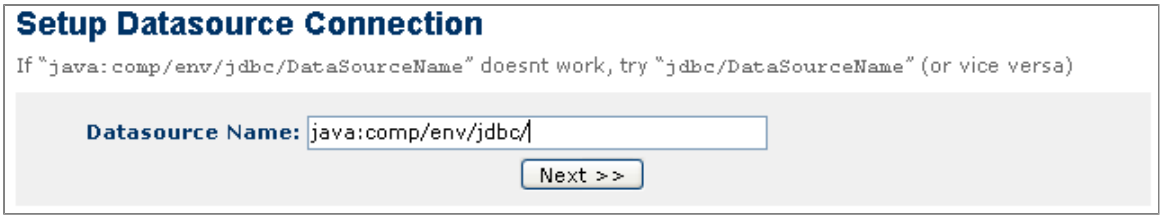
- **Driver Class Name** — The Java class name for the appropriate database driver. This will depend on the JDBC driver, and will be found in the documentation for your database. You will also need to put the appropriate database driver 'jar' file in the server's classpath. For the standalone version, this means copying the jar file into the `<confluence-install>/lib` directory.
- **Database URL** — The JDBC URL for the database you will be connecting to. This will depend on the JDBC driver, and will be found in the documentation for your database.
- **User Name** — A valid username which Confluence will use to access your database.
- **Password** — The password corresponding to the above username.

You will also need to know:

- The size of the connection pool Confluence should maintain. If in doubt, just go with the default provided.
- What kind of database you're connecting to, so you can tell Confluence which dialect it needs to use.

Option 2: Datasource Connection — This asks the Java application server for a database connection. You will need to have configured a datasource into your application server.

Screenshot: Datasource Connection



Supply the following information:

- **Datasource Name** — The JNDI name of the datasource, as configured in the application server.

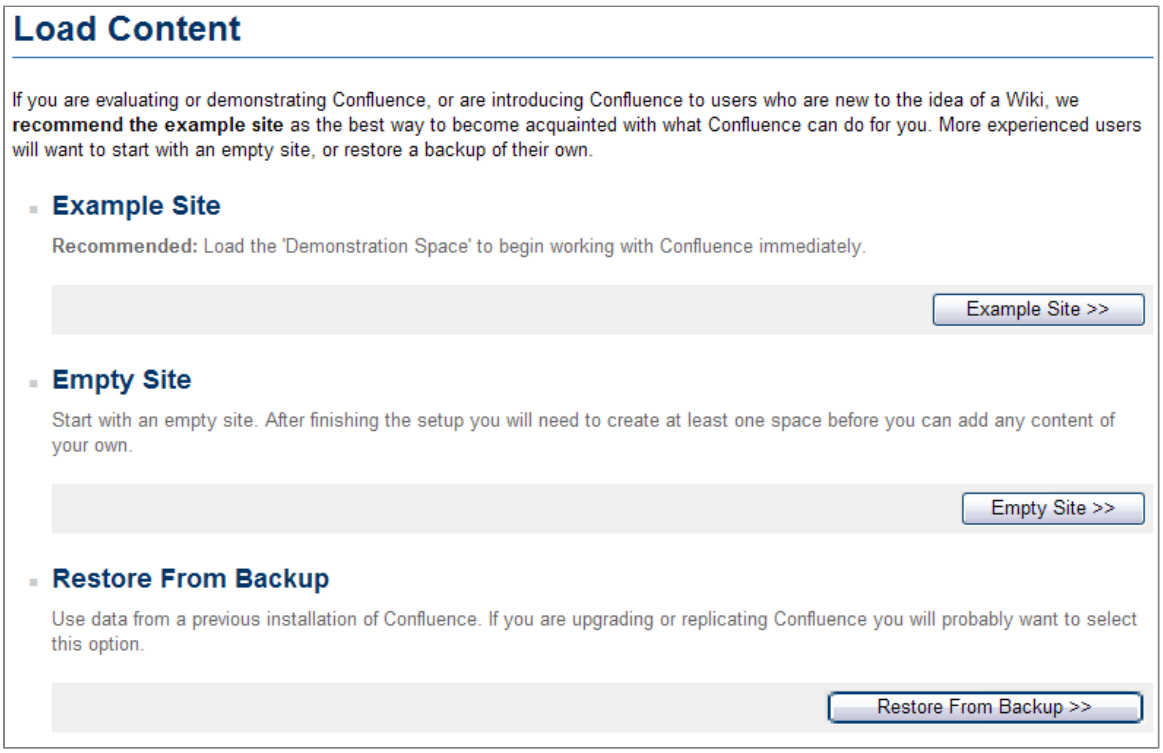
Note: Some servers will have JNDI names like `jdbc/datasourcename`; others will be of the form `java:comp/env/jdbc/datasourcename`. Consult your application-server documentation.

You will also need to know:

- What kind of database you're connecting to, so you can tell Confluence which dialect it needs to use.

6. Custom Installation: Load Content

Screenshot: Load Content

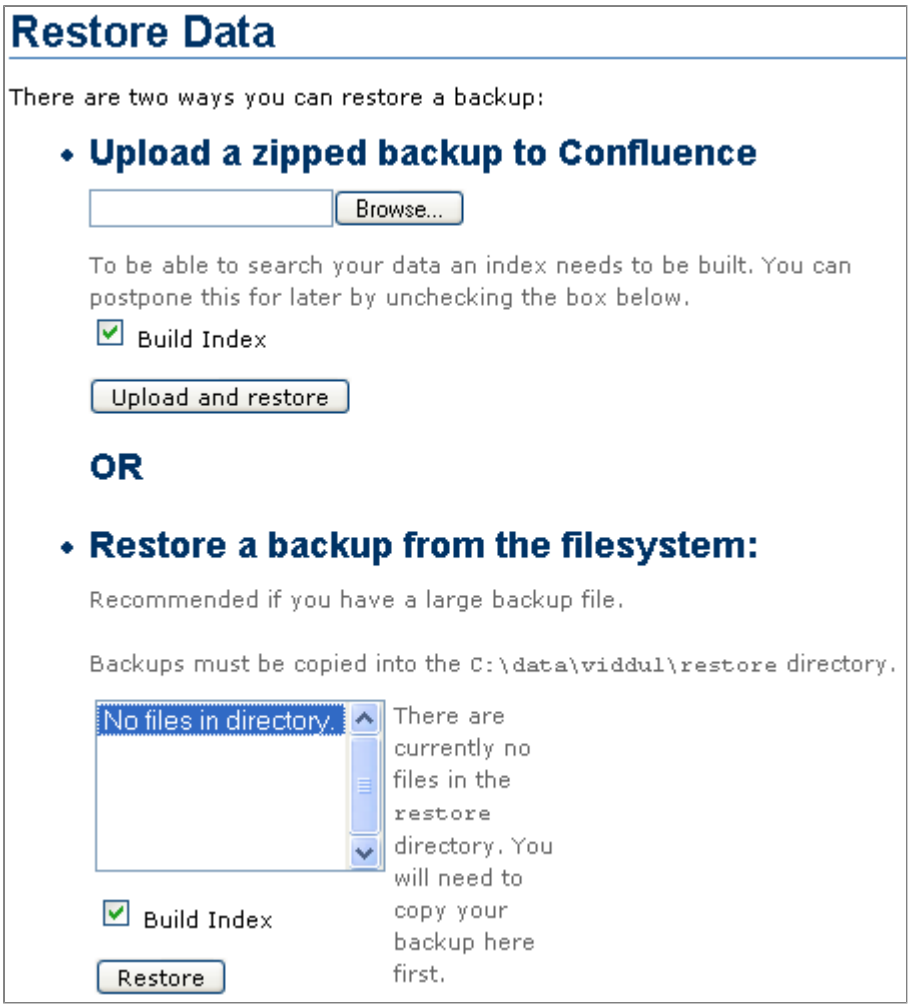


Select one of the following options:

- **Example Site** — This option will load Confluence's 'Demonstration Space'. Select this if you are using Confluence for the first time, or if you want the Demonstration Space for your other Confluence users. The Demonstration Space helps to familiarise you with Confluence and what it can do for you. You can then continue using your Confluence deployment as normal — there's no need to reinstall later.
- **Empty Site** — Select this option if you are already familiar with Confluence. You will need to create at least one space before you can start adding content to the site.
- **Restore from Backup** — Select this option if you want to use Confluence data from a previous installation.

7. Custom Installation: Restore Data from Backup

Screenshot : Restoring Data



This option allows you to reload your data from an existing Confluence installation into your new Confluence site during the initial setup procedure. You can choose to upload data from a zipped backup file, or to restore from a backup file on your file system.

Option 1: Upload a zipped backup to Confluence — This option will load the data from a zipped backup file.

To create a backup file from your existing version of Confluence, go to the '**Backup & Restore**' section of your **Administration Console**.

To restore from a zipped backup:

1. Browse for the relevant daily backup file or a file you have created via a manual backup.
2. Check '**Build Index**' to build the data index, used for the search.
3. Click the '**Upload and Restore**' button.

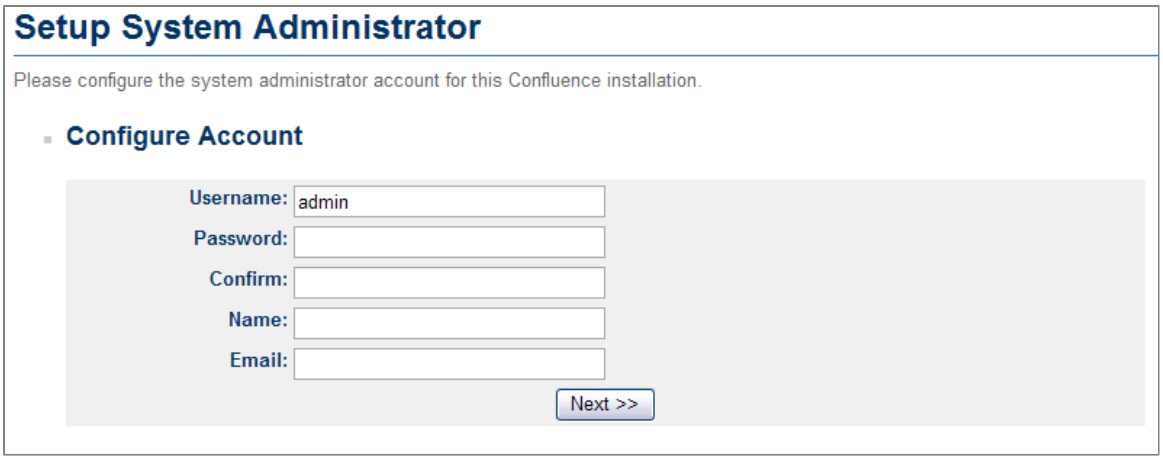
Option 2: Restore a backup from the filesystem — This option is recommended if you have a very large daily backup file (greater than 100MB), or a daily backup file that is already on the server and doesn't require uploading.

1. Copy the backup file into the `restore` directory inside your confluence Home directory and then refresh the page. You should now see your backup file appear on the '**Restore Data**' screen (pictured above), in the box beneath the heading '**Restore a backup from the filesystem**'.
2. Check '**Build Index**' to build the data index, used for the search.
3. Click the '**Restore**' button.

When the restore process has finished, you are ready to log in to Confluence. The system administrator account and all other information has been transferred from your previous Confluence installation.

8. Enter Details of your Confluence System Administrator

Screenshot: System Administrator



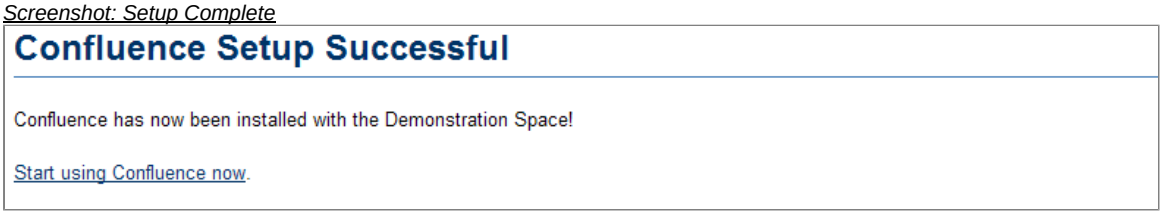
The system administrator has full administrative power over your Confluence instance. This person will be able to add more users, create spaces, and set further Confluence options. Please refer to the overview of global permissions for more information.

Hint: If you are evaluating Confluence, set yourself up as the administrator.

1. Enter the following information to set up your system administrator's user account:
   - **Username** — The username under which the system administrator will log in to Confluence, e.g. 'jsmith'.
   - **Password** — The password which the system administrator will use to log in.
   - **Confirm** — Enter the same password again.
   - **Name** — The system administrator's full name, e.g. 'John Smith'.
   - **Email** — The system administrator's email address, e.g. 'jsmith@example.com'.
2. Click '**Next**'.

9. Setup is Complete

Screenshot: Setup Complete



Congratulations! You have installed and set up Confluence. Click the '**Start using Confluence now**' link to open the '**Demonstration Space**' in your Confluence wiki. This space contains some sample content and ideas, to help you get started quickly.

RELATED TOPICS

Using the Confluence Dashboard
Starting Confluence Automatically on System Startup
Documentation Home

External Database

Custom Installation - Connecting to an External Database

This page is part of the Confluence Setup Guide.

Before you Start

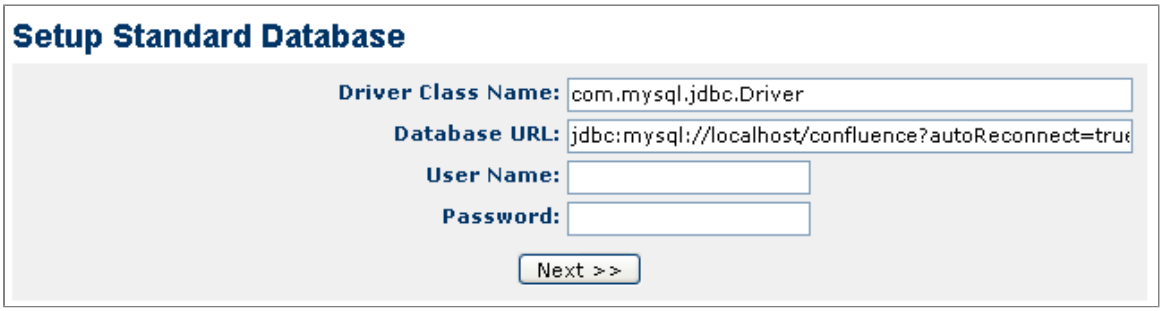
- **Character encoding:**
  - We strongly recommend that character encoding is consistent across your database, application server and web application, and that you use **UTF-8** encoding.
  - Before setting up your database, please read about configuring character encoding.
- **Database name:** When creating a new external database, give it the name '**confluence**'.

You can choose to configure your database via a standard JDBC connection or via a server-managed datasource connection. Choose one of

the two options below.

Option 1: Standard Database Connection — This uses a standard JDBC database connection. Connection pooling is handled within Confluence.

Screenshot: Standard (JDBC) Connection



Supply the following information:

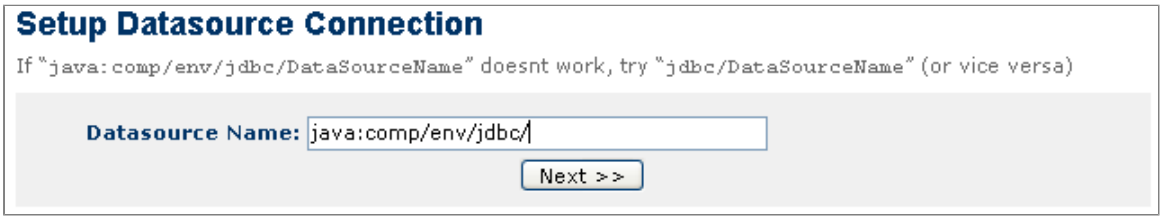
- **Driver Class Name** — The Java class name for the appropriate database driver. This will depend on the JDBC driver, and will be found in the documentation for your database. You will also need to put the appropriate database driver 'jar' file in the server's classpath. For the standalone version, this means copying the jar file into the `<confluence-install>/lib` directory.
- **Database URL** — The JDBC URL for the database you will be connecting to. This will depend on the JDBC driver, and will be found in the documentation for your database.
- **User Name** — A valid username which Confluence will use to access your database.
- **Password** — The password corresponding to the above username.

You will also need to know:

- The size of the connection pool Confluence should maintain. If in doubt, just go with the default provided.
- What kind of database you're connecting to, so you can tell Confluence which dialect it needs to use.

Option 2: Datasource Connection — This asks the Java application server for a database connection. You will need to have configured a datasource into your application server.

Screenshot: Datasource Connection



Supply the following information:

- **Datasource Name** — The JNDI name of the datasource, as configured in the application server.
  Note: Some servers will have JNDI names like `jdbc/datasourcename`; others will be of the form `java:comp/env/jdbc/datasourcename`. Consult your application-server documentation.

You will also need to know:

- What kind of database you're connecting to, so you can tell Confluence which dialect it needs to use.

Next Step

Load content for the site

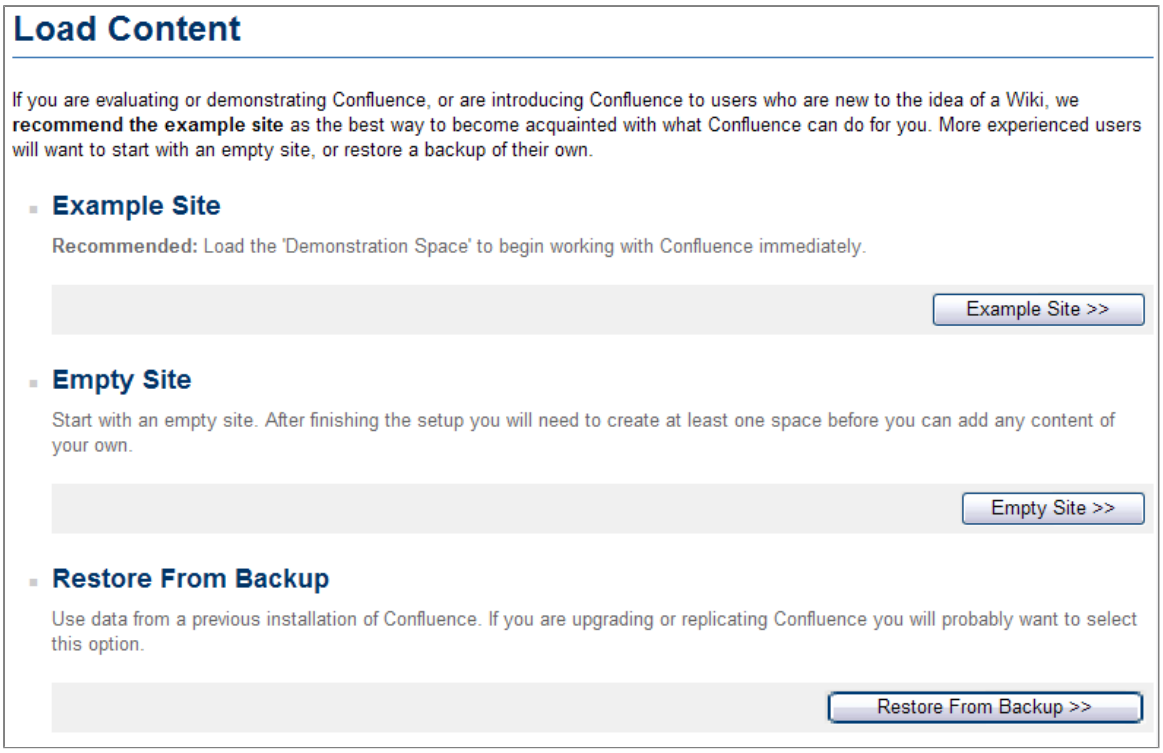
RELATED TOPICS

Database Configuration
Confluence Setup Guide
Confluence User's Guide
Confluence Documentation Home

Load Content for the Site

This page is part of the Confluence Setup Guide.

Screenshot: Load Content



Select one of the following options:

- **Example Site** — This option will load Confluence's 'Demonstration Space'. Select this if you are using Confluence for the first time, or if you want the Demonstration Space for your other Confluence users. The Demonstration Space helps to familiarise you with Confluence and what it can do for you. You can then continue using your Confluence deployment as normal — there's no need to reinstall later.
- **Empty Site** — Select this option if you are already familiar with Confluence. You will need to create at least one space before you can start adding content to the site.
- **Restore from Backup** — Select this option if you want to use Confluence data from a previous installation.

Next Steps

Restore your data from backup, if you have chosen that option.
Start using Confluence — see the Confluence User's Guide.

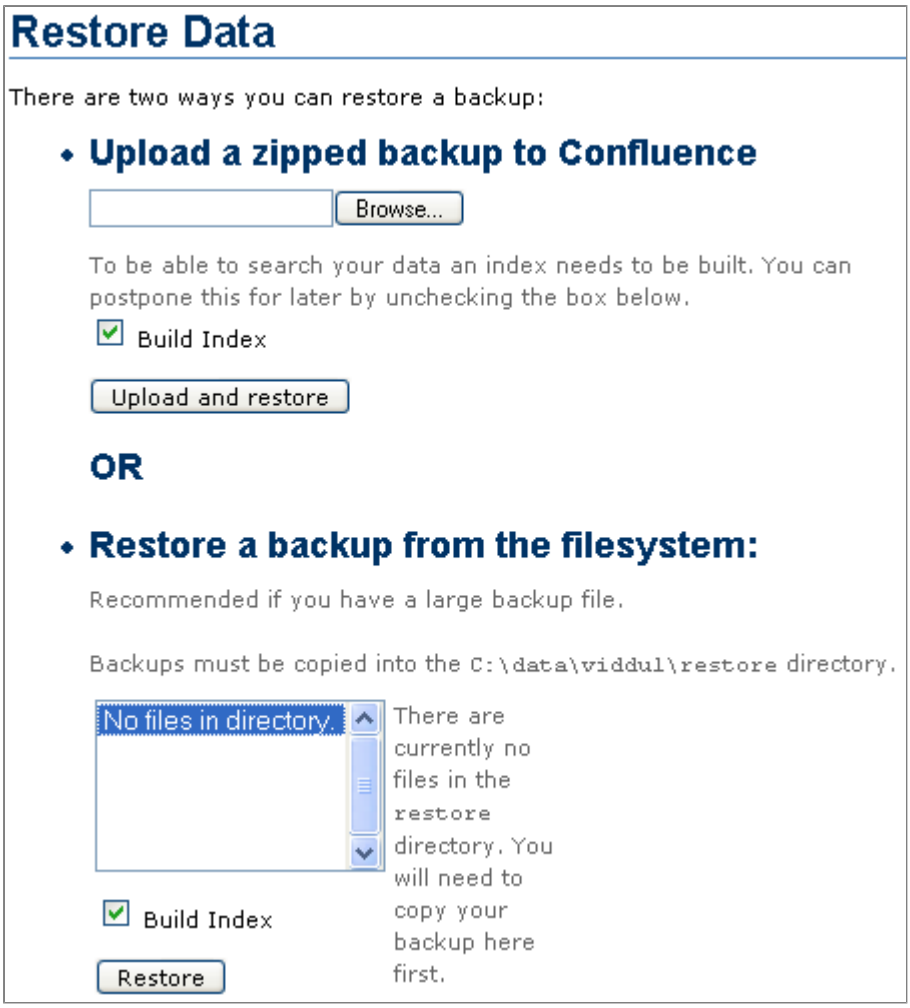
RELATED TOPICS

Confluence Setup Guide
Universal Wiki Converter
Confluence User's Guide
Confluence Documentation Home

Restoring from Backup During Setup

This page is part of the Confluence Setup Guide.

Screenshot : Restoring Data



This option allows you to reload your data from an existing Confluence installation into your new Confluence site during the initial setup procedure. You can choose to upload data from a zipped backup file, or to restore from a backup file on your file system.

Option 1: Upload a zipped backup to Confluence — This option will load the data from a zipped backup file.

To create a backup file from your existing version of Confluence, go to the '**Backup & Restore**' section of your **Administration Console**.

To restore from a zipped backup:

1. Browse for the relevant daily backup file or a file you have created via a manual backup.
2. Check '**Build Index**' to build the data index, used for the search.
3. Click the '**Upload and Restore**' button.

Option 2: Restore a backup from the filesystem — This option is recommended if you have a very large daily backup file (greater than 100MB), or a daily backup file that is already on the server and doesn't require uploading.

1. Copy the backup file into the `restore` directory inside your confluence Home directory and then refresh the page. You should now see your backup file appear on the '**Restore Data**' screen (pictured above), in the box beneath the heading '**Restore a backup from the filesystem**'.
2. Check '**Build Index**' to build the data index, used for the search.
3. Click the '**Restore**' button.

When the restore process has finished, you are ready to log in to Confluence. The system administrator account and all other information has been transferred from your previous Confluence installation.

RELATED TOPICS

Confluence Setup Guide
Confluence User's Guide
Confluence Documentation Home

Upgrading Confluence

This guide will tell you how to upgrade from one version of Confluence to a later version. Choose the type of installation you are upgrading, and follow the link to the installation instructions.

Installation Type	Description
Standalone distribution	The Standalone distribution is the simplest installation, using an Apache Tomcat application server bundled with the Confluence application. Choose this option if your existing Confluence installation is a Standalone deployment. If in doubt, this is the one you want.
EAR/WAR distribution	Use this guide if you have deployed Confluence onto your own existing application server.

Upgrading to Confluence 3.2?
If so, please review the Confluence 3.2 Upgrade Notes for important information about this version of Confluence.

Also, we strongly recommend that you check the upgrade notes for every major version of Confluence that you are skipping, since there might be specific changes between Confluence versions that could affect your Confluence installation. The upgrade notes for recent major versions of Confluence are accessible from the Upgrade Notes Overview page.

RELATED TOPICS

Confluence Installation Guide

Upgrading Confluence Standalone Distribution

This document tells you how to upgrade from one version of Confluence to a later version. These instructions apply to the **Standalone Distribution** of Confluence. The Standalone distribution includes Apache Tomcat as the standalone application server.

If you want to upgrade an **EAR/WAR distribution** deployed on your own existing application server, please refer to Upgrading Confluence EAR-WAR Distribution instead.

Please also check the following before you start using this guide:

- The **version of Confluence** that you will be upgrading to. Refer to the documentation home page to verify the latest Confluence version and to find documentation for older versions.
- The **supported platforms** for the version that you will be upgrading to. Please see the Supported Platforms page for the version of Confluence that you will be upgrading to, as well as the End of Support Announcements for Confluence.
- If you are running **Confluence on a cluster**, please see Upgrading a Confluence Cluster instead of this document.

Upgrading to Confluence 3.2?
If so, please review the Confluence 3.2 Upgrade Notes for important information about this version of Confluence.

Also, we strongly recommend that you check the upgrade notes for every major version of Confluence that you are skipping, since there might be specific changes between Confluence versions that could affect your Confluence installation. The upgrade notes for recent major versions of Confluence are accessible from the Upgrade Notes Overview page.

On this page:

- Before you Start
- Backing Up
- Testing the Upgrade in a Test Environment
- Performing the Upgrade
- Reapplying Customisations to your New Confluence
- Checking for Known Issues and Troubleshooting the Confluence Upgrade

Before you Start

Changing your Database?
If you are planning to change to a different database, we recommend that you complete the Confluence upgrade first. Then follow the instructions on migrating to a different database.

1. Note that you need current software maintenance to perform the upgrade.
2. Confirm that your license support period is still valid before you try to upgrade.
3. If your current license has expired but you have a new license with you, please update your license in Confluence before performing the upgrade.
   If you forget to do this and your license has expired, you will receive errors during the upgrade process. Refer to the instructions on upgrading beyond current license period.
4. Check the release notes for the new version of Confluence you are installing, plus the upgrade notes for any major versions you are skipping. It is important to read these upgrade notes as there might be specific changes between Confluence versions that could

affect your Confluence instance. The upgrade notes pages for recent major versions of Confluence are accessible from the Upgrade Notes Overview page. (Each upgrade notes page is a 'child' of its respective release notes page.)

If you are upgrading from a version of Confluence prior to 2.5.5, the upgrade notes information is located under a heading on the release notes pages.

5. Make sure that your environment (e.g. the database system, the operating system, the application server and so on) still complies with the Confluence System Requirements. A newer version of Confluence may have different requirements than the previous version.
6. If you are using Confluence EAR-WAR edition, check Installing the Confluence EAR-WAR Edition for your specific application server, to see if there is anything extra you will need to do to get Confluence running. For example:
    - Resin 3 users will need to update web.xml.
7. If you are using an external database, familiarize yourself with all known issues for your specific database. Also make sure the Confluence database connector principal (the database user login) has sufficient permissions to modify the database schema.
8. Note which plugins are installed/enabled on your current Confluence instance. Please verify whether a compatible version of the plugin is available in the version of Confluence you are upgrading to. This information is available on the respective home pages for these plugins on the Confluence Extension space. Once you have confirmed the availability of compatible versions, you should upgrade your plugins after successfully upgrading Confluence. This can be done via the **'Plugin Repository'** in your **Administration Console**. Please test these first by applying them to the latest Confluence version in a test environment.
9. If you have made any customisations to Confluence, please verify their compatibility in the latest version. For example, if you have modified any layouts or are using your own custom theme, please test these first by applying them to the latest Confluence version in a test environment.

Backing Up

Before you begin the Confluence upgrade, you must back up the following:

1. **Back up your Confluence Home directory.**
   The Confluence Home directory is the folder where Confluence stores its configuration information, search indexes and page attachments. If you're using the embedded HSQLDB database supplied for evaluation purposes, the database files are also stored in this directory.

   Tip: Another term for 'Home directory' would be 'data directory'.

   The location of the Home directory is stored in a configuration file called `confluence-init.properties`, which is located inside the `confluence/WEB-INF/classes` directory in your Confluence Installation directory.
2. **Back up your database.** Perform a manual backup of your external database before proceeding with the upgrade, and double check that the backup was actually created properly. If you are not a database expert, or unfamiliar with the backup-restore facilities of your database, you should try to restore the backup to a different system to ensure the backup worked, before proceeding. This recommendation is not specific to Confluence usage but just common sense: Surprisingly many companies, even banks, get in trouble for broken database backups.

   The 'embedded database' is the HSQLDB database supplied with Confluence for evaluation purposes, you don't need to back it up since it is stored in the home directory. But you should not use this database for production systems anyway, so if you happen to accidentally still use HSQLDB in a production system, please migrate to a proper database **before** the upgrade.
3. **Back up your Confluence Installation directory** (if you are using Confluence Standalone) **or your Confluence webapp** (if you are using Confluence EAR-WAR edition).
   The 'Confluence Installation directory' is the directory into which the Confluence application files and libraries have been unpacked (unzipped) when Confluence was installed. Confluence does not modify or store any data in this directory. This directory is also sometimes called the 'Confluence Install directory'.

Testing the Upgrade in a Test Environment

Be sure to test the upgrade in a test environment before proceeding on your production server.

1. Create a snapshot of your current production Confluence environment on a test server, as described in the page on Moving Confluence Between Servers.
2. Perform the upgrade on your cloned environment.
3. Test all your unsupported plugins and any customisations with the new version before proceeding on your production server. You can read more about supported and unsupported plugins.

Performing the Upgrade

If you are migrating servers or migrating databases, perform those operations in separate steps.

To install Confluence, unzip the new Confluence installation zip file into a directory of your choice and then edit the configuration files to point your new installation to your existing data files. Follow these instructions:

1. Shut down your existing Confluence instance.
2. Download the Confluence Standalone zip file.
3. **If you are on Windows**, please check your unzip program before extracting the downloaded zip file. Some archive-extract programs cause errors when unzipping the Confluence zip file. You should use a third-party unzip program like 7Zip or Winzip. If you do not have one, please download and install one before continuing:
    - 7Zip — Recommended. If in doubt, download the '32-bit.exe' version
    - Winzip
4. Use your unzip program to unzip the installation file. You should now have a new directory called `confluence-<version>`, e.g.

`confluence-2.9.2-std`.
- In the rest of this document, we will refer to this as the `<Installation-Directory>`.
- If you decide to change the location from the default, make sure that you choose a different location from your existing Confluence installation, because legacy files may cause problems if you install the new Confluence version into an existing directory.
- Do not use spaces in your directory path.
- You can read more about the Confluence Installation directory.

5. Edit the `confluence-init.properties` file found at:
`<Installation-Directory>\confluence\WEB-INF\classes\confluence-init.properties`
and update '`confluence.home`' to point to your **existing** Confluence Home directory.
   - You can read more about the Confluence Home directory.
   - Make sure you have first backed up this directory, as instructed above.
   - Open the `confluence-init.properties` file in a text editor such as Notepad.
   - Scroll to the bottom and find this line:

```
1.# confluence.home=c:/confluence/data
```

   - Remove the '#' and the space at the beginning of this line, so that Confluence no longer regards the line as a comment. The line should now begin with `confluence.home`.
   - Update the directory name after the = sign, to point to your existing Confluence Home directory.
6. If you are running Confluence as a Windows service, use `<Installation-Directory>\bin\tomcat5w.exe` to remove and re-install the Tomcat service. Or use the command prompt and type `<Installation-Directory>\bin\service.bat remove Confluence`.

> ⚠ It is vital that you stop and remove the existing service *prior to uninstalling* the old instance of Confluence! For more information on running Confluence as Windows service, please refer to the Start Confluence automatically on Windows as a Service topic.

⚠ To remove the service installed by Confluence Auto Installer, you need to run the `<confluence auto installer installation folder>\UninstallService.bat`.

7. If you are using an external database (i.e. not the embedded HSQLDB database supplied for evaluation purposes), copy the jdbc driver jar file from your old Confluence Standalone installation to the new Confluence Standalone installation. The jdbc driver jar file in the old Confluence Standalone installation should be located in either the `<Install-Directory>/common/lib` or `<Installation-Directory>/confluence/WEB-INF/lib` directories. Once you have identified this file, copy it to either the `<Install-Directory>/lib` or `<Installation-Directory>/confluence/WEB-INF/lib` directories of your Confluence 3.x installation.
8. If you have delegated your user management to JIRA, LDAP, Crowd, or any other external user management system, copy the following files from your old Confluence installation to your new Confluence installation:
   - `<Installation-Directory>/confluence/WEB-INF/classes/osuser.xml`.
   - `<Installation-Directory>/confluence/WEB-INF/classes/atlassian-user.xml` (if you are upgrading from Confluence 2.2 or later).

> ℹ If you are upgrading from an earlier version of Confluence (2.5.5 and earlier) and are copying your existing `atlassian-user.xml` file from your previous instance, please ensure that the hibernate cache parameter in this file has been enabled, to avoid performance related issues. (NOTE: If you use Crowd for your user management, you do not need to do this.):
>
> ```
> <hibernate name="Hibernate Repository" key="hibernateRepository"
> description="Hibernate Repository" cache="true" />
> ```

9. If you have delegated your user management to Crowd, you will also need to copy the Crowd client library and configuration files from your old Confluence installation to your new Confluence installation:
`<Installation-Directory>/confluence/WEB-INF/lib/crowd-integration-client-X.X.X.jar` and
`<Installation-Directory>/confluence/WEB-INF/classes/crowd.properties`. If you need more information, please refer to the Crowd documentation.
10. Consider any adjustments you need to make to customisations and special configurations, as described below.
⚠ Your new version of Confluence may not function correctly or could encounter problems or errors if these are not implemented.
11. Start your new version of Confluence.
ℹ **Please note** that Confluence will need to re-index attachments and this can take 5-10 minutes. Please wait until Confluence has finished indexing the attachments before trying to access Confluence via your web browser.
12. During the startup process Confluence will create any missing database indexes. If you created any database indexes on your own, please check those afterwards and remove those that duplicate the indexes added by Confluence. Just in case you run into any errors which prevent Confluence from starting up, you can set the system property `hibernate.hbm2ddl.skip_creating_missing_indexes` to `true` to skip automatic index creation.
13. Visit Confluence in your web browser and log in using a username from your previous Confluence installation. You should be able to log in immediately, without seeing the Setup Wizard.
14. Take a quick look around your Confluence site to confirm that all your spaces and pages are present and everything looks normal. You should see the new Confluence version number in the page footer.

Reapplying Customisations to your New Confluence



Hint: The steps below are for advanced Confluence users, who have applied special settings to their Confluence server and/or Confluence look and feel

After upgrading your Confluence installation to a later version of Confluence, you need to consider any customisations you have applied to your system and other special configurations:

- If you had previously installed **Confluence/Tomcat as a Windows service**, uninstall the service (to ensure that the old Confluence cannot start automatically when the server restarts) and reinstall the new one. For details please see Start Confluence automatically on Windows as a Service.
- If you are using a Standalone Edition of Confluence and you have previously defined a **CATALINA_HOME environment variable**, please check that it points to the correct path for the new Confluence Tomcat server.
- If you had previously connected your Confluence installation to an external database via a JNDI datasource or you implemented SSL , edit your new `web.xml` file and and copy over any relevant modifications from your old `web.xml` file, which relate to these customisations.
- If you were previously running **Confluence on a non-standard port**, edit your new `<Installation-Directory>\conf\server.xml` file as described in Change listen port for Confluence Standalone.
- If you had previously defined a **Tomcat datasource**, edit your new `<Installation-Directory>\conf\server.xml` and copy over the datasource definition from your old `server.xml`.
- If you were previously using any **plugins**, install the latest compatible version and disable any plugins that are incompatible with your new version of Confluence. The easiest way to do this is to use the Plugin Repository in the Confluence Administration Console.
- If you are using any **customised themes**, please check that they are displaying as expected. Some further customisation may be required to ensure compatibility with your new version of Confluence.
- If you had previously customised the **default site or space layouts**, you will need to reapply your changes to the new defaults as described here.
- If you had previously modified the Confluence **source code**, you will need to reapply your changes to the new version.
- If you were previously running **Confluence over SSL**, you will need to reapply your configuration as described in Adding SSL for Secure Logins and Page Security.
- If you had previously modified the **memory flags** (`Xms` and `Xmx`) in either the `<Installation-Directory>\bin\setenv.sh` or the `<Installation-Directory>\bin\setenv.bat` file, you may want to make the modifications in your new installation. The parameters are specified in the `JAVA_OPTS` variable.
- If you had changed the Confluence interface text, you will need to pull over the ConfluenceActionSupport.properties file.
- If you were using a custom SSO authenticator or the utility to Automatically Add LDAP users to the confluence-users Group, change seraph-config.xml to the correct authenticator.

Checking for Known Issues and Troubleshooting the Confluence Upgrade

- **Check for known issues**. Sometimes we find out about a problem with the latest version of Confluence after we have released the software. In such cases we publish information about the known issues in the Confluence Knowledge Base. Please check the important technical advisories on the front page of the Knowledge Base and follow the instructions to apply any necessary patches.

- **Did you encounter a problem during the Confluence upgrade?** Please refer to the guide to troubleshooting upgrades in the Confluence Knowledge Base.

RELATED TOPICS

Upgrading Confluence
Upgrading Confluence EAR-WAR Distribution
Confluence Installation Guide
Important Directories and Files
Site Backup and Restore
Database Configuration

Administrators Guide Home Confluence Documentation Home

Upgrading Confluence EAR-WAR Distribution

This document tells you how to upgrade from one version of Confluence to a later version. These instructions apply to the **EAR-WAR Distribution** of Confluence, deployed on your own existing application server.

If you want to upgrade your **Confluence Standalone distribution**, which includes Apache Tomcat as the standalone application server, please refer to Upgrading Confluence Standalone Distribution instead.

Please also check the following before you start using this guide:

- The **version of Confluence** that you will be upgrading to. Refer to the documentation home page to verify the latest Confluence version and to find documentation for older versions.
- The **supported platforms** for the version that you will be upgrading to. Please see the Supported Platforms page for the version of Confluence that you will be upgrading to, as well as the End of Support Announcements for Confluence.
- If you are running **Confluence on a cluster**, please see Upgrading a Confluence Cluster instead of this document.

Upgrading to Confluence 3.2?
If so, please review the Confluence 3.2 Upgrade Notes for important information about this version of Confluence.

Also, we strongly recommend that you check the upgrade notes for every major version of Confluence that you are skipping, since there might be specific changes between Confluence versions that could affect your Confluence installation. The upgrade notes for recent major versions of Confluence are accessible from the Upgrade Notes Overview page.

On this page:

- Before you Start
- Backing Up
- Testing the Upgrade in a Test Environment
- Performing the Upgrade
- Reapplying Customisations to your New Confluence
- Checking for Known Issues and Troubleshooting the Confluence Upgrade

Before you Start

Changing your Database?
If you are planning to change to a different database, we recommend that you complete the Confluence upgrade first. Then follow the instructions on migrating to a different database.

1. Note that you need current software maintenance to perform the upgrade.
2. Confirm that your license support period is still valid before you try to upgrade.
3. If your current license has expired but you have a new license with you, please update your license in Confluence before performing the upgrade.

   If you forget to do this and your license has expired, you will receive errors during the upgrade process. Refer to the instructions on upgrading beyond current license period.
4. Check the release notes for the new version of Confluence you are installing, plus the upgrade notes for any major versions you are skipping. It is important to read these upgrade notes as there might be specific changes between Confluence versions that could affect your Confluence instance. The upgrade notes pages for recent major versions of Confluence are accessible from the Upgrade Notes Overview page. (Each upgrade notes page is a 'child' of its respective release notes page.)

   If you are upgrading from a version of Confluence prior to 2.5.5, the upgrade notes information is located under a heading on the release notes pages.
5. Make sure that your environment (e.g. the database system, the operating system, the application server and so on) still complies with the Confluence System Requirements. A newer version of Confluence may have different requirements than the previous version.
6. If you are using Confluence EAR-WAR edition, check Installing the Confluence EAR-WAR Edition for your specific application server, to see if there is anything extra you will need to do to get Confluence running. For example:
   - Resin 3 users will need to update web.xml.
7. If you are using an external database, familiarize yourself with all known issues for your specific database. Also make sure the Confluence database connector principal (the database user login) has sufficient permissions to modify the database schema.
8. Note which plugins are installed/enabled on your current Confluence instance. Please verify whether a compatible version of the plugin is available in the version of Confluence you are upgrading to. This information is available on the respective home pages for these plugins on the Confluence Extension space. Once you have confirmed the availability of compatible versions, you should upgrade your plugins after successfully upgrading Confluence. This can be done via the '**Plugin Repository**' in your **Administration Console**. Please test these first by applying them to the latest Confluence version in a test environment.
9. If you have made any customisations to Confluence, please verify their compatibility in the latest version. For example, if you have modified any layouts or are using your own custom theme, please test these first by applying them to the latest Confluence version in a test environment.

Backing Up

Before you begin the Confluence upgrade, you must back up the following:

1. **Back up your Confluence Home directory.**
   The Confluence Home directory is the folder where Confluence stores its configuration information, search indexes and page attachments. If you're using the embedded HSQLDB database supplied for evaluation purposes, the database files are also stored in this directory.

   Tip: Another term for 'Home directory' would be 'data directory'.
   The location of the Home directory is stored in a configuration file called `confluence-init.properties`, which is located inside the `confluence/WEB-INF/classes` directory in your Confluence Installation directory.
2. **Back up your database.** Perform a manual backup of your external database before proceeding with the upgrade, and double check that the backup was actually created properly. If you are not a database expert, or unfamiliar with the backup-restore facilities of your database, you should try to restore the backup to a different system to ensure the backup worked, before proceeding. This recommendation is not specific to Confluence usage but just common sense: Surprisingly many companies, even banks, get in trouble for broken database backups.

   The 'embedded database' is the HSQLDB database supplied with Confluence for evaluation purposes, you don't need to back it up since it is stored in the home directory. But you should not use this database for production systems anyway, so if you happen to accidentally still use HSQLDB in a production system, please migrate to a proper database **before** the upgrade.

3. **Back up your Confluence Installation directory** (if you are using Confluence Standalone) **or your Confluence webapp** (if you are using Confluence EAR-WAR edition).
   The 'Confluence Installation directory' is the directory into which the Confluence application files and libraries have been unpacked (unzipped) when Confluence was installed. Confluence does not modify or store any data in this directory. This directory is also sometimes called the 'Confluence Install directory'.

Testing the Upgrade in a Test Environment

Be sure to test the upgrade in a test environment before proceeding on your production server.

1. Create a snapshot of your current production Confluence environment on a test server, as described in the page on Moving Confluence Between Servers.
2. Perform the upgrade on your cloned environment.
3. Test all your unsupported plugins and any customisations with the new version before proceeding on your production server. You can read more about supported and unsupported plugins.

Performing the Upgrade

If you are migrating servers or migrating databases, perform those operations in separate steps.

The upgrade process allows you to unzip the new Confluence installation into a directory of your choice and then edit the configuration files to point your new installation to your existing data files. Follow these instructions:

1. Shut down your existing Confluence instance.
2. Download the Confluence EAR-WAR zip file: Go to the Download Centre, and click '**Show all**' to find the EAR-WAR zip file.
3. **If you are on Windows**, please check your unzip program before extracting the downloaded zip file. Some archive-extract programs cause errors when unzipping the Confluence zip file. You should use a third-party unzip program like 7Zip or Winzip. If you do not have one, please download and install one before continuing:
   - 7Zip — Recommended. If in doubt, download the '32-bit.exe' version
   - Winzip
4. Use your unzip program to unzip the installation file. You should now have a new directory called `confluence-<version>`.
   - In the rest of this document, we will refer to this as the `<Installation-Directory>`.
   - Do not use spaces in your directory path.
   - You can read more about the Confluence Installation directory.
5. Edit the `confluence-init.properties` file found at:
   `<Installation-Directory>\confluence\WEB-INF\classes\confluence-init.properties`
   and update '`confluence.home`' to point to your existing Confluence Home directory.
   - Make sure you have first backed up your Home directory.
   - Open the `confluence-init.properties` file in a text editor such as Notepad.
   - Scroll to the bottom and find this line:

     ```
     1.# confluence.home=c:/confluence/data
     ```

   - Remove the '#' and the space at the beginning of this line, so that Confluence no longer regards the line as a comment. The line should now begin with `confluence.home`.
   - Update the directory name after the = sign, to point to your existing Confluence Home directory.
6. If you are using Tomcat, you need to update either your `confluence.xml` or `server.xml` (depending on where you have defined the Confluence context descriptor) to point to the location of the new Confluence installation (also remember to copy over any customisations such as a **tomcat datasource** if you have one).
7. If you have delegated your user management to JIRA, LDAP or any other external user management system, copy the following files from your old Confluence installation to your new Confluence installation:
   - `<Installation-Directory>/confluence/WEB-INF/classes/osuser.xml`.
   - `<Installation-Directory>/confluence/WEB-INF/classes/atlassian-user.xml` (if you are upgrading from Confluence 2.2 or later).

     If you are upgrading from an earlier version of Confluence (2.5.5 and earlier) and are copying your existing `atlassian-user.xml` file from your previous instance, please ensure that the hibernate cache parameter in this file has been enabled, to avoid performance related issues. (NOTE: If you use Crowd for your user management, you do not need to do this.):

     ```
     <hibernate name="Hibernate Repository" key="hibernateRepository"
     description="Hibernate Repository" cache="true" />
     ```

8. If you have delegated your user management to Crowd, you will also need to copy the Crowd client library and configuration files from your old Confluence installation to your new Confluence installation:
   `<Installation-Directory>/confluence/WEB-INF/lib/crowd-integration-client-X.X.X.jar` and `<Installation-Directory>/confluence/WEB-INF/classes/crowd.properties`. If you need more information, please refer to the Crowd documentation.

9. Restart your application server and start Confluence.
   Please note that Confluence will need to re-index attachments and this can take 5-10 minutes. Please wait until Confluence has finished indexing the attachments before trying to access Confluence via your web browser. (There is no easy and quick way to determine if the indexing process is completed. Please wait for approximately 10 minutes after the server start up before accessing Confluence via a web browser.)
10. During the startup process Confluence will create any missing database indexes. If you created any database indexes on your own, please check those afterwards and remove those that duplicate the indexes added by Confluence. Just in case you run into any errors which prevent Confluence from starting up, you can set the system property `hibernate.hbm2ddl.skip_creating_missing_indexes` to `true` to skip automatic index creation.
11. Visit Confluence in your web browser and log in using a username from your previous Confluence installation. You should be able to log in immediately, without seeing the Setup Wizard.
12. Take a quick look around your Confluence site to confirm that all your spaces and pages are present and everything looks normal. You should see the new Confluence version number in the page footer.
13. Consider any adjustments you need to make to customisations and special configurations, as described below.

Reapplying Customisations to your New Confluence

Hint: The steps below are for advanced Confluence users, who have applied special settings to their Confluence server and/or Confluence look and feel

After upgrading your Confluence installation to a later version of Confluence, you need to consider any customisations you have applied to your system and other special configurations:

- If you had previously installed **Confluence/Tomcat as a Windows service**, uninstall the service (to ensure that the old Confluence cannot start automatically when the server restarts) and reinstall the new one. For details please see Start Confluence automatically on Windows as a Service.
- If you are using a Standalone Edition of Confluence and you have previously defined a **CATALINA_HOME environment variable**, please check that it points to the correct path for the new Confluence Tomcat server.
- If you had previously connected your Confluence installation to an external database via a JNDI datasource or you implemented SSL , edit your new `web.xml` file and and copy over any relevant modifications from your old `web.xml` file, which relate to these customisations.
- If you were previously running **Confluence on a non-standard port**, edit your new `<Installation-Directory>\conf\server.xml` file as described in Change listen port for Confluence Standalone.
- If you had previously defined a **Tomcat datasource**, edit your new `<Installation-Directory>\conf\server.xml` and copy over the datasource definition from your old `server.xml`.
- If you were previously using any **plugins**, install the latest compatible version and disable any plugins that are incompatible with your new version of Confluence. The easiest way to do this is to use the Plugin Repository in the Confluence Administration Console.
- If you are using any **customised themes**, please check that they are displaying as expected. Some further customisation may be required to ensure compatibility with your new version of Confluence.
- If you had previously customised the **default site or space layouts**, you will need to reapply your changes to the new defaults as described here.
- If you had previously modified the Confluence **source code**, you will need to reapply your changes to the new version.
- If you were previously running **Confluence over SSL**, you will need to reapply your configuration as described in Adding SSL for Secure Logins and Page Security.
- If you had previously modified the **memory flags** (`Xms` and `Xmx`) in either the `<Installation-Directory>\bin\setenv.sh` or the `<Installation-Directory>\bin\setenv.bat` file, you may want to make the modifications in your new installation. The parameters are specified in the `JAVA_OPTS` variable.
- If you had changed the Confluence interface text, you will need to pull over the ConfluenceActionSupport.properties file.
- If you were using a custom SSO authenticator or the utility to Automatically Add LDAP users to the confluence-users Group, change seraph-config.xml to the correct authenticator.

Checking for Known Issues and Troubleshooting the Confluence Upgrade

- **Check for known issues**. Sometimes we find out about a problem with the latest version of Confluence after we have released the software. In such cases we publish information about the known issues in the Confluence Knowledge Base. Please check the important technical advisories on the front page of the Knowledge Base and follow the instructions to apply any necessary patches.

- **Did you encounter a problem during the Confluence upgrade?** Please refer to the guide to troubleshooting upgrades in the Confluence Knowledge Base.

RELATED TOPICS

Upgrading Confluence
Upgrading Confluence Standalone Distribution
Confluence Installation Guide
Important Directories and Files
Site Backup and Restore
Database Configuration

Administrators Guide Home Confluence Documentation Home

Upgrading Beyond Current Licensed Period

This page explains the recovery process should you mistakenly try to upgrade your Confluence installation to a version beyond your current license entitlement.

The information on this page applies to Confluence 2.7.2 and later.

License Warnings

During an upgrade an obvious indication that your license has expired can be found in your log file. You will see a 'WARN' level entry similar to this:

```
[confluence.upgrade.impl.DefaultUpgradeManager] isUpgradeAllowed
Your license is now outside of it's support period. You need to renew the license before you can
upgrade to this version of Confluence.
```

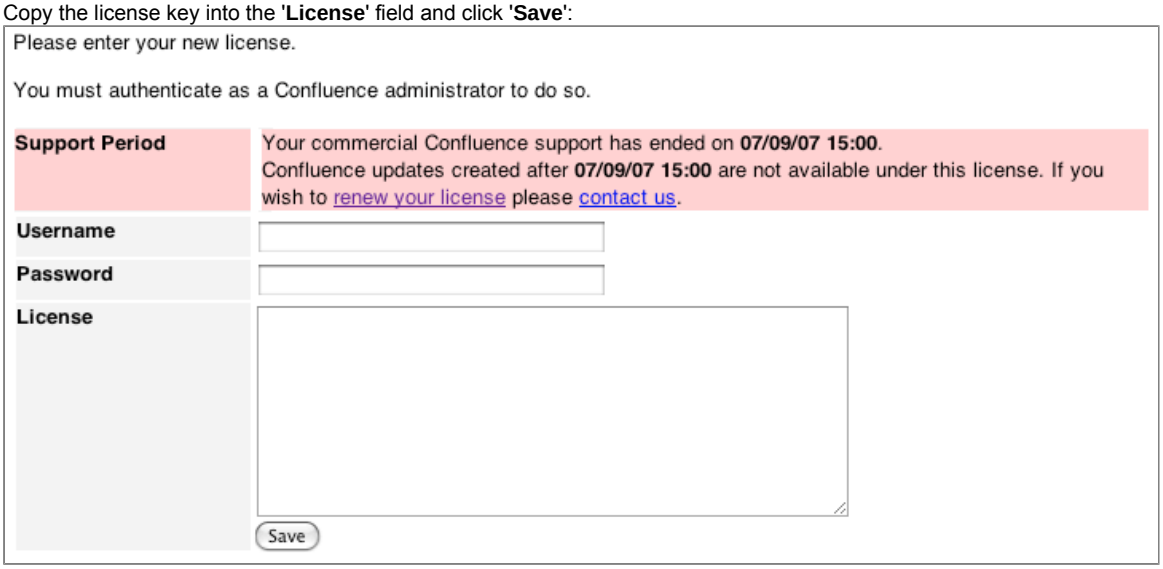
When you try to connect to the Confluence instance, you will see the following warning screen:

You cannot access Confluence at present. Look at the table below to identify the reasons

Time	Level	Type	Description	Exception
2008-02-04 10:51:04	null	(EventType: upgrade)	Cannot proceed with upgrade. Your current license does not entitle you to upgrade to this version of Confluence. Please check that the support period of your license has not expired or that you have the correct partner license. If you wish to renew your license, please contact sales@atlassian.com. If you have a new license, please enter it on this page and **restart**.	`fatal`

Updating the Confluence License

1. Contact Atlassian Sales to arrange for a new license to be issued, as instructed on the warning screen illustrated above.
2. Once you have received a suitable license, supply the license key to Confluence:
    - Click the link given on the license warning screen, illustrated above.
    - You will first be asked to log in as a Confluence administrator.
    - Then you will be presented with a simplified license administration screen. Enter the credentials of a Confluence system administrator.
    - Copy the license key into the '**License**' field and click '**Save**':



3. Restart Confluence to continue the upgrade.

RELATED TOPICS

Upgrading Confluence

Administrators Guide Home Confluence Documentation Home

Confluence Release Cycle

New versions of Confluence are released frequently. Our goals are to:

- Make bug-fixes available to customers sooner
- Give interested customers early access to new features and API changes
- Make Confluence major releases predictable

Feature Releases

We aim to release new versions of Confluence every three to four months. These releases will contain the bulk of new functionality.

Feature releases are numbered by incrementing Confluence's minor version number, so the move from Confluence 2.0 to 2.1 and 2.1 to 2.2 both introduced significant new features to the product. Occasionally we may change to a whole new major version number (Confluence 2.0 was originally slated to be released as 1.5), but that is mostly done for marketing purposes, and shouldn't be considered to have any practical meaning.

Feature releases may not be API-compatible with the previous release. This means that you should test RPC clients, macros and plugins before running them on a newer version of Confluence.

You can find the time line history of our major releases at the downloads archive.

Bug-Fix Releases

Confluence bug-fix releases are scheduled every three to four weeks, depending on the number and urgency of the bugs that have been fixed during that particular development cycle. We aim to minimise the time between a bug being reported and a fix being available, without either us or our customers having to manage clumsy sets of manual patches.

Bug-fix releases will contain mostly bug-fixes, plus the occasional minor new feature or enhancement. Enhancements will be limited, however, as the main aim of these point releases is to improve stability, and make no significant API changes.

Bug-fix releases are numbered by incrementing the patch-level. So the first bug-fix release after Confluence 2.2 is 2.2.1, followed by 2.2.2. Occasionally, we will re-issue a bug-fix release because something was faulty with the original download. In that case we will create a "re-issue" release number, for example 2.1.5a or 2.2.1a.

Obviously, we don't expect anyone to upgrade Confluence every two weeks, administrators should keep their own schedule, based on how much of an inconvenience is being caused by any bugs that may have been fixed since. Sometimes, however, a security issue or serious application bug will arise that we feel it is in everyone's best interests to fix as soon as possible. In such cases, we will recommend in the Release Notes that all customers upgrade to the latest version.

Milestone Releases

Occasionally, when possible, we will release preview "milestone releases" of the next major Confluence version. How often and when we do so depends on the particulars of the current development cycle. In situations where we are working on a number of disparate features we may be able to do a number of progressive development releases, whereas in iterations where we are making significant changes to the Confluence internals, we may not have anything suitable for public consumption until quite late in the release cycle.

Milestone releases will be announced on the Development Releases page, and to the confluence-developer mailing list. Milestone releases are published for testing plugins and early feedback about our work, please don't use them on production systems.

The version number of a Milestone Release will be the version number of the *next* major release, suffixed with -m. So Confluence 2.3-m1 will be followed by 2.3-m2, and so on until the ultimate release of the finished Confluence 2.3.

Release Notes

Confluence 3.1.2

Confluence 3.1.2 is a recommended upgrade which fixes a number of bugs. Please see the 'Updates and Fixes in this Release' section below for details.

We identified a bug (see CONF-18437) that caused errors to be thrown whenever the macro browser was used in a theme other than the default theme. This bug also caused the preview mode to hang when the preview tab was clicked (i.e. the loading icon would spin forever without the preview mode loading). This has been fixed for all themes bundles with Confluence. If you are using a custom theme, you will need to implement a custom workaround, as described in the related JIRA issue.

Read the full release notes here.

All Release Notes

Confluence 3.1

- Confluence 3.1.2 Release Notes
- Confluence 3.1.1 Release Notes
- Confluence 3.1 Release Notes

Confluence 3.0

- Confluence 3.0.2 Release Notes
- Confluence 3.0.1 Release Notes

- Confluence 3.0 Release Notes

Confluence 2.10

- Confluence 2.10.4 Release Notes
- Confluence 2.10.3 Release Notes
- Confluence 2.10.2 Release Notes
- Confluence 2.10.1 Release Notes
- Confluence 2.10 Release Notes

Confluence 2.9

- Confluence 2.9.3 Release Notes
- Confluence 2.9.2 Release Notes
- Confluence 2.9.1 Release Notes
- Confluence 2.9 Release Notes

Confluence 2.8

- Confluence 2.8.3 Release Notes
- Confluence 2.8.2 Release Notes
- Confluence 2.8.1 Release Notes
- Confluence 2.8 Release Notes
- Confluence 2.8 Beta Release Notes

Confluence 2.7

- Confluence 2.7.4 Release Notes
- Confluence 2.7.3 Release Notes
- Confluence 2.7.2 Release Notes
- Confluence 2.7.1 Release Notes
- Confluence 2.7 Release Notes

Confluence 2.6

- Confluence 2.6.3 Release Notes
- Confluence 2.6.2 Release Notes
- Confluence 2.6.1 Release Notes
- Confluence 2.6 Release Notes

Confluence 2.5

- Release Notes 2.5.8
- Release Notes 2.5.7
- Release Notes 2.5.6
- Release Notes 2.5.5
- Release Notes 2.5.4
- Release Notes 2.5.3
- Release Notes 2.5.2
- Release Notes 2.5.1
- Release Notes 2.5

Confluence 2.4

- Release Notes 2.4.5
- Release Notes 2.4.4
- Release Notes 2.4.3
- Release Notes 2.4.2

Confluence 2.3

- Release Notes 2.3.3
- Release Notes 2.3.2
- Release Notes 2.3.1
- Release Notes 2.3

Confluence 2.2

- Release Notes 2.2.10
- Release Notes 2.2.9
- Release Notes 2.2.8
- Release Notes 2.2.7
- Release Notes 2.2.6a
- Release Notes 2.2.5
- Release Notes 2.2.4
- Release Notes 2.2.3
- Release Notes 2.2.2

- Release Notes 2.2.1
- Release Notes 2.2

Confluence 2.1

- Release Notes 2.1.5
- Release Notes 2.1.4
- Release Notes 2.1.3
- Release Notes 2.1.2
- Release Notes 2.1.1
- Release Notes 2.1

Confluence 2.0

- Release Notes 2.0.3
- Release Notes 2.0.2
- Release Notes 2.0.1
- Release Notes 2.0

Confluence 1.4

- Release Notes 1.4.4
- Release Notes 1.4.3
- Release Notes 1.4.2
- Release Notes 1.4.1
- Release Notes 1.4

Confluence 1.3

- Release Notes 1.3.6
- Release Notes 1.3.5
- Release Notes 1.3.4
- Release Notes 1.3.2
- Release Notes 1.3.1
- Release Notes 1.3

Confluence 1.2

- Release Notes 1.2.3
- Release Notes 1.2.2
- Release Notes 1.2.1
- Release Notes 1.2

Confluence 1.1

- Release Notes 1.1.2
- Release Notes 1.1.1
- Release Notes 1.1

Confluence 1.0

- Release Notes 1.0.3
- Release Notes 1.0.1
- Release Notes 1.0



Hint: Finding a list of known issues
To find a list of known issues in a particular Confluence version, you can create a filter in the Atlassian issue tracker and use t
the filtered report. The following example filter is the list of bugs reported for Confluence 2.7:

```
1.http:
//jira.atlassian.com/secure/IssueNavigator.jspa?reset=true&&type=1&pid=10470&fixfor=13115&ve
```

Read the JIRA documentation on creating filters.

Confluence Release Summary

This page shows the highlights of the major Confluence releases.

Current Release

For information about the latest release, please go to the Release Notes.

Confluence 3.2 — 24 March 2010

- Autocomplete for Inserting Links
- Autocomplete for Embedding Images and Documents
- A Link Browser that's Smarter, Smoother, Faster
- New Documentation Theme
- New Easy Reader Theme
- Template Bundles
- Reordering while Moving a Page
- New Keyboard Shortcuts and Editor Hints
- User Interface Enhancements
- And Even More Improvements

Confluence 3.1 — 8 December 2009

- Introducing Gadgets
- Drag-and-Drop
- Office 2007 Support
- New 'Move Page' Feature
- Enhanced Image Browser
- Draft Comparisons
- Page Restrictions Dialog Box
- Other Editor Enhancements
- New Web Browser Versions Supported
- Other Improvements

Confluence 3.0 — 1 June 2009

- Introducing the Macro Browser
- Enhanced User Profiles
- Introducing Your Network
- New User Status
- New Hover Profile Feature
- Customisable Enhanced PDF Exports
- Improved Rich Text Editor
- Performance Improvements
- Engine Room and Developer Community
- Administration Improvements
- More in the release notes

Confluence 2.10 — 3 December 2008

- Introducing the Widget Connector
- Improved Office Connector Now Bundled
- Introducing Quick Navigation
- 'Did You Mean', OpenSearch and More
- Custom Stylesheets for Confluence Spaces
- Updated JIRA Issues Macro with Custom Fields and Dynamic Display
- Enhanced User and Group Management
- Upgraded Rich Text Editor
- Universal Wiki Converter now with SharePoint Import and More
- Improved Activity Macros
- Plugin Framework 2
- More in the release notes

Confluence 2.9 — 7 August 2008

- Streamlined Search
- Auto Save
- Charts
- Page Tree
- Gallery
- New Tutorial
- More in the Menus
- Alphabetical Page Ordering
- Better Spam Prevention
- Plugin Repository
- Engine Room and Developers' Community
- More in the release notes

Confluence 2.8 — 10 April 2008

- Dynamic menus and simplified screen design
- Page ordering
- Collapsible comments
- Multiple-label filter
- Confluence installer
- Task list
- Performance enhancements
- Administration, management and monitoring

- More in the release notes

Confluence 2.7 — 12 December 2007

- JIRA Issues and Portlet macros use new trusted authentication
- Two-tier administrator permissions
- Inserting images and attaching files during page creation
- Sorting of images in Gallery macro
- Simplified and improved logging
- Performance, maintainability and administration
- More in the release notes

Confluence 2.6 — 27 September 2007

- Fresh look for the Default theme
- Personalised comments and Dashboard
- Space description on Dashboard
- Labels on templates
- Default content for space home pages
- Social Bookmarking plugin now bundled with Confluence
- Back-dating and renaming news items
- More in the release notes

Confluence 2.5 — 29 April 2007

- Introducing flexible page restrictions
- Dynamic task list JRE incompatibilities
- contentbylabel macro supports AND condition
- More in the release notes

Confluence 2.4 — 14 March 2007

- Editable comments
- Page mailing
- More in the release notes

Confluence 2.3 — 5 January 2007

- Confluence Massive — cluster support
- People directory
- Activity plugin — usage statistics
- Blogging RPC plugin — manage news in Confluence using blogger-compatible desktop clients
- WebDAV client support via WebDAV plugin — create, edit, move pages, attachments, etc via WebDAV
- More in the release notes

Confluence 2.2 — 27 April 2006

- Personal spaces
- Localisation/internationalisation — drop-in language packs (similar to JIRA)
- CAPTCHA support — spam protection
- Improved searching
- Improved LDAP performance
- Confluence Standalone ships with Tomcat 5.5
- More in the release notes

Confluence 2.1 — 20 December 2005

- Autosave
- Concurrent edit warnings
- LDAP integration with Atlassian User/POLIS
- More in the release notes

Confluence 2.0 — 17 November 2005

- Rich Text Editing — WYSIWYG editor
- Labels
- Dashboard tabs — All, My, Team, New
- RSS builder
- Export pages as Word documents
- Copy pages
- More in the release notes

Confluence 1.4 — 23 May 2005

- New user interface
- Enhanced editing — doing more in the edit interface
- Page permissions

- New plugin types
- Configurable themes
- Completely rewritten Wiki to HTML conversion engine
- More in the release notes

Confluence 1.3 — 30 November 2004

- Mail archiving
- Themes
- Trash can
- More granular space permissions
- More in the release notes

Confluence 1.2 — 23 August 2004

- Page list views — alphabetical, directory view and search view of all pages in a space
- Image thumbnails and thumbnail galleries
- Threaded comments
- Enhanced Search - indexing attachment comments and file names and contextual searching
- New permissions interface
- More in the release notes

Confluence 3.2 Release Notes

24 March 2010

With great pleasure, Atlassian presents the seemingly omniscient, drop dead gorgeous **Confluence 3.2**.

Highlights of this Release:

- Autocomplete for Inserting Links
- Autocomplete for Embedding Images and Documents
- A Link Browser that's Smarter, Smoother, Faster
- New Documentation Theme
- New Easy Reader Theme
- Template Bundles
- Reordering while Moving a Page
- New Keyboard Shortcuts and Editor Hints
- User Interface Enhancements
- And Even More Improvements

Responding to your Feedback:

100+ votes satisfied



- Thank you for all your issues and votes. Keep logging issues to help us keep improving!
- Below is a list of the highlights in this release.
- Attached is the full list of issues resolved in this release.

Upgrading from a previous version of Confluence

- Upgrading Confluence should be fairly straightforward. ***We strongly recommend that you back up your Confluence Home directory and your database before upgrading.***
- Please refer to the Confluence 3.2 upgrade notes for further essential information about **plugins** and other factors affecting your upgrade.

Highlights of Confluence 3.2



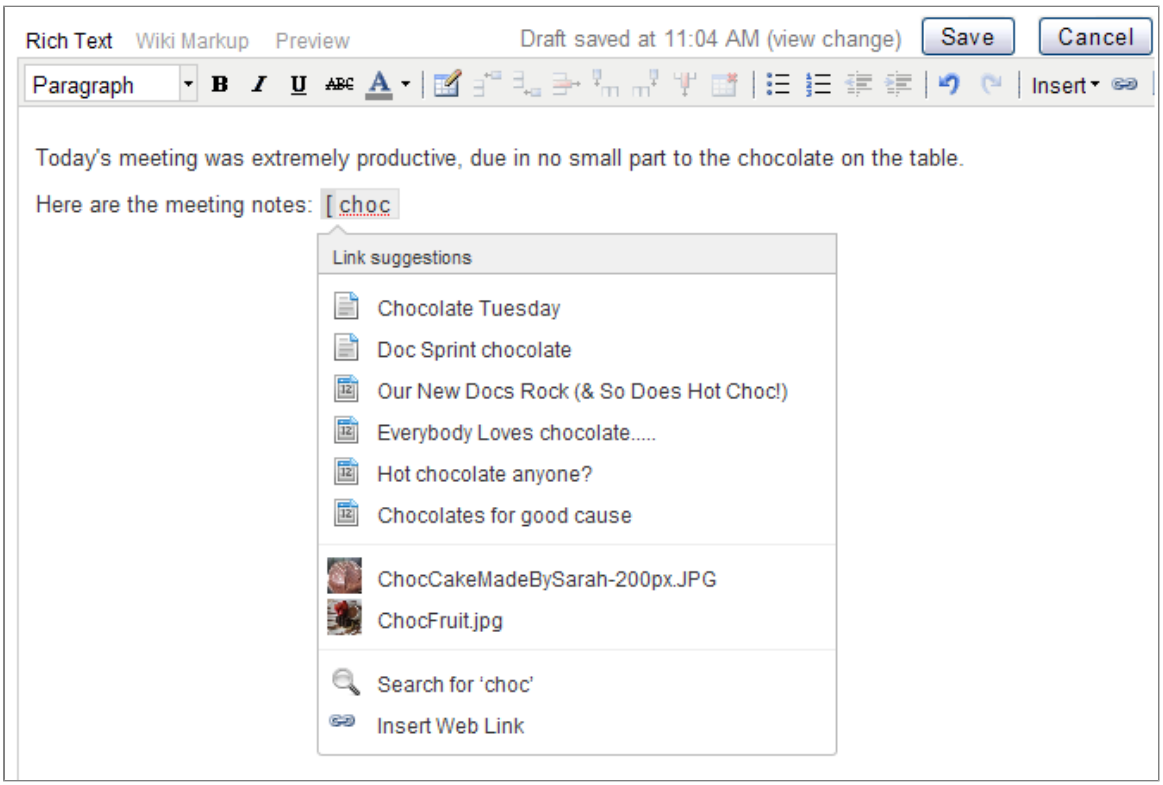
Autocomplete for Inserting Links

Wouldn't it be awesome to have the WYSIWYG ease of a rich text editor combined with the speed of a wiki markup editor? We're working

towards that sweetness. In the rich text editor you can now enter a trigger character to call up a list of suggested links to add to your page. A keyboard shortcut does it too.

Quick guide:

- Enter '**[**' and start typing to see the suggested links that match your text.
- Or press Ctrl+Shift+K to convert text to a link.



More....

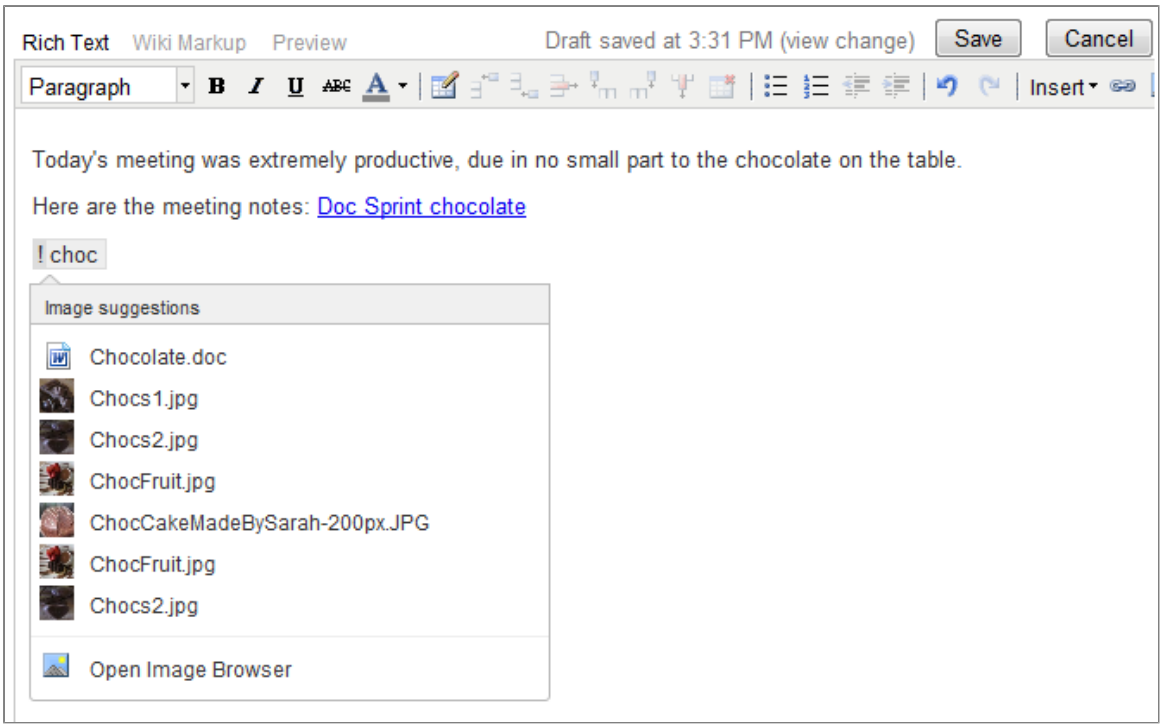


Autocomplete for Embedding Images and Documents

The new autocomplete also offers a quick way of adding images and documents to your page.

In the rich text editor:

- Enter '**!**' and start typing, to see a matching list of images and documents.
- Or press Ctrl+Shift+M immediately after a word or highlighted phrase.

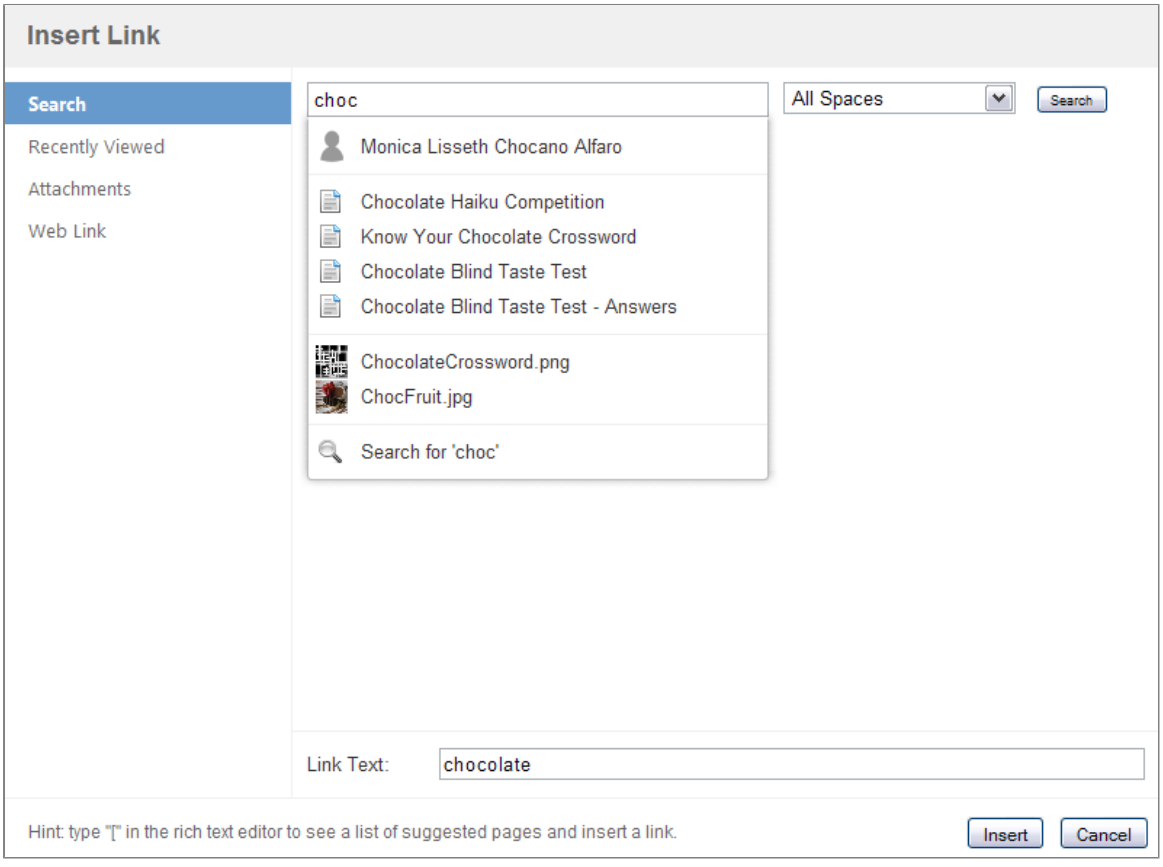


More....

3

A Link Browser that's Smarter, Smoother, Faster

Do you think the autocomplete is sweet? Good, because it's in the new link browser too. We've also made it easier and faster to link to images, attachments and recently-viewed pages. The link browser pops up when you click the 'Insert Link' icon in the editor toolbar.



More...

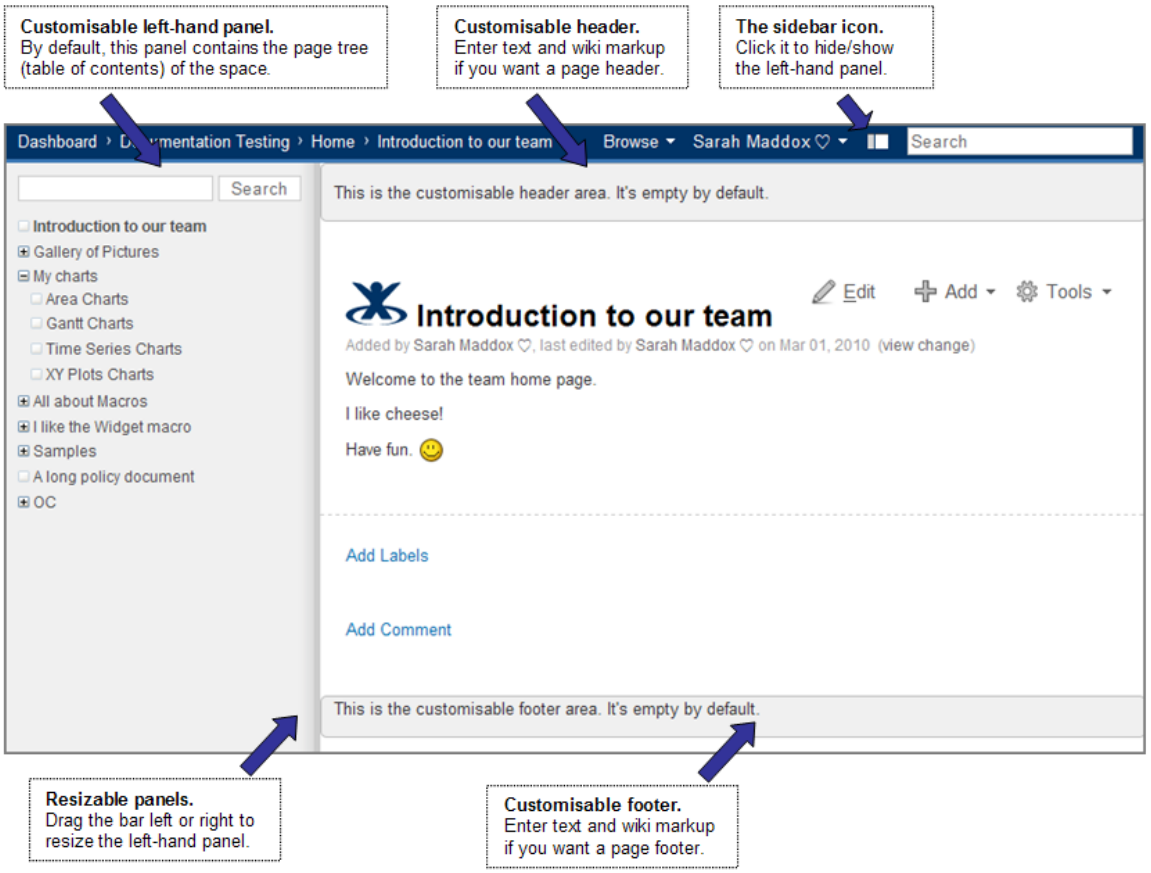


New Documentation Theme

Want an inbuilt table of contents for your wiki space? Lusting after a configurable header and footer? Hankering for sophisticated styling? You got it!

Features in a nutshell:

- Configurable left-hand panel. By default, the panel contains a search box and a table of contents (page tree).
- Resizable panels. People viewing the page can drag the thick bar between the left-hand panel and the content. They can also remove the panel altogether, by clicking the sidebar icon at top right.
- Customisable page header and footer.
- Text styles designed to enhance the content typically found in a documentation space.
- The Space Jump macro for linking from a page in one wiki space to a page with the same name in another space.
- Easy upgrade path from now on. Because the left-hand panel is part of the theme, it will be upgraded whenever Confluence is upgraded. There is no need to remove and then re-apply your customisations on each upgrade, as you would do if you added your own left-hand navigation bar.

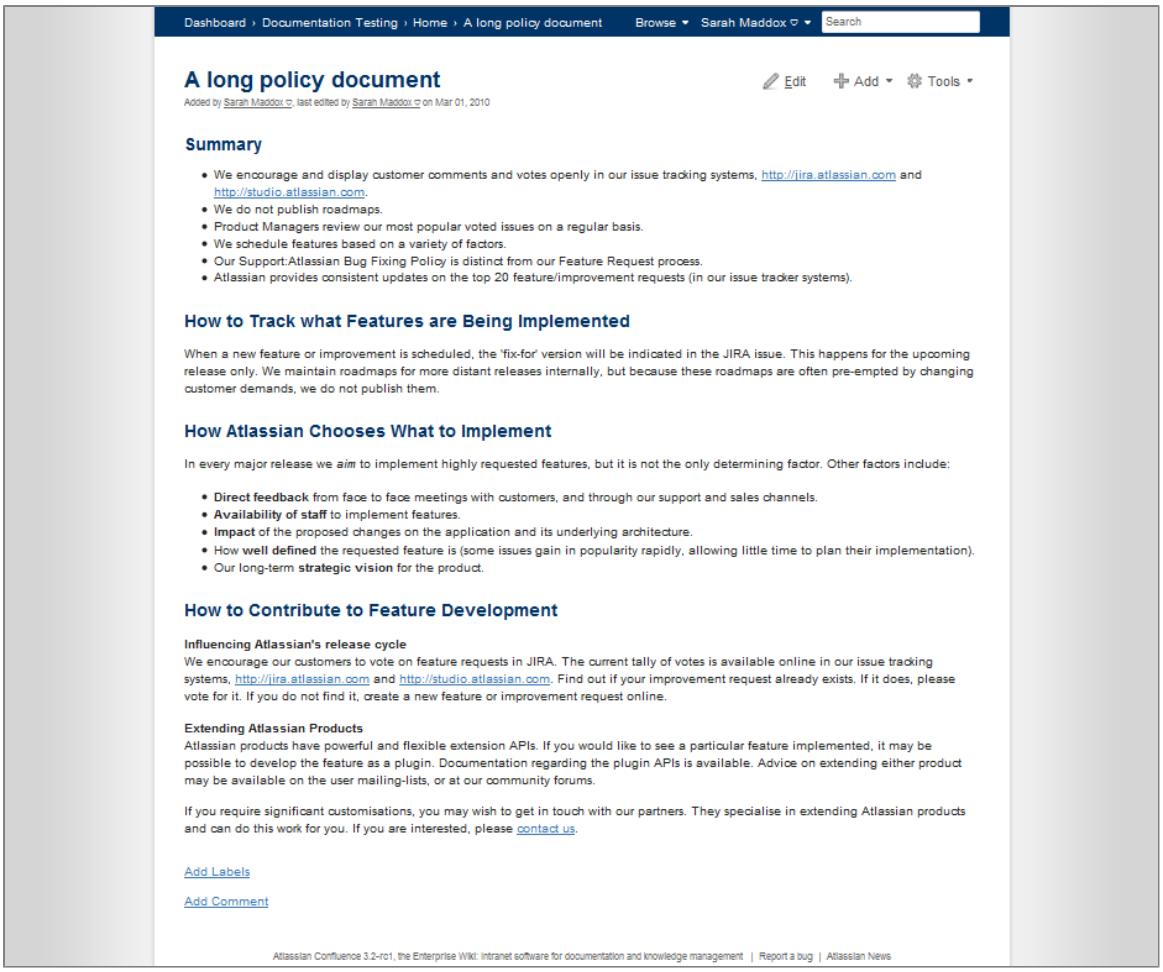


More...



New Easy Reader Theme

With today's huge monitors, it can become hard to read text that spans the width of the screen. Confluence 3.2 introduces the Easy Reader theme that uses only a portion of the screen, to make reading easier. Many websites are formatted this way. The Easy Reader theme is a fixed-width variation of the default Confluence theme. Its larger fonts, smooth gradient background and comfortable line length make it ideal for displaying and reading longer documents.



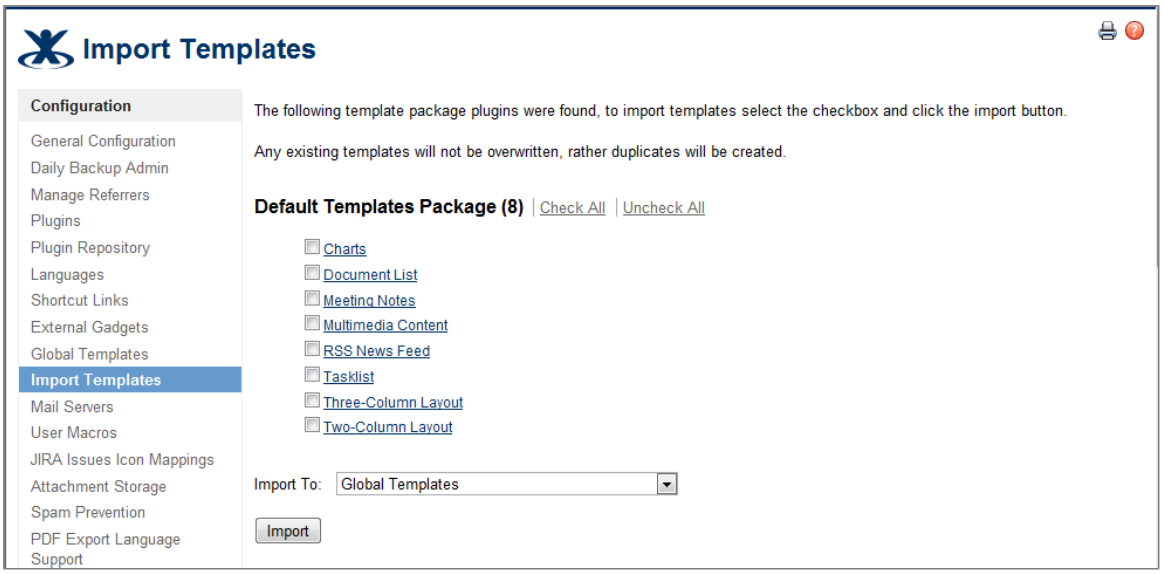
More...



Template Bundles

Confluence page templates make it easy for people to collaborate, yet maintain a consistent document format. We've extended the template functionality in this release so that you can import templates from template bundles via the Confluence Administration Console. Confluence administrators will be able to preview templates before importing them to a specific space or as global templates.

Template bundles are built as plugins, so developers should find it easy to whip up a new template bundle. Confluence 3.2 also ships with a default template bundle. Just import the templates from the default template bundle and your users will have access to a number of handy page templates.

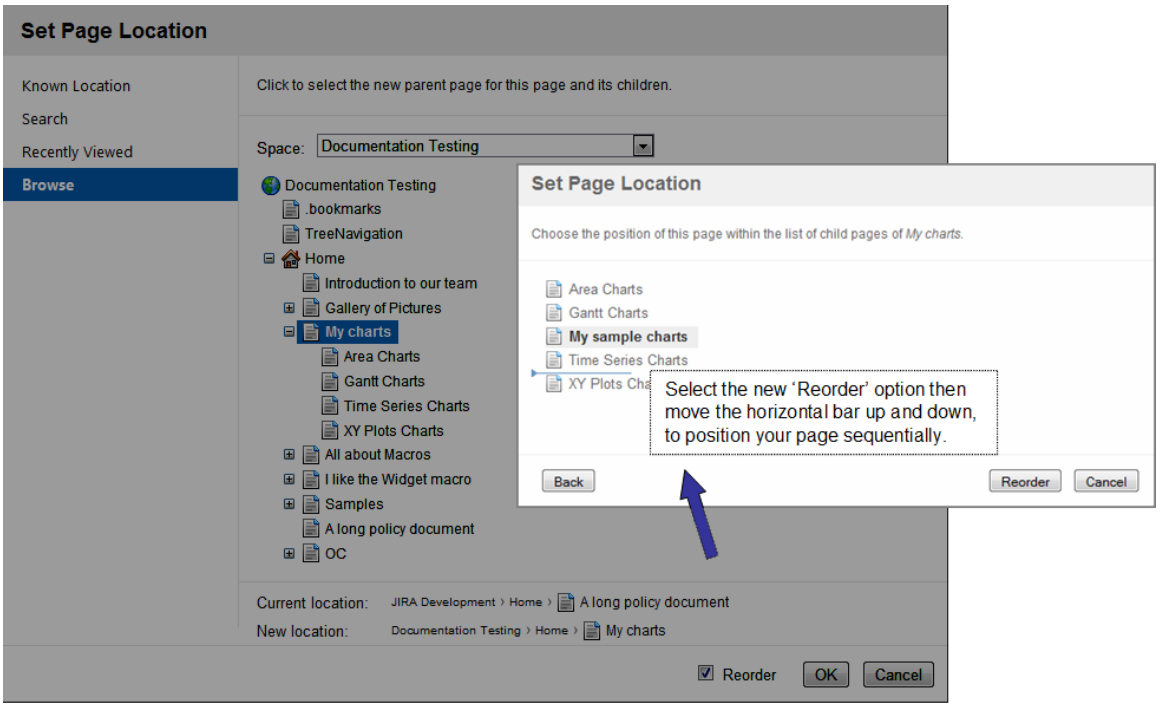


More...



Reordering while Moving a Page

You can now move a page and position it sequentially amongst its siblings at the same time. In earlier versions of Confluence this was only possible on the space's page tree view, accessed via the browse pages menu. The page tree is problematic in big spaces, so some people could not reorder pages at all. Now you can select the new 'Reorder' option when moving a page, and then drag a horizontal bar to put your page in the right spot amongst the other child pages.



More...

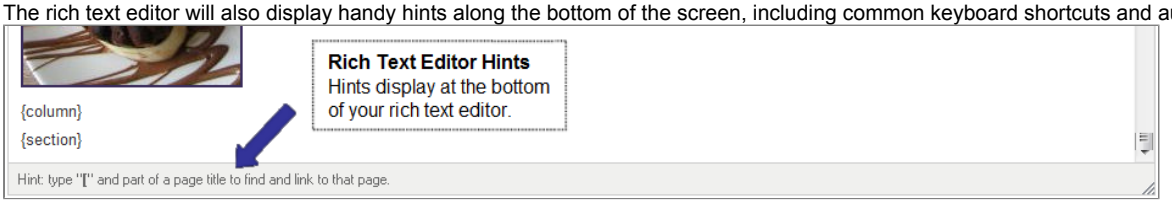


New Keyboard Shortcuts and Editor Hints

We're on a mission to improve your editing experience. We've added a number of new keyboard shortcuts for the rich text editor in this release. Format text into bulleted/numbered lists, manipulate tables, open the macro browser and more, with a few simple key presses. Check out the new keyboard shortcuts below:

Keystroke	Action (Rich Text Editor)
Ctrl+Shift+A	Opens the macro browser
Ctrl+Shift+B	Formats text as a bullet list
Ctrl+Shift+K	Autocomplete for links. Calls up a list of suggested pages or other locations to link to from your page. More...
Ctrl+Shift+M	Autocomplete for embedding images and files. Calls up a list of suggested images, documents and other files to embed in your page. More...
Ctrl+Shift+N	Formats text as a numbered list
Ctrl+Shift+S	Formats text with a strike through
Ctrl+Shift+C	Copies a table row
Ctrl+Shift+X	Cuts a table row
Ctrl+Shift+V	Pastes a table row
Ctrl+Shift+I	Inserts a table

The rich text editor will also display handy hints along the bottom of the screen, including common keyboard shortcuts and autocomplete tips.





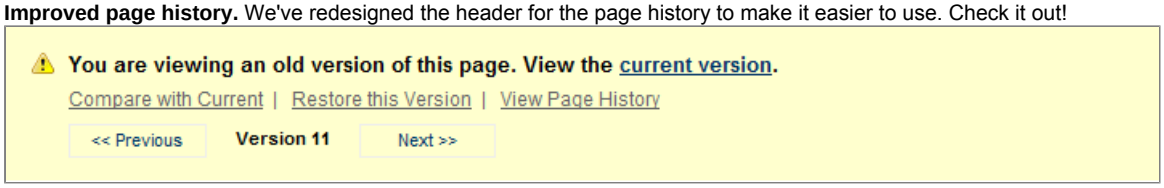
User Interface Enhancements

- **Image Previews in Search Results.** Confluence 3.2 helps you search more efficiently by displaying thumbnail image previews in your search results. You can also opt to display images attached to pages/blog posts in your search results. We've also made the rendering of search results in Confluence pluggable. If you are a developer, you may wish to write your own search result renderer to change how the search results are displayed.



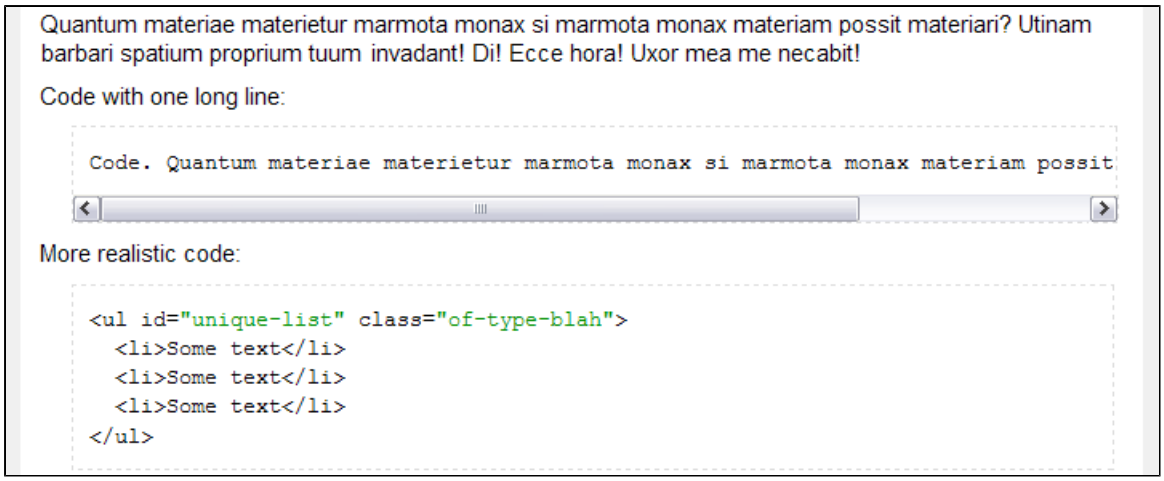
More...

- **Improved page history.** We've redesigned the header for the page history to make it easier to use. Check it out!



More...

- **Handling of oversized content.** Taking advantage of modern browsers, the new Easy Reader theme includes CSS to handle common types of oversized content. Specific sections of a page will now have a localised scroll bar instead of making the whole page scroll. You will see this in action in the code macro, for example.



- **Better RSS feeds.** We have improved and simplified the RSS feed builder. As before, you can choose to include either pages or blog posts or both in your feed. Now you can also choose to include comments from pages and/or comments from blog posts independently. Similarly, you can choose attachments from pages and/or blog posts. In addition, if you filter by label you can now track updates to labelled pages and comments on those pages.



More...

- **Fixed colour scheme issue.** We have removed the hard-coded colours that prevented colour schemes from applying correctly.
- **Fixed editor issue.** We've also fixed an issue that caused white space to overlap the right-hand side of the editor. This has now gone, making more space available to the editor.



And Even More Improvements

- **Hot installation of language packs.** You can now install language packs without restarting the server.
More...
- **Purging the trash is now incremental and logged.** Purging the deleted items from the Confluence trash has caused problems for people with a large number of items in the trash. The operation is slow and prevents other database updates while in progress. Confluence now breaks the operation into separate transactions and tracks progress in the logs.
- **More REST resources.** If you are a plugin developer or use our APIs in some other way, you'll find the REST API improvements in this release useful. The previous release of Confluence introduced a prototype REST API. With Confluence 3.2, we have added the following REST resources: attachments, recently viewed pages, search and user history. There are other small improvements too.
More...
- **Important Bug Fixes:**
  - **Custom PDF stylesheets are unique to spaces.** We've addressed an issue that caused changes made in one space's custom stylesheet to be reflected in the stylesheets for all other spaces. For more information, please see CONF-18843.
  - **Full index rebuilds now work correctly.** A Confluence bug caused a database exception to be thrown during full index rebuilds, which prevented further indexing jobs from being run correctly. This has now been fixed. For more information, please see CONF-18452.
  - **Confluence now upgrades correctly when using Crowd integration.** If you have integrated your Confluence instance with Atlassian's Crowd, you can now upgrade to this version of Confluence without the manual workaround required for Confluence 3.1. For more information, please see CONF-18150.

- **Circular references are no longer possible with the {excerpt-include} macro.** The {excerpt-include} macro could previously be used to create circular references that would crash Confluence. This has been fixed. For more information, please see CONF-15247.
- **Confluence upgrades do not disable enabled plugins.** Plugins that are disabled by default would previously be disabled during a Confluence upgrade, even if they had been enabled prior to the upgrade. This has been fixed. For more information, please see CONF-18417.

Known Issues in this Release

We have an enthusiastic and dedicated group of testers and customers who jump in there, try out the new Confluence release and report any problems so that we can fix them quickly.

We value this feedback, which means that we can tell you about some minor known issues in Confluence 3.2. Sometimes we find out about a problem with the latest version of Confluence after we have released the software. In such cases we publish information about the known issues in the Confluence Knowledge Base. Please check the important technical advisories on the front page of the Knowledge Base.

A big **thank you** to everyone who helps us ensure that Confluence keeps getting better and better.

The Confluence 3.2 Team

Development

Bugfixing and Maintenance
Brian Nguyen
Andrew Lynch
Xu-Heng Tjhin
Ryan Thomas

Engine Room and REST APIs
Anatoli Kazatchkov
Matthew Jensen
Daniel Kjellin

New Links Dialog and Editor Enhancements
Agnes Ro
Ryan Ackley

Autocomplete
Dmitry Baranovskiy
David Taylor

Themes
Ben Buchanan
Jens Schumacher

Build and Release Engineering
Don Willis

Small Improvements
Gerry Claps

Plugin Updates
David Chui

Special Projects (not shipping in this release)
Matt Ryall
Paul Curren
Charles Miller
David Loeng
Chris Kiehl
Jonathan Gilbert

Team Lead
Per Fragemann

Support

Kuala Lumpur
Sashidaran Jayaraman
Zed Yap
Arie Murdianto
Azwandi Mohd Aris
Ming Giet Chong

San Francisco
Jeremy Largman
Maleko Taylor

Tim Wong
Vincent Chang
Peter White
Marian Finch

Sydney
Partha Kamal
Gurleen Anand
Roy Hartono
Michael Seager
Ivan Benko

Others

Design
Stephen Russell

Performance Engineering
George Barnett

Product Management
Audra Eng
Jens Schumacher
Sherif Mansour

Product Marketing Management
Bill Arconati
Matthew Hodges

Quality Assurance
Mark Hrynczak
Penny Wyatt
Andrew Prentice

Technical Writing
Andrew Lui
Sarah Maddox

Confluence 3.2 Upgrade Notes

Below are some important notes on upgrading to **Confluence 3.2**. For details of the new features and improvements in this release, please read the Confluence 3.2 release notes.

On this page:

- Upgrade Notes
  - Clickr Theme and Left Navigation Theme are now Deprecated
  - End of Support for JBoss, Firefox 2 and Safari 2
  - Advance Notice of End of Support for Oracle Weblogic, IBM Websphere, Caucho Resin, DB2 8.2, Java Platform 5 (JDK/JRE 1.5) and Internet Explorer 6
  - Link Browser will Remove Tooltips
- Upgrade Procedure
- Troubleshooting the Upgrade
- Checking for Other Known Issues and Troubleshooting the Confluence Upgrade

Upgrade Notes

Clickr Theme and Left Navigation Theme are now Deprecated

Confluence 3.2 introduces two new themes, the Documentation theme and the Easy Reader theme. At the same time, we are announcing the deprecation of the following two themes:

Deprecated Theme	Description	Suggested Replacement
Clickr theme	This theme was inspired by the Flickr user interface, with Confluence content centred on the page. Please note that some features of Confluence 3.x are not fully supported by this theme.	Easy Reader theme
Left Navigation theme	This theme provides a navigation bar on the left hand side of the screen. Please note that some features of Confluence 3.x are not fully supported by this theme.	Documentation theme

We recommend that you move to a different theme as soon as possible. In the table above, we have suggested new themes that will give you similar and better-supported functionality.



Advance Notice — Clickr Theme and Left Navigation Theme will not be bundled with Confluence 3.3
Please note, the Clickr Theme and Left Navigation Theme will not be bundled with Confluence 3.3. We will not support these themes from Confluence 3.3 onwards.

End of Support for JBoss, Firefox 2 and Safari 2

As previously announced, we are no longer providing support for:

- JBoss (all versions) **from this release onwards**.
- Firefox 2 and Safari 2 **from this release onwards**.

Please see End of Support Announcements for Confluence for further details.

Advance Notice of End of Support for Oracle Weblogic, IBM Websphere, Caucho Resin, DB2 8.2, Java Platform 5 (JDK/JRE 1.5) and Internet Explorer 6

As previously announced, we are planning on ending support for:

- Oracle Weblogic, IBM Websphere and Caucho Resin **in Confluence 3.3**.
- DB2 version 8.2 **in Confluence 3.3**. DB2 9.7 will still be supported.
- Java Platform 5 (JDK/JRE 1.5) **in Confluence 3.3**.
- Internet Explorer 6 **in Confluence 3.3 or 13 July 2010, whichever is sooner**.

Please see End of Support Announcements for Confluence for further details.

Link Browser will Remove Tooltips

In Confluence 3.2 and later, the link browser no longer offers the option to include a tooltip for your link. If you have existing links with tooltips, the tooltip will disappear if you edit the link with the link browser. The tooltip will remain if you edit the link using wiki markup. See issue CONF-18668.

Upgrade Procedure



Upgrade a test environment first
As always, please test your upgrades in your test environment before rolling into production.

If you are already running a version of Confluence, please follow these instructions to upgrade to the latest version:

1. Before you upgrade, we strongly recommend that you back up your Confluence Home Directory and database. See the documentation on backing up your Confluence site. If you are using an external database, perform a database backup.
2. If your version of Confluence is earlier than 3.1.x, then please read the Upgrade Notes Overview and the Upgrade Notes for each version of Confluence listed on that page. (There are hyperlinks to each one.) Also:
    - If you are upgrading from 2.1 or earlier, please read the 2.2 release notes.
    - If you are upgrading from 2.2 or earlier, you will need to upgrade to Confluence 2.7.x first, confirm the upgrade was successful, then upgrade again from version 2.7.x to the latest. For more details, please refer to CONF-11767.
3. Download the latest version of Confluence.
4. Follow the instructions in the upgrade guide.

Troubleshooting the Upgrade

This section lists some specific issues that may occur during or as a result of the upgrade process, and guidelines on fixing the problem if it does happen to you.

- **Left-hand navigation bar in Documentation theme is empty.** If your existing Confluence installation already has the Documentation theme plugin installed, you may find that after upgrading to Confluence 3.2 the left-hand navigation bar is empty in the spaces that use the theme. The fix is to enable all modules of the Documentation theme plugin. See the knowledge base article.

Checking for Other Known Issues and Troubleshooting the Confluence Upgrade

- **Check for known issues**. Sometimes we find out about a problem with the latest version of Confluence after we have released the software. In such cases we publish information about the known issues in the Confluence Knowledge Base. Please check the important technical advisories on the front page of the Knowledge Base and follow the instructions to apply any necessary patches.
- **Did you encounter a problem during the Confluence upgrade?** Please refer to the guide to troubleshooting upgrades in the Confluence Knowledge Base.
- If you encounter a problem during the upgrade and cannot solve it, please create a support ticket and one of our support engineers will help you.

RELATED TOPICS

Confluence 3.2 Release Notes

Issues Resolved in Confluence 3.2

Below are the issues resolved in Confluence 3.2, ordered by number of votes. For the full details of the fixes, improvements and new features, please take a look at our issue tracker. Please also take a look at the Confluence 3.2 release notes for the new features in Confluence 3.2.

JIRA Issues (76 issues)

Type	Key	Summary	Status	Resolution	Votes
	CONF-11469	Improve usability of page tree for moving and ordering pages	Resolved	Fixed	65
	CONF-7745	RSS Feed should incude only comments of content included in feed	Resolved	Fixed	22
	CONF-17718	Downloading a .docx file in IE7/WinXP gives it a .zip extension (technically true but the average end-user wouldn't know that).	Resolved	Fixed	9
	CONF-7496	livesearch returns no results for complete words - only part words	Resolved	Fixed	6
	CONF-18774	Drag & Drop does not work on Firefox 3.6	Resolved	Fixed	5
	CONF-17971	'Move Page' dialog box cannot rearrange/reorder pages like the old page moving feature in 3.0 did.	Resolved	Fixed	4
	CONF-9312	RSS feeds that filter on labels should include comments of pages that match the label	Resolved	Fixed	4
	CONF-18437	Macro browser and Preview tab appears to be broken for all themes EXCEPT shipped default theme	Resolved	Fixed	3
	CONF-17341	NPE in ListItemConverter.convertNode(ListItemConverter.java:50) when saving page	Resolved	Fixed	2
	CONF-18751	PermittedSpacesScope may decrease the BooleanQuery maxClauseCount	Resolved	Fixed	1
	CONF-18381	NPE setting up Demo content on a custom db	Resolved	Fixed	1
	CONF-17311	Reword theme descriptions	Resolved	Fixed	1
	CONF-15180	Class Cast Exception being thrown when error encountered during mail queue flushing	Resolved	Fixed	1
	CONF-12134	Group templates together	Resolved	Fixed	1
	CONF-19069	Children pages don't get returned in the correct order when requested through rest	Resolved	Fixed	0
	CONF-18647	UnsupportedOperationException thrown when saving a page edit after logging out in another tab when a draft already exists.	Resolved	Fixed	0
	CONF-18588	Move Page Dialog has a memory leak in IE on the browse tab	Resolved	Fixed	0
	CONF-18586	Add more keyboard shortcuts for the RTE	Resolved	Fixed	0
	CONF-18554	ChildPositionComparator deprecated without alternative	Resolved	Fixed	0
	CONF-18544	Page with 286 attachments only shows first 20 on Attachment View	Resolved	Fixed	0
	CONF-18538	IE stylesheets defined in Themes don't get served with IE conditional comments	Resolved	Fixed	0
	CONF-18536	Add service for getting i18n from plugins	Resolved	Fixed	0

	CONF-18489	The Download All as zip is using the wrong mime type	Resolved	Fixed	0
	CONF-18485	Webdavloader.xpi not compatible with FireFox 3.6	Resolved	Fixed	0
	CONF-18452	Site indexing stops when a db exception occurs	Resolved	Fixed	0
	CONF-18446	Display editor hints in the status bar of the RTE	Resolved	Fixed	0
	CONF-18417	Enabled plugins that are disabled by default don't stay enabled after Confluence upgrades	Resolved	Fixed	0
	CONF-18409	Logs missing from installer edition	Resolved	Fixed	0
	CONF-18390	Rest Plugins have no Transactions	Resolved	Fixed	0
	CONF-18309	NPE when clicking 'PDF Stylesheet' in space admin when not global administrator	Resolved	Fixed	0
	CONF-18303	Rest Search Service	Resolved	Fixed	0
	CONF-18284	Pressing tab on autcomplete drop downs (such as quick nav) in Firefox causes the text to go white	Resolved	Fixed	0
	CONF-18271	cachecontents.jsp page is broken	Resolved	Fixed	0
	CONF-18240	New Fixed Width Theme	Resolved	Fixed	0
	CONF-18235	Left Nav Theme is not showing the correct text	Resolved	Fixed	0
	CONF-18216	{plugins-supported} macro is not working correctly after upgraded from 2.10 to 3.1	Resolved	Fixed	0
	CONF-18206	Need to remove hard-coded instances of blue	Resolved	Fixed	0
	CONF-18199	Dropdown doesn't work nicely with long usernames	Resolved	Fixed	0
	CONF-18195	Impossible to select a user using mouse from an autocomplete dropdown box in a search filter "Who"	Resolved	Fixed	0
	CONF-18194	viewxls macro in the macro browser does not work correctly when inserting attachments from a comment on a blog	Resolved	Fixed	0
	CONF-18150	Cannot upgrade to Confluence 3.1 when using Crowd integration	Resolved	Fixed	0
	CONF-18098	Fix license action requires authentication, can't be used if user migration fails	Resolved	Fixed	0
	CONF-18072	Improve caching of language preference when user is using site default language	Resolved	Fixed	0
	CONF-18036	Add plugin extension point for search result rendering	Resolved	Fixed	0
	CONF-18027	Missing internationalisation when number of user licenses is exceeded	Resolved	Fixed	0
	CONF-17898	Index Queue flush can miss entries	Resolved	Fixed	0
	CONF-17860	Extra white space on the right when adding a page @ 1024 x768	Resolved	Fixed	0
	CONF-17659	Excerpt generation misses first character	Resolved	Fixed	0

	CONF-17382	MacroMetadataParser ignores all attributes for macros without parameters	Resolved	Fixed	0
	CONF-17281	Improve Page History Navigation UI	Resolved	Fixed	0
	CONF-17203	Extraneous <p> tag in space.vmd	Resolved	Fixed	0
	CONF-17153	Space on the end of a database connection URL during setup can cause connection to fail.	Resolved	Fixed	0
	CONF-17077	Message about exceeding the number of licenced users contains a simple counting error	Resolved	Fixed	0
	CONF-17028	Email link should not appear for pages that the user is restricted from viewing	Resolved	Fixed	0
	CONF-17024	Improve instructions on the Confluence error page (500.jsp)	Resolved	Fixed	0
	CONF-16578	When Confluence license is invalid in confluence.cfg.xml, Confluence fails with "Upgrade failed" in the browser	Resolved	Fixed	0
	CONF-16389	Mail page plugin throws null pointer due to invalid LDAP membership	Resolved	Fixed	0
	CONF-16058	CreateRssFeedAction doesn't handle multiple RSS types in request nicely	Resolved	Fixed	0
	CONF-15919	Disabled Users Can Be Followed	Resolved	Fixed	0
	CONF-15915	Macro Browser - Chart Macro - "X-axis" and "Y-axis" labels are not immediately understood	Resolved	Fixed	0
	CONF-15756	Edit Attachment Storage page has escaped tags	Resolved	Fixed	0
	CONF-15528	Spelling errors and missing information in the online documentation (i.e. the notation guide)	Resolved	Fixed	0
	CONF-15317	Edit User Groups unreadable in Firefox	Resolved	Fixed	0
	CONF-15262	Add Keyboard shortcut for the macro browser	Resolved	Fixed	0
	CONF-15031	More useful filenames for PDF Exports	Resolved	Fixed	0
	CONF-14875	Link Browser	Resolved	Fixed	0
	CONF-14669	Tables don't initially look like tables when inserted in the RTE	Resolved	Fixed	0
	CONF-14636	On the search screen, add help text to "Who" field	Resolved	Fixed	0
	CONF-14565	In Browse Labels screen, Recent Labels and Popular Labels boxes on the right should be aligned with the current (future) styling.	Resolved	Fixed	0
	CONF-14558	Change "Children" to "Child Pages"	Resolved	Fixed	0
	CONF-13763	Moving attachments may silently fail to move the actual file	Resolved	Fixed	0
	CONF-13529	New "Configure whitelist" admin screen loses breadcrumbs	Resolved	Fixed	0
	CONF-13423	Insert Link popup needs to be cleaned up	Resolved	Fixed	0
	CONF-13392	Creating a page with RTE has whitespace where the notation guide is in wiki markup mode	Resolved	Fixed	0

	CONF-13196	list of users includes options to view more users per paqe even if the number of users less than 10	Resolved	Fixed	0
	CONF-2577	Auto-completion of likely words in edit dialogs	Resolved	Fixed	0

Above are the issues resolved in Confluence 3.2, ordered by number of votes. For the full details of the fixes, improvements and new features, please take a look at our issue tracker.

Confluence 3.1.2 Release Notes

3 March 2010

Confluence 3.1.2 is a recommended upgrade which fixes a number of bugs. Please see the 'Updates and Fixes in this Release' section below for details.

We identified a bug (see CONF-18437) that caused errors to be thrown whenever the macro browser was used in a theme other than the default theme. This bug also caused the preview mode to hang when the preview tab was clicked (i.e. the loading icon would spin forever without the preview mode loading). This has been fixed for all themes bundles with Confluence. If you are using a custom theme, you will need to implement a custom workaround, as described in the related JIRA issue.

Don't have Confluence 3.1 yet?

Take a look at the new features and other highlights in the Confluence 3.1 Release Notes.

Download Latest Version

Upgrading from a Previous Version of Confluence

Upgrading Confluence should be fairly straightforward. Please read the Confluence 3.1.2 Upgrade Notes. We *strongly recommend* that you back up your `confluence.home` directory and database before upgrading.

Updates and Fixes in this Release

There's a complete list of fixes below. Click a specific issue to see details of the fix. Click here to open a report on http://jira.atlassian.com for Resolved or Closed issues in Confluence 3.1.2.

JIRA Issues (21 issues)					
Type	**Key**	**Summary**	**Priority**	**Status**	**Resolution**
	CONF-18381	NPE setting up Demo content on a custom db		Resolved	Fixed
	CONF-18218	LDAP users with uppercase usernames cannot add favourites		Resolved	Fixed
	CONF-18437	Macro browser and Preview tab appears to be broken for all themes EXCEPT shipped default theme		Resolved	Fixed
	CONF-18207	Global colour scheme can not be applied to space-theme		Resolved	Fixed
	CONF-18202	German Umlauts break the QuickSearch		Resolved	Fixed
	CONF-17311	Reword theme descriptions		Resolved	Fixed
	CONF-16389	Mail page plugin throws null pointer due to invalid LDAP membership		Resolved	Fixed
	CONF-14820	Inherited page permissions are not wrapped correctly in UI		Resolved	Fixed
	CONF-13627	Modz Detector for Confluence - track your file system changes automatically		Resolved	Fixed
	CONF-13529	New "Configure whitelist" admin screen loses breadcrumbs		Resolved	Fixed
	CONF-7473	Confluence uses the server's locale for dates and not the default language		Resolved	Fixed

	CONF-18271	cachecontents.jsp page is broken		Resolved	Fixed
	CONF-18235	Left Nav Theme is not showing the correct text		Resolved	Fixed
	CONF-18194	viewxls macro in the macro browser does not work correctly when inserting attachments from a comment on a blog		Resolved	Fixed
	CONF-18144	Macro browser macro alias		Resolved	Fixed
	CONF-17533	Uploading new version of attachment creates duplicate options in the Insert Image Dialog		Resolved	Fixed
	CONF-17213	Warning on new license version still appears after license upgrade to v2		Closed	Cannot Reproduce
	CONF-16585	Confluence allows you to change the title of a page to whitespace, which prevents editing and access unless you know the pageid.		Resolved	Fixed
	CONF-15919	Disabled Users Can Be Followed		Resolved	Fixed
	CONF-17028	Email link should not appear for pages that the user is restricted from viewing		Resolved	Fixed
	CONF-13832	Invalid content label entries are populated in the team label entry field on the Space Admin -> Edit Space Labels screen when validation fails (& vice versa)		Resolved	Fixed

Confluence 3.1.2 Upgrade Notes

Below are some important notes on upgrading to **Confluence 3.1.2**. Confluence 3.1.2 is a recommended upgrade which fixes some security flaws as well as other bugs. For more details, please read the Confluence 3.1.2 Release Notes.

On this page:

- Upgrade Notes
- Upgrade Procedure

Upgrade Notes

There are no upgrade tasks specific to Confluence 3.1.2. Please refer to the 'Upgrade Procedure' section below.

Upgrade Procedure

If you are already running a version of Confluence, please follow these instructions to upgrade to the latest version:

1. Before you upgrade, we strongly recommend that you **back up your `confluence.home` directory and database**.
2. If your version of Confluence is earlier than 3.1.2, read the release notes and upgrade guides for **all releases** between your version and the latest version. In particular:
    - Please read the Confluence 3.1.2 Release Notes.
    - If you are upgrading from 2.1 or earlier, please also read the 2.2 release notes.
3. Download the latest version of Confluence.
4. Follow the instructions in the Upgrade Guide.

Confluence 3.1.1 Release Notes

27 January 2010

Confluence 3.1.1 is a recommended upgrade which fixes a number of issues.

Performance Enhancements and Fixes

This release fixes two bugs that affected configuration of Confluence's internal caches. The first bug caused an error to be thrown when a Confluence internal cache was either increased to over 5000 elements (from under 5000) or resized when it was already over 5000 elements, via the Administration Console. The second bug caused changes made to Confluence's internal caches via the Administration Console to be lost after a Confluence restart. This meant that caches could only be resized by manually editing the cache configuration files (XML).

We identified an issue where Hibernate was not cleaning up its session correctly when closed, causing memory leaks. We have changed the way that Hibernate SessionImpl objects are stored which should fix this problem and reduce memory consumption.

A system property (`atlassian.user.experimentalMapping`) was introduced in Confluence 2.10 to fix performance problems noted when adding a local user to a local user group. We have fixed an bug with this system property that prevented users from being removed from user groups via the Administration Console, when the system property is set.

Other Fixes

We have found a bug that periodically caused an error to display when trying to view the 'Confluence Gadgets' window. This has now been fixed.

There's a complete list of fixes below. Click a specific issue to see details of the fix.

Don't have Confluence 3.1 yet?

Take a look at the new features and other highlights in the Confluence 3.1 Release Notes.

Download Latest Version

Upgrading from a Previous Version of Confluence

Upgrading Confluence should be fairly straightforward. Please read the Confluence 3.1.1 Upgrade Notes. We *strongly recommend* that you back up your `confluence.home` directory and database before upgrading.

Updates and Fixes in this Release

JIRA Issues (46 issues)

Type	Key	Summary	Priority	Status	Resolution
	CONF-18137	Updating Caches via Administration > Cache Statistics, will update the ehcache.xml file, but upon restart all the caches are set back to default.		Resolved	Fixed
	CONF-14989	Possible net.sf.hibernate.impl.SessionImpl Memory Leaks		Resolved	Fixed
	CONF-18095	macros that do not provide parameter information are rendered poorlt		Resolved	Not a bug
	CONF-18030	NullPointerException thrown while browsing/searching for users (LDAP configuration only)		Resolved	Fixed
	CONF-17898	Index Queue flush can miss entries		Resolved	Fixed
	CONF-17884	QuickNav gadget login button is missing		Resolved	Fixed
	CONF-17246	Resizing caches does not correctly update keyArray in EhCache's MemoryStore:keyArray		Resolved	Fixed
	CONF-17039	Confluence upgrade failed due to database case sensitivity		Resolved	Fixed
	CONF-12319	Adding users to a large group is slow with default (Hibernate) user management		Resolved	Fixed
	CONF-18164	The log4j properties file included in builds is incorrect		Resolved	Fixed
	CONF-18093	DefaultUserProfileService tries to modify an unmodifiable collection (PATCH)		Resolved	Fixed
	CONF-18024	Auto support requests from 3.1 are setting "affects version" to 2.5		Resolved	Fixed
	CONF-17973	Link dialog history can list blog posts as pages instead		Resolved	Fixed
	CONF-17946	Pdf export - wrong content type.		Resolved	Fixed
	CONF-17932	Pull down menus do not appear in IE6 when the username is longer than 26 characters		Resolved	Fixed
	CONF-17864	Drag and drop plugin on edit page causes javascript memory leak in IE		Resolved	Fixed
	CONF-17847	When inserting a link to an attachment in IE without specifying alias it inserts \^attachment as an alias		Resolved	Fixed
	CONF-16643	Thumbnails not working in 3.0 for the same JPEG file as still does in 2.7.3		Resolved	Fixed

	CONF-16176	NPE given by tasklist user picker when not logged in		Resolved	Fixed
	CONF-14005	First install page links to email and forums when it should link to support.		Resolved	Fixed
	CONF-11838	Extra <html> in Custom HTML field breaks Custom HTML edit page		Resolved	Fixed
	CONF-7735	Cannot edit or remove attachments with en dash or em dash in the filename		Resolved	Fixed
	CONF-18103	Userpicker loses search term when changing pages in advance search		Resolved	Fixed
	CONF-18100	Change 500page.jsp so it informs users where to zip log files and submit to us		Resolved	Fixed
	CONF-18043	Fix misleading text on referrers admin screen		Resolved	Fixed
	CONF-18038	Long Breadcrumbs are displayed on the next line in IE8 Standards mode		Resolved	Fixed
	CONF-18007	Debug logging contains the message 'AND YERRRRRR MAWWWWWWWWWWWWWWW' which isn't meaningful enough to be helpful		Resolved	Fixed
	CONF-17998	Unnecessary 20px Macro Browser icons in images subfolder		Resolved	Fixed
	CONF-17986	Browse: Attachments: Current sorted column is not indicated for space attachments, which is inconsistent with page attachments		Resolved	Fixed
	CONF-17984	Browse: Attachments: Sort by Creation Date is incorrect		Resolved	Fixed
	CONF-17960	Content added after setup but before the initial reindex finishes will not be in the index		Resolved	Fixed
	CONF-17896	Confluence Administrators are unable to access the support utility		Resolved	Fixed
	CONF-17877	RTE has whitespace in drafts mode where the wiki markup help used to be		Resolved	Duplicate
	CONF-17784	Include the support zip utility in the raise support request form		Resolved	Fixed
	CONF-17631	No spinner on page tree indicating progress when opening the tree in the move page dialog		Resolved	Fixed
	CONF-17273	Markup editor doesn't work nicely with IE8's "Find In This Page" feature		Resolved	Fixed
	CONF-17114	Mouse Select in the User Follow Box doesn't always work		Resolved	Fixed
	CONF-16974	Advanced macros do not accept a space as a separator for the author parameter		Resolved	Fixed
	CONF-16852	"Now following Anonymous" instead of actual user		Resolved	Fixed
	CONF-16815	Incorrect docs for {float} macro		Resolved	Not a bug
	CONF-16512	Impossible to select a user using mouse from an autocomplete dropdown box in a search filter "Who"		Resolved	Fixed
	CONF-16088	Content by label macro is missing both the 'author' and 'operator' parameters from the macro browser.		Resolved	Fixed
	CONF-16060	Space details macro doesn't display its description correctly - instead, showing some internal code variable.		Resolved	Fixed
	CONF-14494	Borders are not applied to image thumbnails		Resolved	Fixed

	CONF-10671	export throws null pointer exception when servlet context contains capital letter		Resolved	Fixed
	CONF-10653	Confluence is not able to move pages from one space to another that contain asterisks		Resolved	Fixed

If you like, you can open a report on jira.atlassian.com showing all resolved and closed issues in this release.

Confluence 3.1.1 Upgrade Notes

Below are some important notes on upgrading to **Confluence 3.1.1**. Confluence 3.1.1 is a recommended upgrade which fixes a number of issues. For more details, please read the Confluence 3.1.1 Release Notes.

On this page:

- Upgrade Procedure

Upgrade Procedure

If you are already running a version of Confluence, please follow these instructions to upgrade to the latest version:

1. Before you upgrade, we strongly recommend that you **back up your `confluence.home` directory and database**.
2. If your version of Confluence is earlier than 3.1.0, read the release notes and upgrade guides for **all releases** between your version and the latest version. In particular:
    - Please read the Confluence 3.1.1 Release Notes.
    - If you are upgrading from 2.1 or earlier, please pay careful attention to the 2.2 release notes.
3. Download the latest version of Confluence.
4. Follow the instructions in the Upgrade Guide.

Confluence 3.1 Release Notes

Confluence 3.1 fixes some security flaws. Please refer to the security advisory for details of the security vulnerabilities, risk assessment and mitigation strategies.

8 December 2009

With great pleasure, Atlassian presents Confluence 3.1.

Confluence 3.1 is a major release which presents a number of new features and enhancements. With Confluence 3.1, we introduce **OpenSocial Gadgets**, which allow you to add functionality from other web applications or websites, such as JIRA 4.0+, iGoogle or Gmail, directly into a Confluence page or blog post. Confluence also provides its own Gadgets that can be embedded into other web applications.

The new **Drag-and-Drop** feature is a major step forward in Confluence file attachment. With Drag-and-Drop, you can attach one or more files to a page or blog post simply by dragging them from your desktop onto the Confluence page. You can also embed images, Office documents and PDF files into a page or blog post while in edit mode simply by dragging and dropping them into the rich text editor.

Office 2007 documents are now fully supported in Confluence so features like document import, edit in Word and embedding documents in a page now work with Office 2007 files. With the new **Move Page** feature, you can more quickly and reliably move the page you are editing or viewing to a new parent page, even to one within a new space. An enhanced **Image Browser** lets you preview images before embedding them into a page. You can also attach image files to a page by dragging and dropping them into the Image Browser window itself. A new **Page Restrictions** dialog lets you view and apply page restrictions without editing the page. **Draft Comparisons** allow you to view changes made to a page or blog post before they are actually saved.

Other editor enhancements include **Exit Notifications** to prevent the loss of unsaved content in a page, blog post or comment, a new **Insert Menu** for the rich text editor, fields that auto-complete in the macro browser and improved rich text editor speed.

Confluence 3.1 now fully supports Internet Explorer 8, Safari 4 and Firefox 3.5. Other smaller improvements include the ability to add pages or blog posts directly from the Confluence dashboard, a facility to link to Confluence pages more easily and the ability to expand activity streams at the click of a button.

Finally, we have also introduced a simplified installation process for evaluators to help get Confluence up and running quickly. Simplified Confluence evaluation installers have been created for MacOS X and Windows operating systems.

Highlights of this Release:

- Introducing Gadgets
- Drag-and-Drop
- Office 2007 Support
- New 'Move Page' Feature
- Enhanced Image Browser
- Draft Comparisons
- Page Restrictions Dialog Box
- Other Editor Enhancements
    - Edit Mode Exit Notification
    - New Rich Text Editor Insert Menu

    - Macro Browser Smart Fields
    - Rich Text Editor Speed
- New Web Browser Versions Supported
- Other Improvements
    - Add Pages or Blog Posts from the Dashboard
    - New 'Link to this page' feature
    - Get 'More' from your Activity Streams
    - User Interface Performance Improvements
    - Other Small Enhancements and Improvements to Confluence

Responding to your Feedback:

300+ votes satisfied

DOWNLOAD latest version

- Thank you for all your issues and votes. Keep logging, to help us keep improving!
- Below is a list of the highlights in this release.
- Attached is the full list of issues resolved in this release.

Upgrading from a previous version of Confluence
• Upgrading Confluence should be fairly straightforward. ***We strongly recommend that you back up your Confluence Home directory and your database before upgrading.*** • Please refer to the Confluence 3.1 Upgrade Notes for further essential information about **plugins** and other factors affecting your upgrade.

Highlights of Confluence 3.1

1

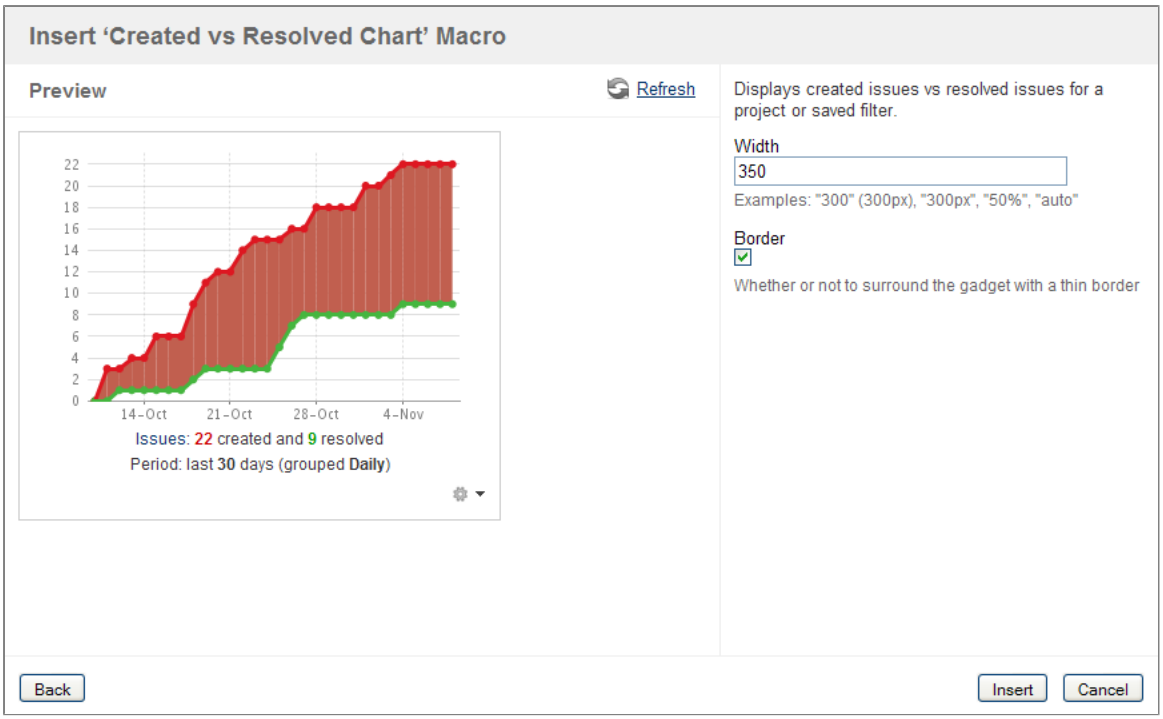
Introducing Gadgets

Gadgets are small objects that offer dynamic content and functionality which may be served by any OpenSocial-compliant web application, such as JIRA 4.0+, Confluence or non-Atlassian applications such as iGoogle and Gmail. Confluence can interact with any gadgets that support the OpenSocial specification.

- Confluence supports the use of gadgets in pages and blog posts, which are accessible through the macro browser.
- Confluence can also serve its own gadgets, for use in any other OpenSocial-compliant web application. Gadgets bundled with Confluence include:
    - **Activity Stream** — This gadget shows a list of recent activities that have occurred on the Confluence server, such as the addition of new pages, blog posts or comments, content edits, status updates and so on.
    - **Quick Navigation Aid** — This gadget provides heading and content search capabilities on a Confluence server.

      Your Confluence installation can also serve these gadgets in any of its own pages or blog posts.
      For more information on using these gadgets, refer to Confluence Gadgets.

Inserting a JIRA Gadget onto a Confluence Page





Drag-and-Drop

The new 'drag-and-drop' feature allows you to drag one or more file(s) which are accessible from your computer and drop them directly into a Confluence page or blog post.

- Files can be attached to a page or blog post by dropping them directly onto the page view or the 'Attachments' list associated with the page.
- Image files can be attached to a page or blog post by dragging them from your computer directly onto the Image Browser.
- Image and Office files can be added directly into your Confluence page or blog post content by dropping them into the rich text editor's editor window.

For more information about this feature and on how to set it up, refer to the Drag-and-Drop documentation.

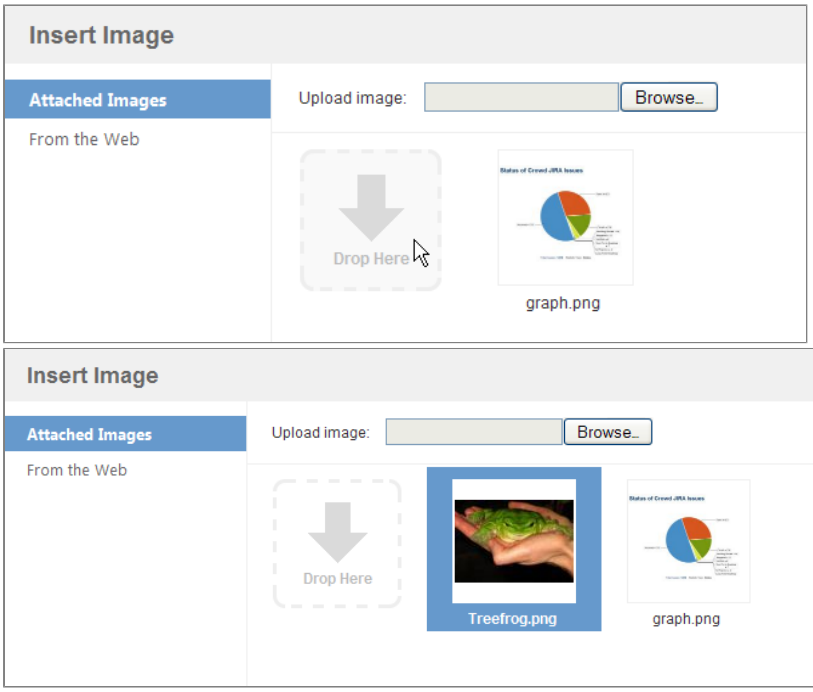
Screenshot: 'Drag-and-Drop' Images or Other Files Directly onto a Page



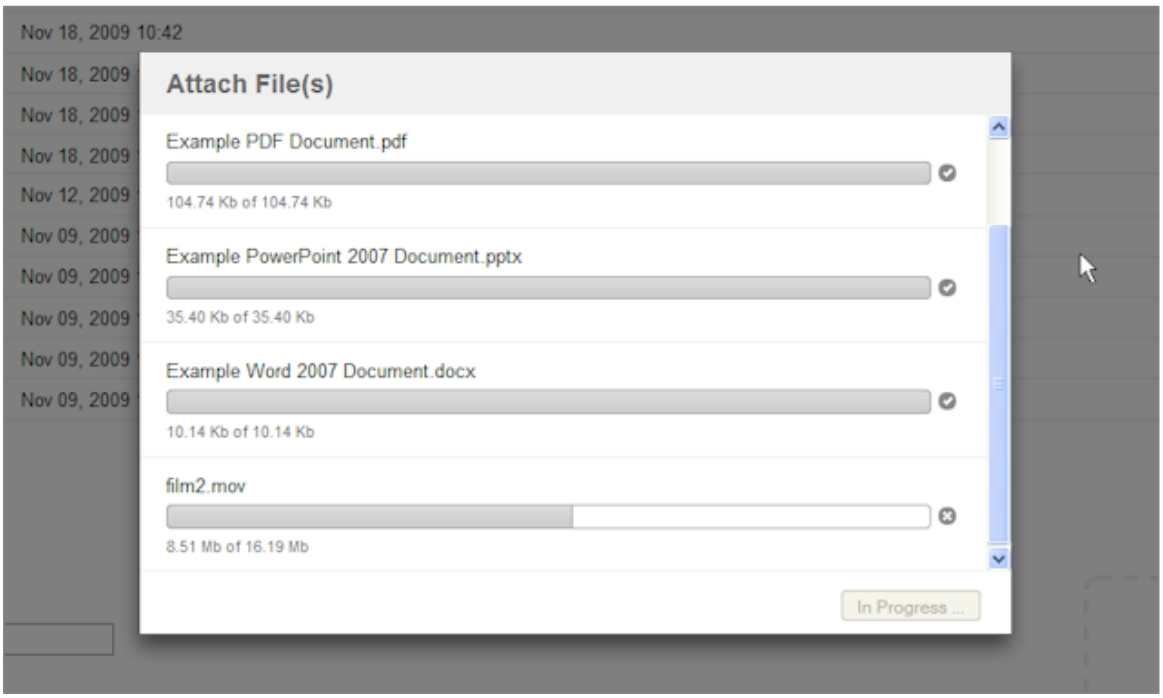
Video: Using Drag-and-Drop

Download Video

Screenshots: Attaching an Image to the Image Browser



Screenshot: Attaching Multiple Files to an 'Attachments' list





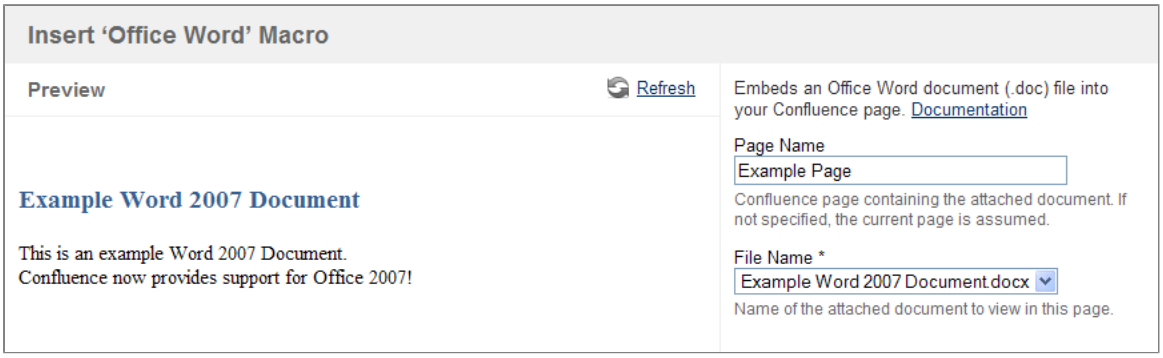
Office 2007 Support

Confluence now provides full support for the new Office 2007 file formats, allowing you to view and edit content from Microsoft Word 2007 (`.docx` and `.dotx`), PowerPoint 2007 (`.pptx` and `.potx`) and Excel 2007 (`.xlsx`) files.

- Along with existing Microsoft Office versions, Confluence now fully indexes Microsoft Office 2007 files and their content can be searched by Confluence.
- Using Confluence's Office connector macros, you can insert Word, PowerPoint or Excel 2007 files directly into your Confluence page or blog post.

- Office files can be edited directly from any page or blog post or their list of attachments.
  If you use the Firefox browser to work with Confluence, don't forget to reconfigure the Firefox add-on (WebDAV Launcher options) to handle the new Office 2007 file extensions. Otherwise, you will not be able to edit these new Office 2007 file formats from Confluence.

Screenshot: Embedding an Office 2007 Document



This feature uses technology licensed from Aspose.

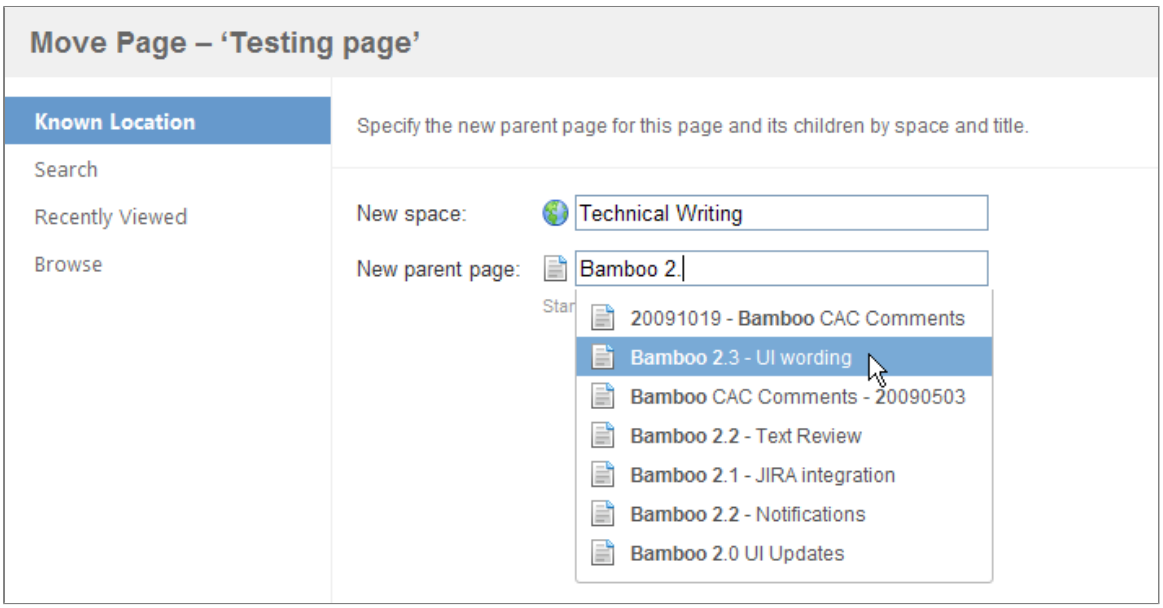


New 'Move Page' Feature

Confluence introduces a new page moving feature, that easily allows you to move the page you are currently viewing, adding or editing, to another page in the same or another space within your Confluence site. This feature is available through a new 'Move Page' dialog box, which provides the following flexible methods for moving pages:

- **Known Location** – Allows you to type the name of a space and within that space, the 'parent' page under which to move your page.
- **Search** – Allows you to search for a 'parent' page (within a selected space or set of spaces) under which to move your page.
- **Recently Viewed** – Allows you to select one of your recently viewed pages to be the 'parent' of your page to be moved.
- **Browse** – Allows you to select a space and page that will be the 'parent' of your page to be moved. Pages are browsed via a tree view.
  For more information, refer to Moving a Page.

Screenshot: Moving a Page to a Known Parent Page



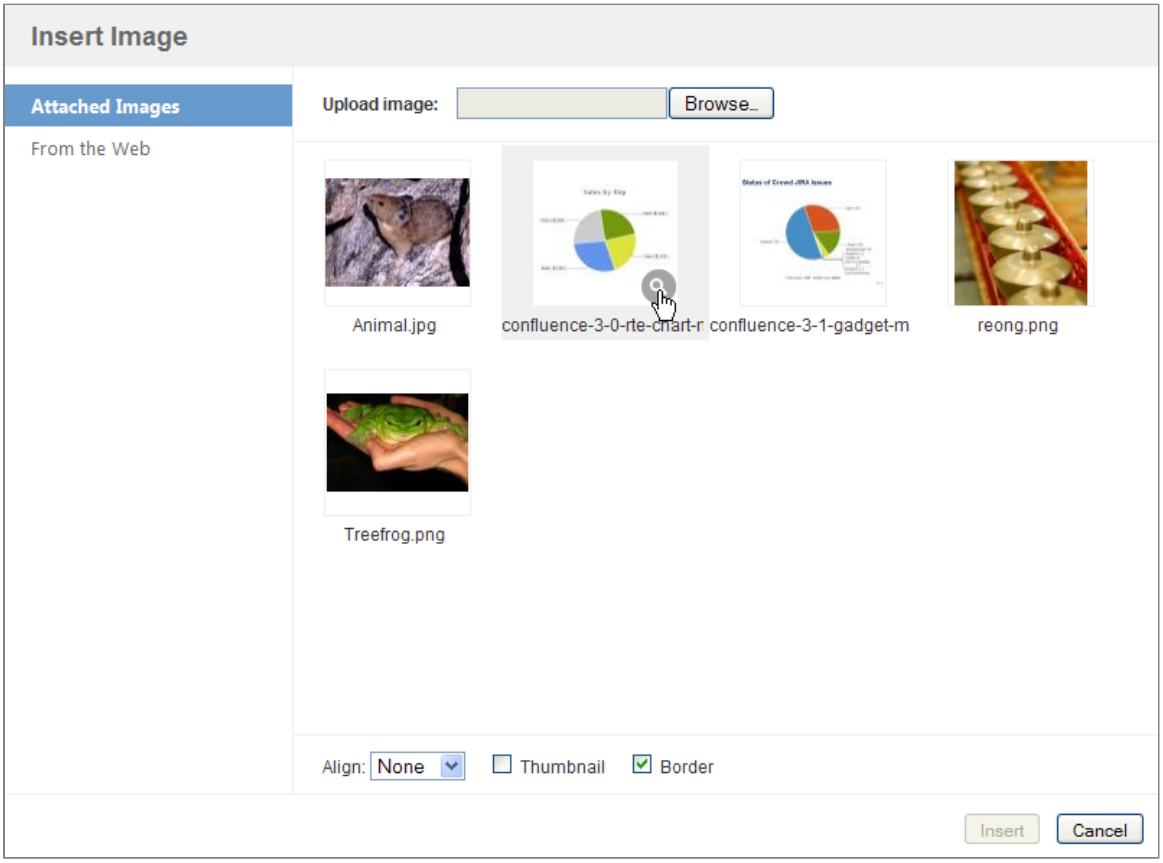


Enhanced Image Browser

A new 'Image Browser' has been introduced to replace the old 'Insert Image' window. The image browser provides a less-cluttered and enhanced interface that allows you to:

- Preview an image in detail before inserting it into a page. This is done by hovering over any image in the browser and clicking the 'magnifying glass' icon in the lower-right corner.
- Preview an image elsewhere on the web via its URL before inserting it into a page.

Screenshot: Previewing an Image in the Image Browser



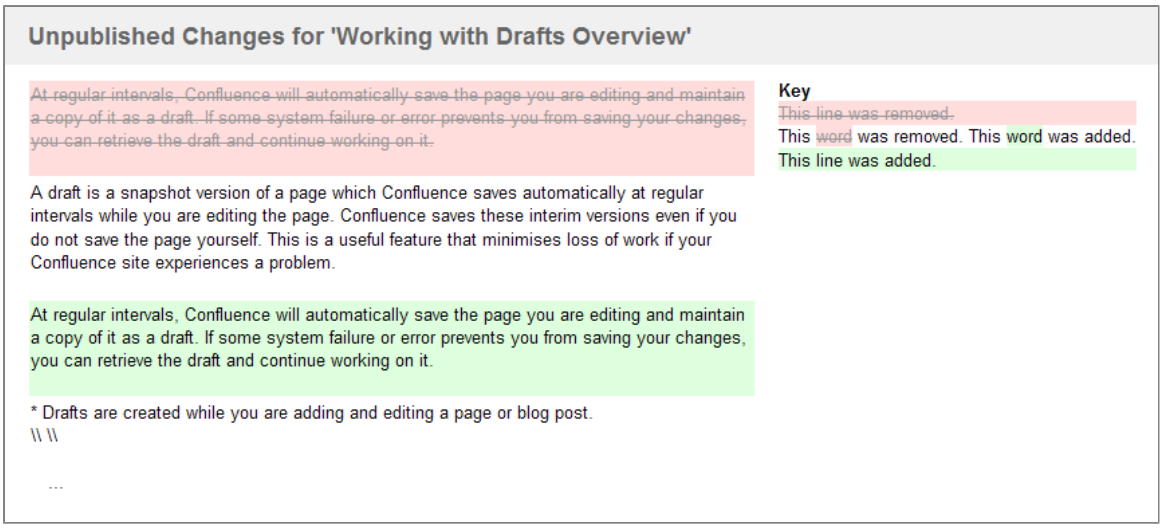


Draft Comparisons

Confluence's drafts feature has been enhanced, such that you can now view unsaved changes in your drafts as a 'diff' before you decide to resume editing them. This nifty 'draft comparison' feature comes in handy, particularly when other people have made subsequent changes to a page or blog post in your drafts list and you need to merge changes or resolve a conflict.

Screenshots: Accessing and Viewing Changed Content with Draft Comparisons





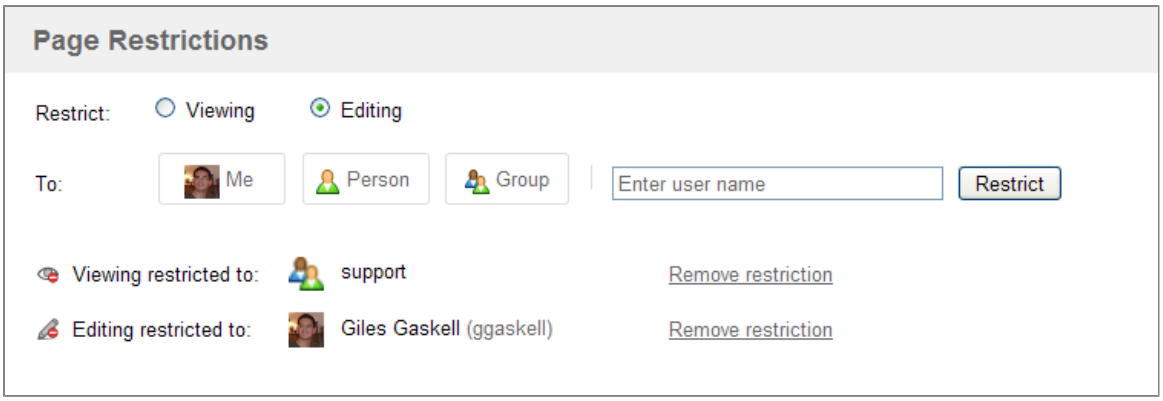


Page Restrictions Dialog Box

Confluence's page restrictions feature has been incorporated into a convenient and accessible dialog box, which is now easier to use than before.

- The page restrictions dialog box can be accessed from the padlock icon or the '**Tools**' -> '**Restrictions**' menu item whilst viewing any Confluence page. From this dialog box, you can see all viewing and editing restrictions associated with the current page. You no longer need to view the page's associated 'Info' page to see the page's restrictions.
- You no longer have to edit a page to modify its page restrictions. You can edit all page restrictions from this easily accessible dialog box.
  The page restrictions dialog box is still accessible when a page is in edit mode.
- In addition to user and group names, the name field also accepts a user's full name. Full names are 'auto-completed' to help you find the relevant person more rapidly.

Screenshot: The Page Restrictions Dialog Box

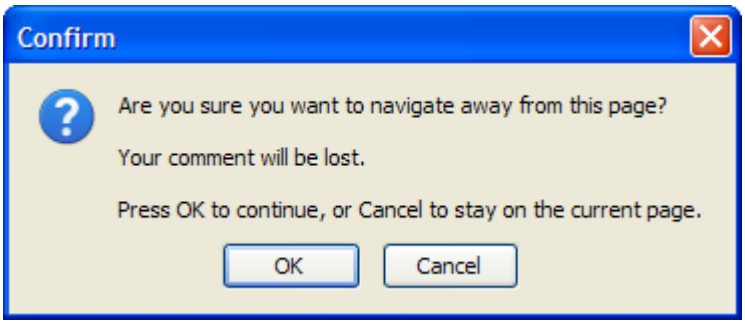




Other Editor Enhancements

Edit Mode Exit Notification

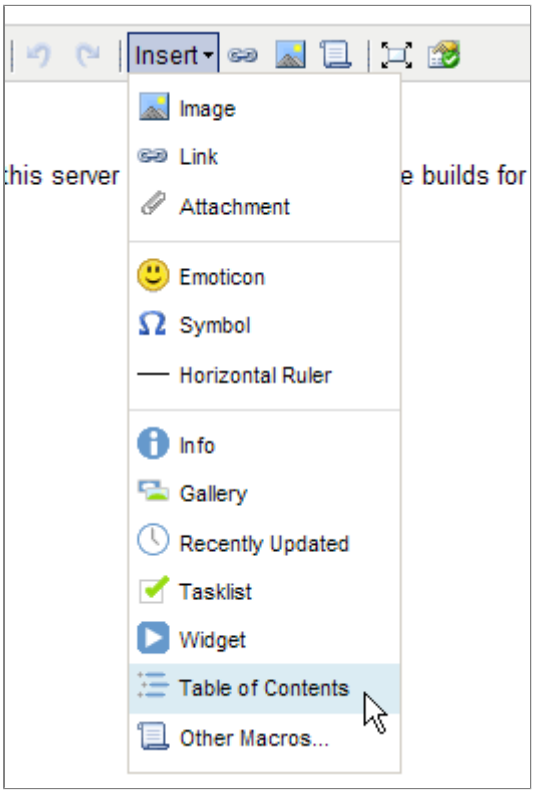
Whenever you add or edit a page, comment or blog post and then click onto another Confluence feature that navigates away from your unsaved content, a message box appears, warning that your content will be saved as a draft (if it is a page) or lost (if a comment). This allows you to cancel out of this action if it was accidental.



New Rich Text Editor Insert Menu

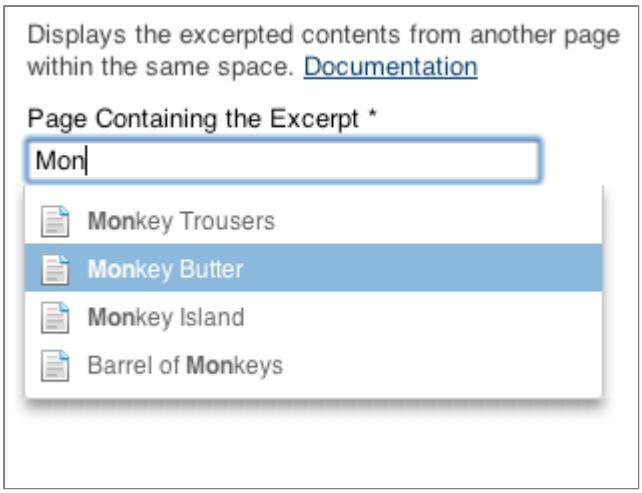
Confluence's rich text editor now combines a number of its commonly used editing features into a new convenient '**Insert**' menu.

- The Horizontal line, Insert Symbol and Insert Emoticon Toolbar icons have been moved into the new insert menu.
- The functionality to insert images, links or attachments into a page can now also be accessed from this menu.
- The macro browser, as well as a number of commonly-used macros are conveniently accessible from this menu too.



Macro Browser Smart Fields

When using the Macro Browser, an 'auto-complete' feature is now provided on any parameters that require the entry of a single item, such as a page title, username or space key. This greatly facilitates the customisation of macros and minimises the need to know the exact item names in advance.



Rich Text Editor Speed

Thanks to many individual technical improvements, the rich text editor opens up a lot faster than in previous Confluence releases. In a local network environment, the rich text editor is accessible almost instantly. When accessing a Confluence server on a different continent, the rich text editor still opens up rapidly. In our Sydney office for instance, accessing the rich text editor from our Confluence server in the US takes less than 3 seconds.



New Web Browser Versions Supported

Confluence 3.1 now fully supports the following recent web browser versions:

- Internet Explorer 8
- Safari 4
- Firefox 3.5

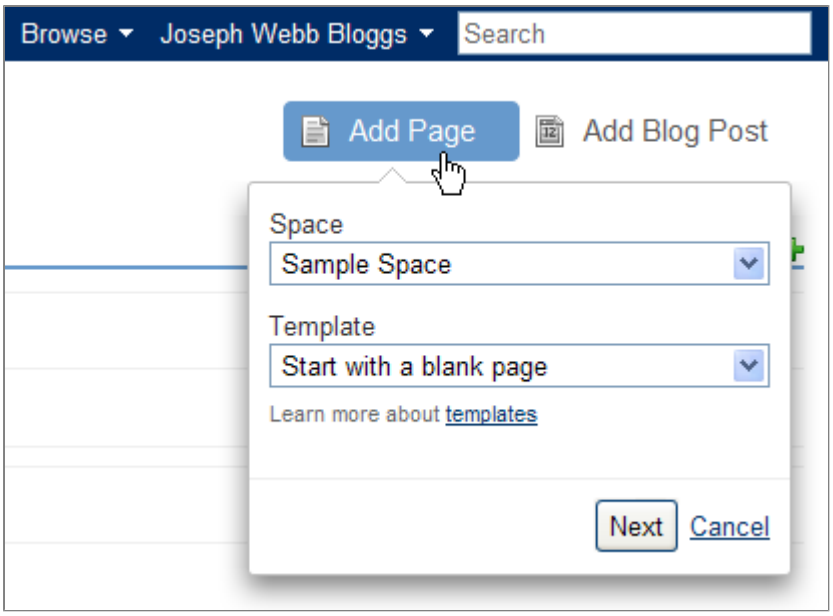
Confluence now provides full usability with these recent browser versions as many bugs associated with these browsers have been fixed.



Other Improvements

Add Pages or Blog Posts from the Dashboard

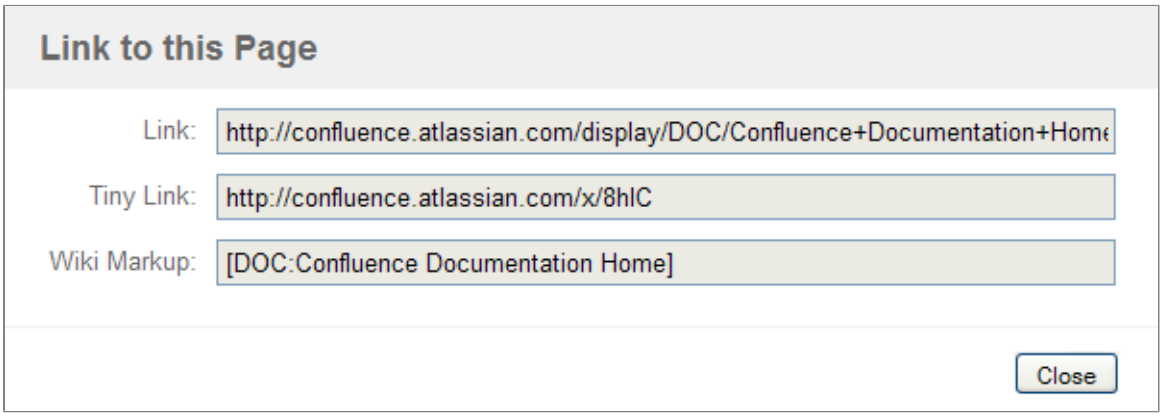
You can now add pages or blog posts directly from the Dashboard without having to browse to a specific space first. To do this, click on either the '**Add Page**' or '**Add Blog Post**' buttons to open the pop-up balloon, which allows you to choose the space in which to add the new page or blog post and in the case of pages, a template on which to base the page content.



New 'Link to this page' feature

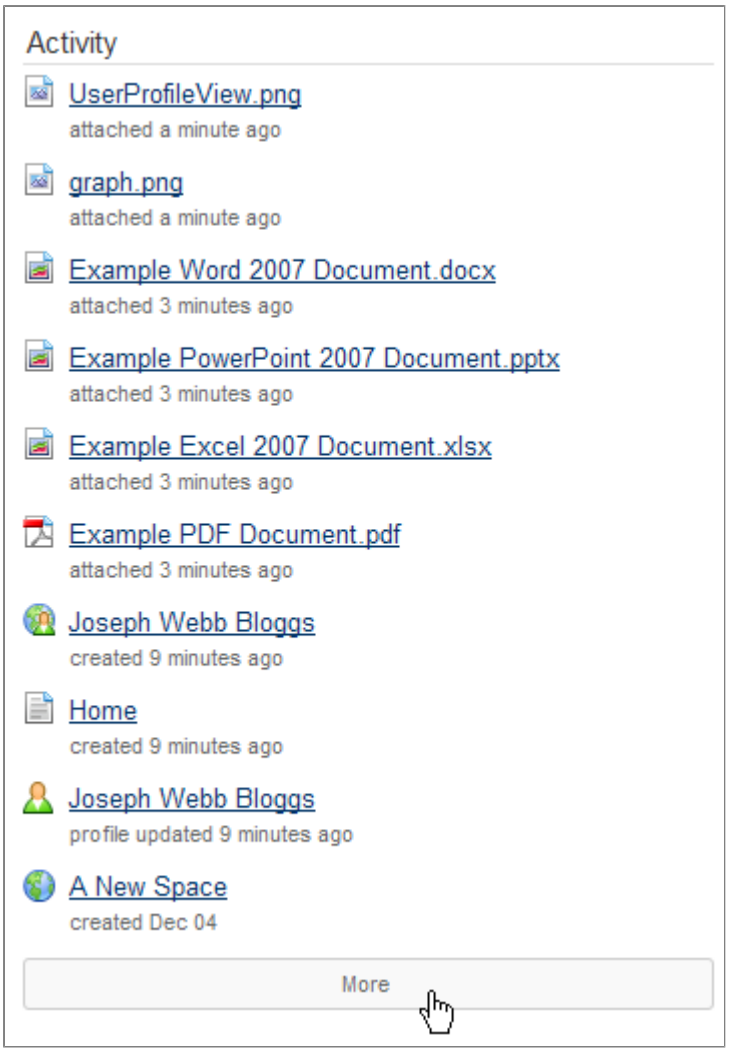
If you wish to link to a Confluence page from any other location on the web, use the convenient '**Link to this Page**' feature (available from any page's or blog post's '**Tools**' menu). Upon selecting this feature, the 'Link to this Page' dialog box opens, from which you can copy three versions of the link to embed elsewhere:

- **Link** – Standard URL which should work from any other accessible location on the web.
- **Tiny Link** – A reduced-length version of the 'Link', which can be used in text fields of limited length, such as tweets or Confluence Status Updates.
- **Wiki Markup** – A wiki markup version of the link, which can be used in any other location within your Confluence site.



Get 'More' from your Activity Streams

A 'More' feature has been added to various activity streams throughout the Confluence interface, including the user profile sidebar, user profile page and recently updated macro. Clicking 'More' expands the list of results, providing a convenient means of accessing progressively more distant user activities.



User Interface Performance Improvements

Most JavaScript and Cascading Style Sheet (CSS) files are now downloaded in one batch, greatly improving the performance of Confluence's editing features and general page rendering.

Other Small Enhancements and Improvements to Confluence

- Support for OAuth — With the introduction of gadgets (above) in this release, Confluence 3.1 now allows you to establish OAuth relationships with other web applications such as JIRA 4.0+, iGoogle, Gmail etc., thereby allowing them to share resources via gadgets.
- New Log In and Log Out screens.
- In an aim to minimise confusion, 'News Items' are now consistently called 'Blog Posts' throughout the Confluence interface and a list of blog posts is collectively referred to as a 'Blog'.
- Macro developers are now able to specify whether the macro body should or should not be displayed in Rich Text editor. For more information, please refer to CONF-12149.
- Other minor interface improvements.

Known Issues in this Release

We have an enthusiastic and dedicated group of testers and customers who jump in there, try out the new Confluence release and report any problems so that we can fix them quickly.

We value this feedback, which means that we can tell you about some minor known issues in Confluence 3.1. Sometimes we find out about a problem with the latest version of Confluence after we have released the software. In such cases we publish information about the known issues in the Confluence Knowledge Base. Please check the important technical advisories on the front page of the Knowledge Base.

A big **thank you** to everyone who helps us ensure that Confluence keeps getting better and better.

The Confluence 3.1 Team

Development

Bugfixing, Maintenance and Drag-and-Drop
Anatoli Kazatchkov
David Loeng

David Taylor

Engine Room
Agnes Ro
Charles Miller
Matthew Jensen

Gadgets and Office Connector
Chris Kiehl
Andrew Lynch
Xu-Heng Tjhin
Ryan Ackley
Jonathan Gilbert

Editor Enhancements
Paul Curren
Don Willis
Matt Ryall
Dmitry Baranovskiy

Small Improvements
Brian Nguyen
Chris Broadfoot

Plugin Updates
David Chui

Team Lead
Per Fragemann

Support

Kuala Lumpur
Sashidaran Jayaraman
Zed Yap
Arie Murdianto
Azwandi Mohd Aris
Ming Giet Chong

San Francisco
Jeremy Largman
Maleko Taylor
Tim Wong
Vincent Chang
Peter White
Marian Finch

Sydney
Partha Kamal
Gurleen Anand
Roy Hartono
Michael Seager
Ivan Benko

Others

Design
Stephen Russell

Performance Engineering
George Barnett

Product Management
Audra Eng
Jens Schumacher

Product Marketing Management
Bill Arconati
Matthew Hodges

Quality Assurance
Mark Hrynczak
Penny Wyatt
Andrew Prentice

Technical Writing
Giles Gaskell

Edwin Dawson

Confluence 3.1 Known Issues

Below are some known issues associated with Confluence 3.1.

On this page:

- JIRA/Crowd and Confluence deployment
- JIRA Gadgets in Confluence
- Bamboo integration
- PDF exports only render gadgets as links
- Problematic Confluence Gadgets window when running Confluence on Java 6
- Other issues

JIRA/Crowd and Confluence deployment

Confluence will not start up or will display strange behaviour (drop down menus not working) if JIRA 4.0/4.0.1 or Crowd 2.0.x is running on the same application server installation, for example, by attempting to run Confluence and JIRA 4.0 in the same Apache Tomcat server installation. This problem results from a bug in JIRA (tracked as JRA-19894) that is scheduled to be fixed in JIRA 4.0.2. Crowd will be fixed in the 2.1 release. In the meantime, please read our KB article on how to resolve this issue.

In the meantime, you can run JIRA or Crowd and Confluence in different 'instances' of the same application server, for example two separate Apache Tomcat server installations. If you do so already or have installed standalone distributions of both Confluence and JIRA 4.0 or Crowd 2.0.x, you can ignore this known issue.

JIRA Gadgets in Confluence

As of Confluence 3.1, users can embed gadgets into Confluence pages. When integrating JIRA gadgets into Confluence pages, you may encounter UI problems like a missing "Login"-button which is required to make the gadget authenticate with the JIRA server. This problem can be circumvented by setting up JIRA and Confluence to use Trusted Apps communication (since it removes the need for manual authentication). See KB article for details.

Bamboo integration

Our continuous integration product Bamboo exposes gadgets which can be embedded into Confluence pages. However, some of these exhibit problems once embedded onto a Confluence page:

- BAM-4900 : Unable to edit Bamboo gadgets in Confluence
- BAM-4890 : Bamboo gadget added in JIRA dashboard is not saving the preferences

These bugs are being fixed in Bamboo 2.5, which will ship in January 2010.

PDF exports only render gadgets as links

If you place any gadget on a Confluence page and export the page to PDF, the gadget output will not be rendered in the PDF output. Instead, each gadget is rendered on a page as a box containing the name of the gadget, the latter of which is hyperlinked. Clicking this hyperlink, opens the gadget contents itself in a new browser window or tab.

Problematic Confluence Gadgets window when running Confluence on Java 6

The Confluence Gadgets window may indicate that '**An error has occurred while trying to load the Gadget Directory**' and prevent you from accessing the URLs of your Confluence gadgets. This problem can occur if you are running Confluence on Java 6. After you install Confluence 3.1 or upgrade an existing Confluence installation to this version, please check the Confluence Gadgets window immediately after starting the Confluence server.

If you see this error message and cannot access your Confluence gadgets, it can be resolved by restarting Confluence. (You may need to do this more than once.)

For more information about this issue, please refer to CONF-17417.

Other issues

Refer to our JIRA site for a list of Confluence 3.1-specific bugs.

Confluence 3.1 Upgrade Notes

Below are some important notes on upgrading to **Confluence 3.1**. For details of the new features and improvements in this release, please read the Confluence 3.1 Release Notes.

On this page:

- Upgrade Notes
  - New License Key Requirements for Confluence 3.1
  - Upgrading an Existing Confluence License for Confluence 3.1 Compatibility
  - Custom layouts must be re-implemented after upgrading Confluence
  - Customers running Confluence on Weblogic are required to specify the `prefer-web-inf-classes` element in the `weblogic.xml` file prior to upgrading Confluence
  - The Drag-and-Drop feature disables the ability to drag and drop links or text in Firefox 3.0
  - Clarification of supported user management configurations in Confluence

- Upgrade Procedure
- Checking for Known Issues and Troubleshooting the Confluence Upgrade

Upgrade Notes

New License Key Requirements for Confluence 3.1

> **In order to continue using Confluence, you must upgrade your Confluence license to Atlassian's new license key format.**

We have undertaken this change to enhance and improve the support we provide our customers. Bear in mind that this license upgrade will not incur any additional costs and does not change Confluence's functionality in any way.

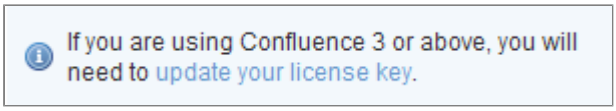
Upgrading an Existing Confluence License for Confluence 3.1 Compatibility

This procedure can only be performed by Confluence Administrators.

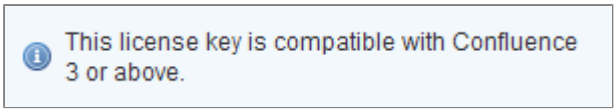
To upgrade your existing license to the new license key format, which will be required for continued use of Confluence 3.1:

1. Visit the license upgrade area in your account at my.atlassian.com.
2. Enter your Atlassian account details (email address and password) to access and manage your Atlassian product licenses.
3. If your Confluence license is already associated with a Confluence Server ID (that is, most customers running recent versions of Confluence), follow procedure **b** below. If my.atlassian.com prompts you to enter a Server ID before upgrading your license, follow procedure **a**:
    a. **Procedure for associating a Confluence license with your Confluence Server ID:**
        i. If you have not already upgraded your Confluence installation to version 3.1, go to your Confluence installation's License Details page to access the Server ID associated with your Confluence license. Make a note of this Server ID.
        ii. If you have already upgraded your Confluence installation to version 3.1 and cannot access your Confluence installation's License Details page, open the `confluence.cfg.xml` file (located inside the Confluence Data Directory) in a text editor and make a note of your Confluence installation's Server ID from the `confluence.setup.server.id` property in this file.
        iii. If you cannot access the Server ID in your Confluence installation using either of the two methods above, then please contact our customer support team for further assistance. This issue may occur when upgrading from a very old version of Confluence.
        iv. Follow the remaining prompts on my.atlassian.com to enter your Server ID and then upgrade your Confluence license.
    b. **Procedure for upgrading a Confluence license associated with a Server ID:**
        i. Select the appropriate Confluence license to expand its details.
        ii. In the 'info' note below your license on the right-hand side of the page, click the '**update your license key**' link (as shown in screenshot 1 below). Once this is done, the note changes to that shown in screenshot 2 below.

        Screenshot 1: License Key Upgrade Function

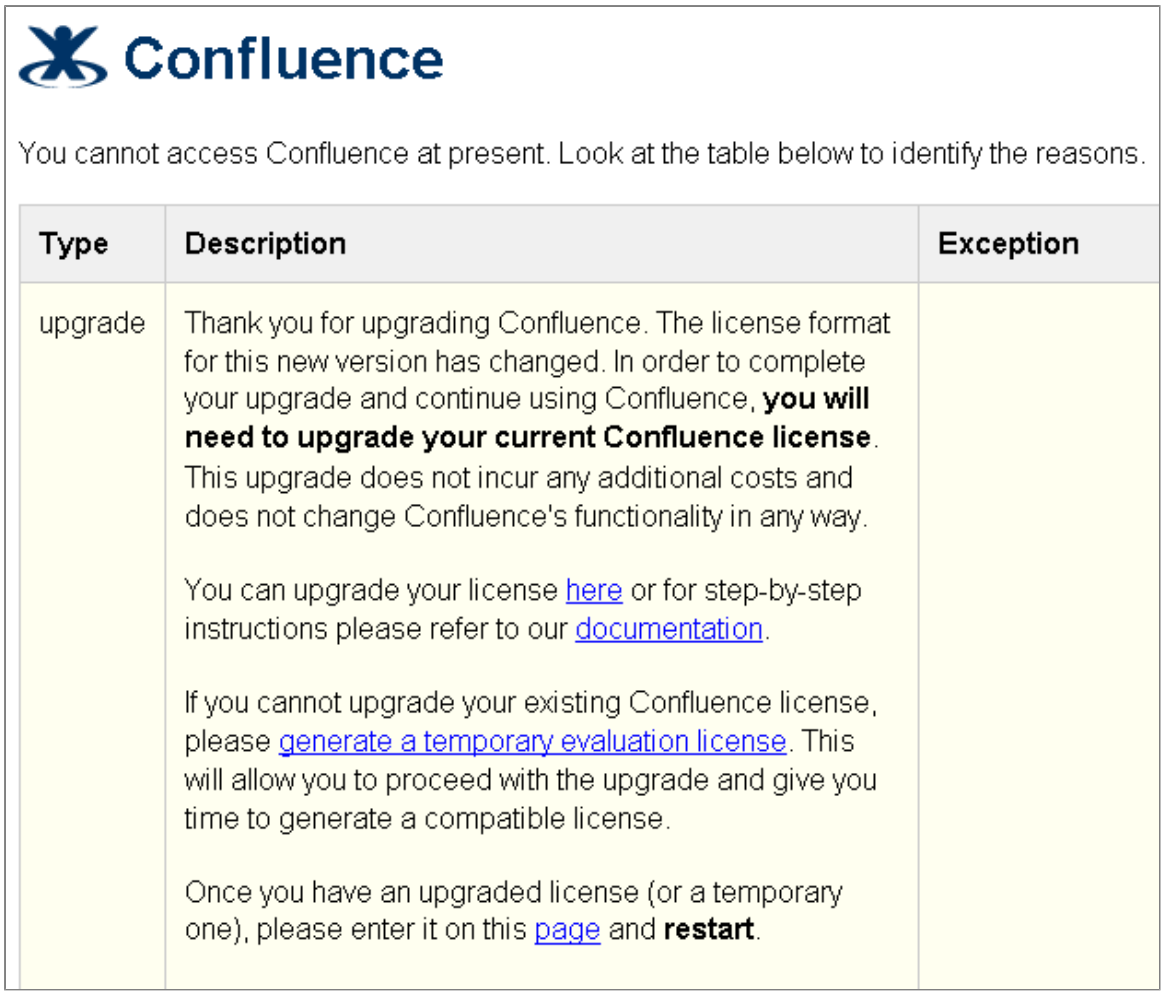
        

        Screenshot 2: Upgraded License Key Note

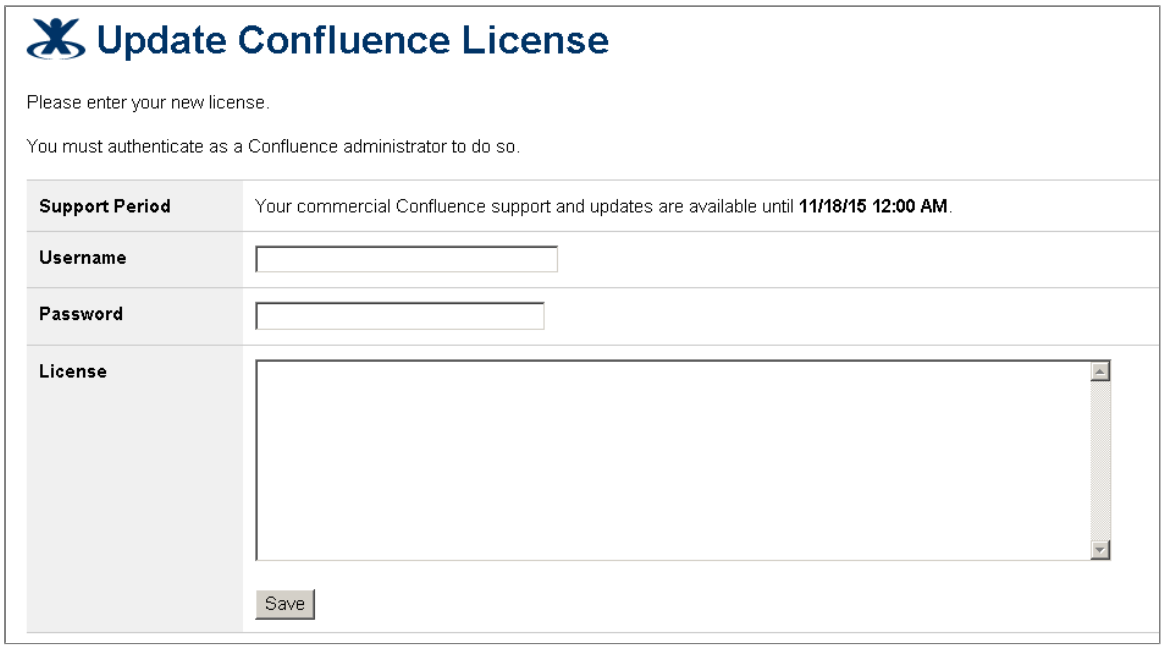
        iii. Copy the new license from the text box above this message to your clipboard.
4. Ensure your upgraded Confluence installation has been started, enter your Confluence site's URL into a web browser and on the Confluence license upgrade screen, click the '**page**' link in the 'Description' field to begin updating Confluence with your upgraded Confluence license to open the 'Update Confluence License' page.

    Screenshot: Confluence License Upgrade Screen



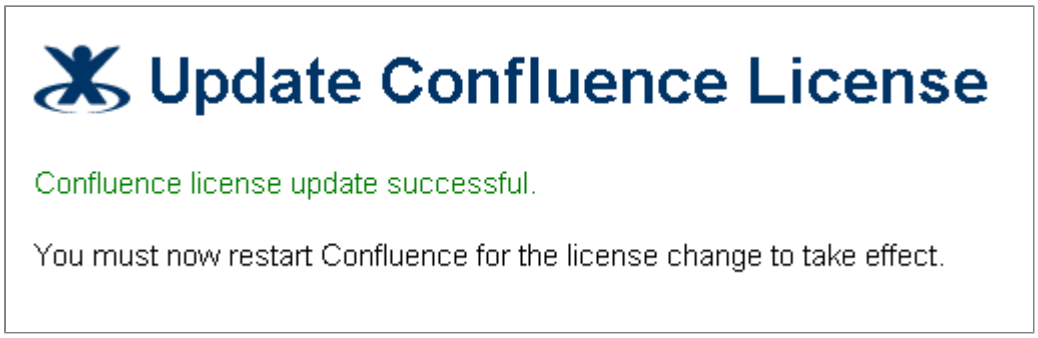
5. On the 'Update Confluence License' page, paste your new license details into the 'License' field.
6. Enter the *Confluence administrator* account's username and password details and click the '**Save**' button.

   Screenshot: Update Confluence License Page



7. If your license update was successful, you will be prompted to restart Confluence.

Screenshot: Successful Confluence License Update Screen



Custom layouts must be re-implemented after upgrading Confluence

If you have customised your Main Layout on either the space or the global level, or if you have a custom theme plugin, the new Link to this Page, Page Restrictions and 'Move Page' dialog box features in Confluence 3.1 will not immediately work for you.

To resolve these issues, you will need to re-implement your custom layouts or reset these custom layouts back to the Confluence default settings.

Other Issues:

The following other issues may occur prior to resolving these issues:

- The Atlassian Confluence footer may appear fixed on the page and obscure content that extends below the length of the web page.
- Some pages may not render at all.

Depending on the version of Confluence you upgraded from and the customisations that had been implemented, other user interface problems or problems with Confluence's functionality may be found.

Re-implementing custom layouts:

To re-implement your custom layouts, please refer to Upgrading Custom Layouts for details on retrieving the customisations made to your layouts and re-implementing them into your upgraded version of Confluence.

Customers running Confluence on Weblogic are required to specify the `prefer-web-inf-classes` element in the `weblogic.xml` file prior to upgrading Confluence

If you are a customer running Confluence on Weblogic, then before upgrading to Confluence 3.1, you must ensure that the `prefer-web-ing-classes` element in the `weblogic.xml` file has been specified with the content value of `true`. For more information, please refer to Installing Confluence EAR-WAR on Weblogic.

To do this:

1. Ensure that Confluence and Weblogic have been stopped.
2. Open the `weblogic.xml` file in a text editor. (This file is located in the `<confluence install directory>\confluence\WEB-INF` directory.)
3. Ensure that `<prefer-web-inf-classes>true</prefer-web-inf-classes>` has been added as a child element of the `<container-descriptor>` element, such that your `<container-descriptor>` element looks something like:

   ```
   <container-descriptor>
   <prefer-web-inf-classes>true</prefer-web-inf-classes>
   </container-descriptor>
   ```

   Your particular `weblogic.xml` may have other child elements of the `<container-descriptor>` element, so leave these intact.

4. Save any changes made to the `weblogic.xml` file.
5. Follow the upgrade procedure.

The Drag-and-Drop feature disables the ability to drag and drop links or text in Firefox 3.0

Some browsers like Firefox and Safari allow Confluence users to create links easily by dragging and dropping hyperlinks from other web pages directly into the rich text editor window. This browser-specific feature also allows the rearrangement of text when editing Confluence wiki page content, by highlighting text and dragging and dropping it elsewhere.

However, the Confluence 3.1 Drag-and-Drop feature is not compatible with these link creation and text rearrangement features of Firefox 3.0 and setting up Confluence's Drag-and-drop feature will disable these Firefox 3.0 features.

To allow the Confluence 3.1 Drag-and-Drop feature to work together with these link creation and text rearrangement features of Firefox, upgrade your Firefox browser to version 3.5.

Clarification of supported user management configurations in Confluence

In Confluence 3.1, Atlassian is clarifying our support for user management configurations and code-level customisations. This is being done so that we can deliver significant improvements in our user management performance and configuration in a future release.

Confluence is fully supported with the following configurations provided and documented by Atlassian:

- Built-in user management with Atlassian-User Hibernate managers ("default user management")
- Built-in user management, with OSUser LDAP authentication (deprecated since 2.7, moving to Atlassian-User LDAP is recommended)
- External user management with read-only JIRA JDBC providers ("JIRA delegated user management")
- External user management with Atlassian-User LDAP providers ("standard LDAP user management")
- External user management with Atlassian Crowd.

Unfortunately, we cannot offer complete support for code-level customisations in our user management system. The following caveats apply to customers who are using extensions to our user management systems:

- Custom implementations of Atlassian-User managers are supported for problems which are not related to user management.
- Custom Seraph authenticators are supported for problems which are not related to user management or authentication.
- Custom implementations of OSUser providers are not supported with Confluence. The OSUser APIs required for implementing custom providers were deprecated in Confluence 2.7, so customers with custom OSUser implementations will need to migrate their code to the Atlassian-User API to have support for their Confluence instance.

Advance warning: In Confluence 3.2, Atlassian will be removing functionality required for custom OSUser providers. From this point, custom OSUser providers will not work correctly with Confluence 3.2, so we recommend porting any custom OSUser providers to the Atlassian-User interfaces as part of your Confluence upgrade process.

Upgrade Procedure

> **Upgrade a test environment first**
>
> As always, please test your upgrades in your test environment before rolling into production.

If you are already running a version of Confluence, please follow these instructions to upgrade to the latest version:

1. Before you upgrade, we strongly recommend that you back up your Confluence Home Directory and database. See the documentation on backing up your Confluence site. If you are using an external database, perform a database backup.
2. If your version of Confluence is earlier than 3.0.x, then please read the Upgrade Notes Overview and the Upgrade Notes for each version of Confluence listed on that page. (There are hyperlinks to each one.) Furthermore:
    - If you are upgrading from 2.1 or earlier, please read the 2.2 release notes.
    - If you are upgrading from 2.2 or earlier, you will need to upgrade to Confluence 2.7.x first, confirm the upgrade was successful, then upgrade again from version 2.7.x to the latest. For more details, please refer to CONF-11767.
3. Download the latest version of Confluence.
4. Follow the instructions in the Upgrade Guide.

Checking for Known Issues and Troubleshooting the Confluence Upgrade

- **Check for known issues**. Sometimes we find out about a problem with the latest version of Confluence after we have released the software. In such cases we publish information about the known issues in the Confluence Knowledge Base. Please check the important technical advisories on the front page of the Knowledge Base and follow the instructions to apply any necessary patches.
- **Did you encounter a problem during the Confluence upgrade?** Please refer to the guide to troubleshooting upgrades in the Confluence Knowledge Base.
- If you encounter a problem during the upgrade and cannot solve it, please create a support ticket and one of our support engineers will assist you through the process.

RELATED TOPICS

Confluence 3.1 Release Notes

Issues Resolved in Confluence 3.1

Below are the issues resolved in Confluence 3.1, ordered by number of votes. For the full details of the fixes, improvements and new features, please take a look at our issue tracker. Please also take a look at the Confluence 3.1 Release Notes for the new features in Confluence 3.1.

JIRA Issues (184 issues)

Type	Key	Summary	Status	Resolution	Votes
	CONF-15413	Officially Support IE8	Resolved	Fixed	92
	CONF-6888	Some word docs don't get correctly indexed		Fixed	75

			Resolved		
	CONF-12006	Index the content of the new MS Office File Formats - docx xlsx etc.	Resolved	Fixed	67
	CONF-16079	Support RTE in Safari 4	Resolved	Fixed	36
	CONF-8472	Change 'News' to 'Blogs'	Resolved	Fixed	27
	CONF-12302	When inserting a Link in a Wiki Markup Mode in IE, the new link is included at the end of the page instead of where the cursor is	Resolved	Fixed	23
	CONF-14090	Notification mail for page-move displays wrong user when moving via Browse, Pages, Tree	Resolved	Fixed	15
	CONF-4455	'Watch' icon on and off states are not clearly defined	Resolved	Fixed	13
	CONF-6085	Can't find group members of group when DN does not include username	Resolved	Fixed	12
	CONF-1054	Drag & drop file uploading	Resolved	Fixed	12
	CONF-8504	Page permissions should somehow display the display name, not just the username or have a link to the user profile.	Resolved	Fixed	10
	CONF-15502	Edit Office documents from the attachments page and attachments macro	Resolved	Fixed	9
	CONF-13364	Upgrade to TinyMCE 3.2.x	Resolved	Fixed	9
	CONF-12864	PageNotFound action can render Confluence inoperable	Resolved	Fixed	8
	CONF-15827	Edit In Word is not working with OpenOffice or NeoOffice	Resolved	Fixed	7
	CONF-16989	Address frequency of cluster panic errors	Resolved	Fixed	6
	CONF-15055	Doc import doesn't work on websphere	Resolved	Fixed	6
	CONF-14681	New child pages are not sorted in the order specified for the space...	Resolved	Fixed	6
	CONF-12009	users with ldap usernames that are UPPERCASE fail isViewingMyProfile() thus edit tabs and createPersonalSpace links do not work when viewing your perferences	Resolved	Fixed	6
	CONF-7591	List installed plugins and install date under System Information	Resolved	Fixed	6
	CONF-9496	Some links in notification mail are in the language of the user who edited the page	Resolved	Fixed	5
	CONF-8584	Create a 'Documentation' theme	Resolved	Fixed	5
	CONF-13754	HibernateGroupManager.hasExternalMembership() is slow for group with thousands of users	Resolved	Fixed	4
	CONF-12329	Move included JavaScript tags to the bottom of HTML pages	Resolved	Fixed	4
	CONF-10199	Adding emotion (smiley characters) in tables automatically appends an undesirable line break	Resolved	Fixed	4
	CONF-16306	Macros are undesirably rendered in their own paragraph in 3.0	Resolved	Fixed	3
	CONF-16169	Macros that contain non-wiki-markup bodies are broken by rich text editor	Resolved	Fixed	3
	CONF-16168	File downloads over SSL don't work with cache control in IE browser		Fixed	3

			Resolved		
	CONF-15887	Queries in SQL macro are corrupted by editing with Rich Text Editor	Resolved	Fixed	3
	CONF-9604	Rich text editor removes XML within code tags when editing comment	Resolved	Fixed	3
	CONF-17742	Using IE7 to dowmload an attachments containing spaces in the filenames will be replace with underscore	Resolved	Fixed	2
	CONF-16518	Left Hand navigation theme is still using the old footer.vmd	Resolved	Fixed	2
	CONF-15988	Axis MultiRef should be disabled	Resolved	Fixed	2
	CONF-15902	Space Admin - PDF Stylesheet Page has no mention of HTML	Resolved	Fixed	2
	CONF-14239	blog-posts macro does not render "Read More..." link after excerpt	Resolved	Fixed	2
	CONF-17841	The 'enabled' state of modules that are disabled by default in bundled plugins can't survive (plugin) upgrades	Resolved	Fixed	1
	CONF-17600	Editor jumps to top of page when saving draft in IE8	Resolved	Fixed	1
	CONF-17398	INDEXQUEUEENTRIES table not being cleaned despite the index queue clean job being running at 2am every day with no errors	Resolved	Fixed	1
	CONF-17322	IE8 complains that some parts of the edit page are not secure. Suspecting the hamster	Resolved	Fixed	1
	CONF-17136	NonClusterManager does not prevent concurrent execution of jobs	Resolved	Fixed	1
	CONF-17016	Searches for users are randomly ordered with Crowd version 1.6.2 or greater	Resolved	Fixed	1
	CONF-16743	Add Powerpoint 2007 support to the viewfile macro	Resolved	Fixed	1
	CONF-16731	Indexing language "Custom Japanese" cannot be selected	Resolved	Fixed	1
	CONF-16717	Add Excel 2007 support to viewfile macro	Resolved	Fixed	1
	CONF-16707	Word 2007 support	Resolved	Fixed	1
	CONF-16660	Implement Prototype of the Confluence REST API	Resolved	Fixed	1
	CONF-16543	Implement the fix for CONF-16348 on trunk: Attachment File Not Found - in children pages when a page is moved to another space	Resolved	Fixed	1
	CONF-16393	Resuming a blog post draft caches the object as a Page	Resolved	Fixed	1
	CONF-16380	RaiseSupportRequestAction throws NullPointerException due to improperly loaded plugin	Resolved	Fixed	1
	CONF-16321	Edit User Profile does not work with Captcha	Resolved	Fixed	1
	CONF-16308	Add ability to view the change before resuming or discarding	Resolved	Fixed	1
	CONF-16227	Recently-updated not filtering emails	Closed	Fixed	1
	CONF-16164	Some strings in space.vmd are not internationalised	Resolved	Fixed	1
	CONF-16114	Provide easier access to Tiny URL and Wiki Markup link via modal dialog		Fixed	1

			Resolved		
	CONF-16075	When editing, prompt the user for confirmation if they close the tab or window or navigate to a different page.	Resolved	Fixed	1
	CONF-16056	Profiling is hard to understand - doesn't show times for macro rendering	Resolved	Fixed	1
	CONF-15086	Cannot create a page with a duplicate title, even though the page should be saved to a different space which does not contain the duplicate	Resolved	Fixed	1
	CONF-14319	Excpetion "not legal for a JDOM character content: 0x1d is not a legal XML character." in RSS feed	Resolved	Fixed	1
	CONF-12349	Remove news link still visible for those user/group which is not granted with remove news permission	Resolved	Fixed	1
	CONF-11488	Page title is lost after selecting a template (Default theme)	Resolved	Fixed	1
	CONF-9719	Add information about (outbound HTTP) proxy settings to System Info	Resolved	Fixed	1
	CONF-9580	Resources with same path in different plugins override one another	Resolved	Fixed	1
	CONF-3551	Consecutive divs in PDF export not separated by whitespace	Resolved	Fixed	1
	CONF-18823	Some powerpoint files can potentially hang indexing threads causing a full reindex to sit at 99% indefinitely	Resolved	Fixed	0
	CONF-18129	Rich Text Editor does not work in Safari 4.0.4	Resolved	Fixed	0
	CONF-17967	XSS vulnerability in pagetree and page macros	Resolved	Fixed	0
	CONF-17943	Confluence Source Distribution fails to build without coherence jars	Resolved	Fixed	0
	CONF-17930	Some Word docs containing tables appear blank when using the Office Connector	Resolved	Fixed	0
	CONF-17827	Gadgets do not work with Confluence running on WebLogic	Resolved	Fixed	0
	CONF-17761	User does not receive email notification when a page gets deleted if he is a last modifier of the page	Resolved	Fixed	0
	CONF-17669	Improve styles of footer	Resolved	Fixed	0
	CONF-17572	Can't display particular Excel files in the \{viewfile\} macro	Resolved	Fixed	0
	CONF-17570	Page content is not visible in edit wiki markup mode in IE6	Resolved	Fixed	0
	CONF-17517	Add a log entry in atlassian-confluence.log when importing a word document or display the attachments that the plugin is processing in the admin page	Resolved	Fixed	0
	CONF-17496	New version of attachment added via RPC changes case of file name	Resolved	Fixed	0
	CONF-17463	DraftId is not set for pages created using templates	Resolved	Fixed	0
	CONF-17453	Move Page dialog is giving an incorrect error message when trying to move a page to its own descendant	Resolved	Fixed	0
	CONF-17389	Table resizing controls in FF are misaligned	Resolved	Fixed	0
	CONF-17370	Login/Log In grammar incorrect	Resolved	Fixed	0
	CONF-17361	XSS vulnerability can be exploited using the Gallery macro		Fixed	0

			Resolved		
	CONF-17360	Add Dashboard widget to allow easy creation of pages and posts right from the Dashboard	Resolved	Fixed	0
	CONF-17318	Page menus on IE8 lack icons and the spacing is too dense	Resolved	Fixed	0
	CONF-17310	Allowing drag and drop attachments onto Confluence and its editor	Resolved	Fixed	0
	CONF-17282	Tweak attachment upload UI and add an extension zone	Resolved	Fixed	0
	CONF-17256	Wrap Func Test RPC Impl with a proxy to check for administrator access	Resolved	Fixed	0
	CONF-17255	Viewfile macro slide viewer should support full screen keyboard interactivity available in Flash 10	Resolved	Fixed	0
	CONF-17248	Bundled Plugin dependency model causing problems with source build	Resolved	Fixed	0
	CONF-17225	Viewxls macro causes out of memory error and then fails to render	Resolved	Fixed	0
	CONF-17196	Improve SystemErrorInformationLogger	Resolved	Fixed	0
	CONF-17191	Allow macro icons of different dimensions in macro browser	Resolved	Fixed	0
	CONF-17171	Make the whole row in page history view clickable instead of just tiny checkbox	Resolved	Fixed	0
	CONF-17166	Bundle Confluence News gadget with Confluence 3.1	Resolved	Fixed	0
	CONF-17160	Confluence fails on Turkish system locale	Resolved	Fixed	0
	CONF-17151	CSS doesn't load and JavaScript errors due to too many CSS files in IE	Resolved	Fixed	0
	CONF-17133	Instructions in atlassian-user.xml for Crowd integration are incorrect	Resolved	Fixed	0
	CONF-17127	Consecutive Italic lines break in RTE round-trip	Resolved	Fixed	0
	CONF-17118	Link to missing space loses destination using Wysiwyg Editor	Resolved	Fixed	0
	CONF-17117	Make personal blog more visible	Resolved	Fixed	0
	CONF-17105	ExportDownloadResourceManager can't be autowired in atlassian-plugin-sdk-3.0-beta5	Resolved	Fixed	0
	CONF-17100	Expose ServletModuleManager as a host component	Resolved	Fixed	0
	CONF-17096	Default bundled languages not available after startup	Resolved	Fixed	0
	CONF-17094	Create windows evaluation Installer	Resolved	Fixed	0
	CONF-17086	Background images/icons for menu elements and don't display in IE8	Resolved	Fixed	0
	CONF-17057	Create "Insert" menu for RTE	Resolved	Fixed	0
	CONF-17045	Add admin config to serve javascript in the header	Resolved	Fixed	0
	CONF-17035	Query for determing alternate pages out of a space is failing on oracle with:		Fixed	0

		"ORA-00907 missing right parenthesis"	Resolved		
	CONF-17030	Query to find page alternatives fails on DB2 because we are invoking SELECT DISTINCT on the VERSIONCOMMENT clob column	Resolved	Fixed	0
	CONF-16983	ConfluenceCachingBandanaPersister logs not serializable warnings for null values	Resolved	Fixed	0
	CONF-16982	Change JsonResult content type to text/javascript	Resolved	Fixed	0
	CONF-16957	Minor footer styling improvements	Resolved	Fixed	0
	CONF-16954	Force footer to appear at the bottom of the page.	Resolved	Fixed	0
	CONF-16944	Add the ability to subscribe to the success or failure of a draft save	Resolved	Fixed	0
	CONF-16919	Make rich text editor more extensible by introducing a plugin point to full screen editor initialisation and exposing references to editors	Resolved	Fixed	0
	CONF-16907	Index job fails due to HibernateObjectRetrievalFailureException	Resolved	Fixed	0
	CONF-16893	New improved lightweight image dialog	Resolved	Fixed	0
	CONF-16888	indexbrowser.jsp displays documents but links to details display nothing	Resolved	Fixed	0
	CONF-16885	Provide more than one hook into point at which the editor has been initialised	Resolved	Fixed	0
	CONF-16857	Servlet packages exported with the no version	Resolved	Fixed	0
	CONF-16817	Allow more updates to be retrieved via the "get more" link in the recently updated macro	Resolved	Fixed	0
	CONF-16802	AttachmentCreateEvent not published by Rpc method	Resolved	Fixed	0
	CONF-16793	Export Page PDF menu item does not have ID	Resolved	Fixed	0
	CONF-16790	Log in screen UI improvements	Resolved	Fixed	0
	CONF-16784	Promote PAC from the plugins management screen	Resolved	Fixed	0
	CONF-16776	Smart fields in the Macro Browser for single parameter types	Resolved	Fixed	0
	CONF-16719	Add index flush debug logging	Resolved	Fixed	0
	CONF-16702	Make the Macro Browser not rely on the macro name as a unique key	Resolved	Fixed	0
	CONF-16687	Help Tips - external links open in new tab/window	Resolved	Fixed	0
	CONF-16679	Image upload mechanism within AUI dialog	Resolved	Fixed	0
	CONF-16677	Backend for listing and providing JSON representation of image attachments	Resolved	Fixed	0
	CONF-16659	Extract Confluence decoration process out of ProfilingPageFilter	Resolved	Fixed	0
	CONF-16654	When saving attachment data Confluence should indicate if the data provided is incomplete	Resolved	Fixed	0
	CONF-16601	"Create Personal Space" menu item missing for some users, particularly LDAP		Fixed	0

		users	Resolved		
	CONF-16555	PDF Export is calculating the width of pre blocks incorrectly if it doesn't contain spaces	Resolved	Fixed	0
	CONF-16533	Add "Get More" to updates on profile page and on the personal space sidebar	Resolved	Fixed	0
	CONF-16524	Memory leak if Confluence can't create a log file in the home directory	Resolved	Fixed	0
	CONF-16516	Cannot remove all admin space permissions and add new one in one step	Resolved	Fixed	0
	CONF-16504	Display caps-lock indicator icon for password fields in Safari	Resolved	Fixed	0
	CONF-16501	Safari search placeholder is lowercase and not internationalised	Resolved	Fixed	0
	CONF-16475	User status operations (add/remove/clear) should produce events	Resolved	Fixed	0
	CONF-16463	Improve "Browse > Attachments" Page	Resolved	Fixed	0
	CONF-16441	Replace usages of Prototype.js with jQuery	Resolved	Fixed	0
	CONF-16412	Show Page Permissions in a dialog on the View page	Resolved	Fixed	0
	CONF-16363	Confirmation message is incorrectly seen when discarding a draft while editing	Resolved	Fixed	0
	CONF-16314	Inconsistency between using space name and space key in error messages reporting duplicate page titles	Resolved	Fixed	0
	CONF-16298	Improve System Information reporting	Resolved	Fixed	0
	CONF-16250	The sentence in view personal labels page is misleading	Resolved	Fixed	0
	CONF-16237	Log system info on startup	Resolved	Fixed	0
	CONF-16226	RSS-bug can cause Websphere to run out of file handles	Resolved	Fixed	0
	CONF-16177	Restructure the Bundled plugins so the source can be resolved in IDEA	Resolved	Fixed	0
	CONF-16157	Maximum Attachments per Upload setting not working	Resolved	Fixed	0
	CONF-16148	Correction to the content by label wording on the macro browser is required. What's currently there is misleading.	Resolved	Fixed	0
	CONF-16146	Use of action explicitely in vm files for consistency	Resolved	Fixed	0
	CONF-16123	Attachment links (Download/View) do not apply styleClass	Resolved	Fixed	0
	CONF-16122	Search Results do not apply style to item links	Resolved	Fixed	0
	CONF-16106	Improve Table Styling	Resolved	Fixed	0
	CONF-16101	More "Edit" links for Office Connector documents.	Resolved	Fixed	0
	CONF-16093	Add debugging for DefaultConfluenceIndexManager to trunk	Resolved	Fixed	0
	CONF-16031	Use inline modal dialogs for thumbnail-previews		Fixed	0

			Resolved		
	CONF-16015	DefaultIndexQueueEntryFilter and ChainedSearchFilter use Apache Commons classes from wrong package	Resolved	Fixed	0
	CONF-15663	RebuildIndexUpgradeTask only runs on one cluster node	Resolved	Fixed	0
	CONF-15646	Prevent global settings from being accidentally overwritten	Resolved	Fixed	0
	CONF-15621	Username with ampersand breaks the profile macro, follow macro and personal space sidebar	Resolved	Fixed	0
	CONF-15616	Admin user search does not trim preceeding whitespaces	Resolved	Fixed	0
	CONF-15163	Confluence ships with a javax.xml.namespace.QName class	Resolved	Fixed	0
	CONF-14983	Add sections for wiki editor toolbar web-item plugins	Resolved	Fixed	0
	CONF-14745	Typo in message when trying to move attachment to a non-existent page	Resolved	Fixed	0
	CONF-14609	People Directory link on dashboard has duplicate ID	Resolved	Fixed	0
	CONF-14538	Change left-hand navigation style and add collapse/expand.	Resolved	Fixed	0
	CONF-14530	When commenting on a page in WikiMarkup mode, there is no way for me to access the notation guide.	Resolved	Fixed	0
	CONF-14218	simple usability: in attach image dialog, click to select image should focus on the Submit button	Resolved	Fixed	0
	CONF-13962	Uninstalling a plugin with a servlet filter prevents Confluence from answering any requests.	Resolved	Fixed	0
	CONF-13731	Images inserted in comments do not display in comment preview	Resolved	Fixed	0
	CONF-13399	standard sitemesh filter should be applied to plugin-generated content so that it can be decorated if needed	Resolved	Fixed	0
	CONF-13285	Axis webservice may be asked to deliver invalid XML characters	Resolved	Fixed	0
	CONF-13159	AttachmentPermissionsDelegate does not include a permissions check for drafts	Resolved	Fixed	0
	CONF-13099	SWF file content cannot be played together with its skin in Internet Explorer 6 and 7	Resolved	Fixed	0
	CONF-12886	Members of confluence-administrators receive notifications for changes to restricted pages, including daily update	Resolved	Fixed	0
	CONF-12848	High bit characters causing webservice axisfaults encoding exception	Resolved	Fixed	0
	CONF-12843	Increase Cache value for Bandana Persiter	Resolved	Fixed	0
	CONF-12685	Unable to highlight page in tree to enable it to be moved	Resolved	Fixed	0
	CONF-11699	Titles of news posts cannot be corrected for capitalization	Resolved	Fixed	0
	CONF-10631	Tree refreshes into a collapsed state after a page move	Resolved	Fixed	0
	CONF-9228	Incorrect file permissions on index timestamp are not reported	Resolved	Fixed	0
	CONF-3899	NullPointerException when using {excerpt} in Global Templates		Fixed	0

			Resolved		

Above are the issues resolved in Confluence 3.1, ordered by number of votes. For the full details of the fixes, improvements and new features, please take a look at our issue tracker.

Confluence 3.0.2 Release Notes

6 October 2009

Confluence 3.0.2 is a recommended upgrade which fixes some security flaws and other issues.

Please refer to the security advisory for details of the security vulnerabilities, risk assessments and mitigation strategies.

> **Critical issue affecting non-clustered implementations of Confluence 3.0.2**
> Non-clustered (i.e. you do not have a clustered license) implementations of Confluence 3.0.2 are affected by an issue that can cause Confluence to crash. Please read the Confluence 3.0.2 Upgrade Notes for details on the issue and instructions on how to address it.

Editing and Visual Improvements

A bug in the Rich Text Editor lead to the generation of line break and other character formatting problems after saving a page. This bug has been fixed.

The sizes of some headings were considered too similar to be visually distinguishable on a page, especially when the headings were separated by intervening text content. This was particularly the case for heading sizes 2 and 3. Hence, the sizes of headings were modified to make them visually more distinct.

The format of colours used in Confluence's user profile areas has been modified slightly to make headings more prominent and form labels clearer.

The blog posts macro was missing the 'Restrict to These Authors' (author) parameter from the macro browser. However, this parameter is now available in the macro browser.

Other Enhancements and Fixes

Some customers' users experienced long delays while logging into Confluence, especially when their user accounts belonged to groups containing a large number of other user accounts. This issue was fixed.

Customers were able to generate Confluence page PDF exports directly from external web sites by adding the 'Export to PDF' link (accessible via a Confluence page's 'Tools' menu) to their external web pages. Unfortunately, this function was broken by the introduction of the form token handling security enhancement feature introduced in Confluence 3.0. In Confluence 3.0.2, however, this issue was resolved.

Some customers experienced an issue in which automatic content indexing would stop. This problem has been resolved.

When browsing Active Directory groups in Confluence, it was not possible to view group members if the LDAP Distinguished Names (DN) did not include the username. This bug was fixed.

There's a complete list of fixes below. Click a specific issue to see details of the fix.

Don't have Confluence 3.0 yet?

Take a look at the new features and other highlights in the Confluence 3.0 Release Notes.

Download Latest Version

Upgrading from a Previous Version of Confluence

Upgrading Confluence should be fairly straightforward. Please read the Confluence 3.0.2 Upgrade Notes. We *strongly recommend* that you back up your `confluence.home` directory and database before upgrading.

Updates and Fixes in this Release

JIRA Issues (20 issues)					
Type	**Key**	**Summary**	**Priority**	**Status**	**Resolution**
	CONF-16651	XSS vulnerability can be exploited with the pagetree macro		Resolved	Fixed
	CONF-16644	XSS vulnerability can be exploited with the Userlister macro		Resolved	Fixed
	CONF-15440	XSS vulnerability can be exploited with the contentbylabel macro		Resolved	Fixed

	CONF-15108	Session Fixation attack using JSESSIONID in Confluence		Resolved	Fixed
	CONF-13754	HibernateGroupManager.hasExternalMembership() is slow for group with thousands of users		Resolved	Fixed
	CONF-16459	PDF export link cannot be published to other sites...		Resolved	Fixed
	CONF-16428	Saving a page can lead to round-trip errors that do not occur by just switching tabs.		Resolved	Fixed
	CONF-15585	Use #333 for Headings in Confluence and #666 for labels		Resolved	Fixed
	CONF-14512	Newline lost between panel macro and table or list breaking markup		Resolved	Fixed
	CONF-9954	h2 and h3 are too similar in Confluence 2.6		Resolved	Fixed
	CONF-9575	Content Indexing stops		Resolved	Fixed
	CONF-8148	Cluster safety job should be made more generic and report multiple deployments with same DB as well		Resolved	Fixed
	CONF-6085	Can't find group members of group when DN does not include username		Resolved	Fixed
	CONF-17159	The new {code} macro puts line numbers in text when I copy/paste		Resolved	Duplicate
	CONF-16955	Support Entitlement Number is listed twice on the System Information page		Resolved	Fixed
	CONF-16774	Allow system plugins to be enabled		Resolved	Fixed
	CONF-16089	The blog posts macro is missing the 'author' parameter from the macro browser.		Resolved	Fixed
	CONF-13635	Show permgen, xmx and xms memory settings in the System Info		Resolved	Fixed
	CONF-16745	Change german translation on configuration page: Am --> Ein		Resolved	Fixed
	CONF-16683	superfluous </table> in /includes/common-listdecorators.vm		Resolved	Fixed

Click here to open a report on http://jira.atlassian.com for Resolved or Closed issues in Confluence 3.0.2.

Confluence 3.0.2 Upgrade Notes

Below are some important notes on upgrading to **Confluence 3.0.2**. Confluence 3.0.2 is a recommended upgrade which fixes some security flaws as well as other bugs. For more details, please read the Confluence 3.0.2 Release Notes.

On this page:

- Upgrade Notes
  - ClusterManager for Non-Clustered Licenses causes Critical Errors
  - New License Key Requirements for Confluence 3.1
  - Upgrading an Existing Confluence License for Confluence 3.1 Compatibility
- Upgrade Procedure

Upgrade Notes

ClusterManager for Non-Clustered Licenses causes Critical Errors

The ClusterManager in the non-clustered (i.e. you do not have a clustered license) implementations of Confluence 3.0.1 and Confluence 3.0.2 does not perform any locking (e.g. for job synchronisation). As a result, your Confluence instance may crash due to certain jobs being executed concurrently when they shouldn't be.

We **strongly recommend** that you upgrade to Confluence 3.1 or later to avoid this issue, if possible. If you wish to install or upgrade to a non-clustered implementation of Confluence 3.0.1 or 3.0.2, you must apply the patch attached to CONF-17136 after upgrading.

New License Key Requirements for Confluence 3.1

The next major release of Confluence (version 3.1) will require Confluence administrators to upgrade (or have upgraded) their Confluence license to Atlassian's new license key format. We are undertaking this change to enhance and improve the support we provide our customers. Bear in mind that this license upgrade will not incur any additional costs and does not change Confluence's functionality in any way.

Confluence 3.0.2 will still operate as usual under your existing Confluence license. However, any customers running Confluence 3.0.0 or later will be able to upgrade their Confluence license to the new license key format before Confluence 3.1 is released. Customers whose licenses expire before Confluence 3.1 is released *will be required* to upgrade their license to the new license key format.

As a consequence of these imminent licensing changes, relevant pages of the Administration Console area in Confluence 3.0.2 contain warnings about Atlassian's new license key requirement, which will be mandatory in Confluence 3.1. Links to upgrade old licenses to the new license key format are provided in these warnings.

Upgrading an Existing Confluence License for Confluence 3.1 Compatibility

If you have just upgraded to Confluence 3.1:

Please refer to the License Upgrade procedure for Confluence 3.1 in the Confluence 3.1 Upgrade Notes instead.

The following procedure can only be performed by Confluence Administrators, on Confluence versions 3.0.0 or later.

To upgrade your existing license to the new license key format, which will be required for continued use of Confluence 3.1:

1. Visit the license upgrade area in your account at my.atlassian.com.

   If you are using Confluence 3.0.2, you can also access this site by visiting the opening page of the Administration Console (by selecting '**Browse**' -> '**Confluence Admin**' menu item) or the License Details page of your Confluence installation and clicking on the '**Upgrade this license key now**' link.
   If you are already using an upgraded Confluence license, this link will not appear on these pages.

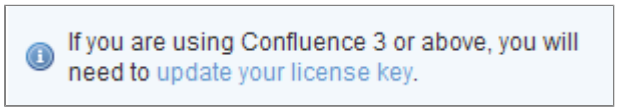
2. Enter your Atlassian account details (email address and password) to access and manage your Atlassian product licenses.

3. Select the appropriate Confluence license to expand its details.

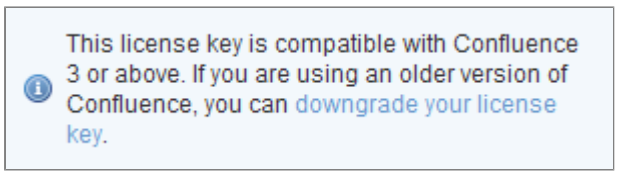
4. In the 'info' note below your license on the right-hand side of the page, click the '**update your license key**' link (as shown in screenshot 1 below). Once this is done, the note changes to that shown in screenshot 2 below.

   If the 'info' note looks like the one in screenshot 2, then your license key has been upgraded and you should not need to take any further action.

   Screenshot 1: License Key Update Function

   Screenshot 2: Updated License Key Note

5. Copy the new license from the text box above this message to your clipboard.

6. Visit the License Details page in your Confluence installation and paste the new license from your clipboard into the '**License**' field.

7. Click the 'Save' button. You will notice two changes:

   - If you are using Confluence version 3.0.2, the links to upgrade your license on the License Details and Administration Console opening pages will vanish.
   - A 'Support Entitlement Number' (SEN) will be assigned to your license on the License Details page.

Upgrade Procedure

If you are already running a version of Confluence, please follow these instructions to upgrade to the latest version:

1. Before you upgrade, we strongly recommend that you **back up your `confluence.home` directory and database**.

2. If your version of Confluence is earlier than 3.0.2, read the release notes and upgrade guides for **all releases** between your version and the latest version. In particular:
    - Please read the 3.0.2 release notes.
    - If you are upgrading from 2.1 or earlier, please also read the 2.2 release notes.
3. Download the latest version of Confluence.
4. Follow the instructions in the Upgrade Guide.

Confluence 3.0.1 Release Notes

20 August 2009

Confluence 3.0.1 is a recommended upgrade which fixes some security flaws and other issues.

Please refer to the security advisory for details of the security vulnerabilities, risk assessments and mitigation strategies.

Critical issue affecting non-clustered implementations of Confluence 3.0.1
Non-clustered (i.e. you do not have a clustered license) implementations of Confluence 3.0.1 are affected by an issue that can cause Confluence to crash. Please read the Confluence 3.0.1 Upgrade Notes for details on the issue and instructions on how to address it.

Attachment Handling Fixes

When a hierarchy of pages was moved from one space to another, the attachments on child or descendent pages of the parent page were not correctly moved. Instead, users would have to move one page at a time between spaces in order to maintain page attachment integrity. This issue has now been resolved.

Sporadic issues associated with attachment migration occurred when upgrading from either Confluence 2.9.x or 2.10.x, to 3.0. These have now been fixed.

Macro Fixes

A bug was identified in which excerpted content would not be rendered in a blog post macro until the source page containing that content had first been viewed. This phenomenon could occur when excerpt include macros were used in a blog post. It could also occur when excerpt macros were used in a blog post in conjunction with the `content=excerpts` blog post macro parameter. This problem has now been fixed.

An issue was found with the tasklist macro whereby special characters used in its title were not correctly escaped. This has now been resolved.

Rich Text Editor Fixes

An issue was found in which the Rich Text Editor's link removal feature (available from the right-click context menu) did not work with external links. This has now been resolved.

In the Firefox web browser, the spell checker is now automatically enabled by default in the Rich Text Editor. Users no longer have to first disable the right-click context menu and then enable and select 'Check Spelling' from the Firefox's own right-click context menu.

An issue was identified, which prevented the ability to escape from the quote text effect once it had been selected. This has now been addressed, such that a paragraph is automatically added after selecting this text effect.

Other Enhancements and Fixes

When a Confluence administrator first installs Confluence, runs through the Confluence Setup Wizard and then reaches the database configuration step, the database password is now hidden and is no longer shown in clear text.

It is now possible to filter network RSS feeds by different content types. This is achieved by implementing parameter modifications to the RSS feed link in your RSS newsreader. For more information, please refer to Subscribing to a Network RSS Feed.

Some customers experienced problems importing their site backup from a previous version of Confluence into version 3.0. This has now been resolved in Confluence 3.0.1.

An issue was identified in which multiple blog posts posted on a single day would be listed out of chronological order. This has now been fixed and multiple blog posts posted on a single day are now ordered according to their time of creation.

A problem was identified when accessing Confluence content in Internet Explorer that caused file downloads to fail over an SSL connection. This problem has now been fixed in this release of Confluence.

An issue was found in the page tree views on instances of Confluence running on Weblogic 10.x. This has now been resolved.

There's a complete list of fixes below. Click a specific issue to see details of the fix.

Don't have Confluence 3.0 yet?

Take a look at the new features and other highlights in the Confluence 3.0 Release Notes.

Download Latest Version

Upgrading from a Previous Version of Confluence

Upgrading Confluence should be fairly straightforward. Please read the Confluence 3.0.1 Upgrade Notes. We *strongly recommend* that you back up your `confluence.home` directory and database before upgrading.

Updates and Fixes in this Release

JIRA Issues (44 issues)					
Type	**Key**	**Summary**	**Priority**	**Status**	**Resolution**
	CONF-16141	Directory traversal in Profile Picture path - leads to privilege escalation in < 3.0		Resolved	Fixed
	CONF-16136	XSS vulnerability can be exploited on the WebDAV Configuration page		Resolved	Fixed
	CONF-16135	XSS vulnerability in space name when page move would create a duplicate		Resolved	Fixed
	CONF-16348	Attachment File Not Found - in children pages when a page is moved to another space		Resolved	Fixed
	CONF-16019	XSS vulnerability when moving page between spaces		Resolved	Fixed
	CONF-16509	Upgrading from any version before 2.9 to 3.0 doesn't migrate attachments and/or breaks custom space logos		Resolved	Fixed
	CONF-16466	Attachment migration from 2.10 to 3.0 fails		Resolved	Fixed
	CONF-16420	Update PDF export plugin to be compatible with new cluster/cache architecture		Resolved	Fixed
	CONF-16311	Build a Standard Edition of Confluence 3.0 (without Coherence)		Resolved	Fixed
	CONF-16225	Some bundled themes don't support web resource injection		Resolved	Fixed
	CONF-16209	XSS in PDF screen		Resolved	Fixed
	CONF-16084	Cannot filter a network feed by contentType		Resolved	Fixed
	CONF-16016	The JIRA Issues Macro in the Macro Browser is missing two parameters - "renderMode" and "baseurl".		Resolved	Fixed
	CONF-16014	Blog Posts Macro only renders excerpts if target page has been rendered		Resolved	Fixed
	CONF-16005	The Favourite Pages Macro in the Macro Browser is missing its "Maximum Number of Results" parameter.		Resolved	Fixed
	CONF-15997	Invalid error message when Updating status and session expired		Resolved	Fixed
	CONF-15970	XSS in user links		Resolved	Fixed
	CONF-15940	Server Base URL not set when sending a support request email...		Resolved	Fixed
	CONF-15923	Unlink in RTE doesn't work for external links		Resolved	Fixed
	CONF-15908	Tasklist macros double escaping titles in IE		Resolved	Fixed
	CONF-15788	Unable to import site backup during set up of Confluence 3.0		Resolved	Fixed
		Cursor jumps to beginning of the page from new paragraph after cancelling the			

	CONF-15751	Macro Browser on Firefox		Resolved	Fixed
	CONF-15722	The DynamicTaskList2 plugin provides an explicit description making it impossible to internationalised		Resolved	Fixed
	CONF-15701	Passwords are visible when configuring database		Resolved	Fixed
	CONF-15610	New look for user email preferences needs its layout fixed		Resolved	Fixed
	CONF-15001	Firefox 3 does not enable the spell checker on the comment text area by default		Resolved	Fixed
	CONF-14276	Unable to create renderer-component module in plugins2		Resolved	Fixed
	CONF-13702	Session must not be invalidated on logout		Resolved	Fixed
	CONF-13482	Can't get out of blockquote format in the rte		Resolved	Fixed
	CONF-12576	Duplicate friendly cache names which result in cache statistics not being visible		Resolved	Fixed
	CONF-10607	Blog posts appear in wrong order		Resolved	Fixed
	CONF-8098	User browser shows duplicate accounts when a user exists both locally and in LDAP		Resolved	Fixed
	CONF-16552	Renaming the ehcache config file fails on windows		Resolved	Fixed
	CONF-16218	Add a note into the admin screen explaining that you can use wiki markup and even an include macro to put an actual page onto the dashboard.		Resolved	Fixed
	CONF-16037	Misleading message when removing page permissions through info page		Resolved	Fixed
	CONF-16017	The JUnit Macro in the Macro Browser is missing one parameter - "debug".		Resolved	Fixed
	CONF-15995	User Profile page does not indicate why name and e-mail address is uneditable when LDAP integrated.		Resolved	Fixed
	CONF-15948	Parameters passed to jQuery extend method are in wrong order		Resolved	Fixed
	CONF-15920	User Hover is not working for a username which contains plus characters		Resolved	Fixed
	CONF-15152	Images with mimetypes that do not match extension cannot be used as thumnails		Resolved	Fixed
	CONF-14592	Macro Browser icons have tooltips showing in IE		Resolved	Fixed
	CONF-12292	Draft Page Titles not displaying		Resolved	Fixed
	CONF-16028	Typo in log4j.properties		Resolved	Fixed
	CONF-15696	Context menu toggle icon loses the tick when an anonymous user switches to full-screen view		Resolved	Fixed

Click here to open a report on http://jira.atlassian.com for Resolved or Closed issues in Confluence 3.0.1.

Confluence 3.0.1 Upgrade Notes

Below are some important notes on upgrading to **Confluence 3.0.1**. Confluence 3.0.1 is a recommended upgrade which fixes some security flaws as well as other bugs. For more details, please read the Confluence 3.0.1 Release Notes.

On this page:

- Upgrade Notes

    - ClusterManager for Non-Clustered Licenses causes Critical Errors
    - Introduction of 'Standard' and 'Clustered' editions
    - Confluence Caching Layer Changes
- Upgrade Procedure

Upgrade Notes

If you have customised the cache settings in your installation of Confluence (e.g. for performance reasons), then please read ALL information on this page before upgrading to Confluence 3.0.1 or to a 'standard edition' of Confluence. This page also contains important information for customers who are upgrading their installation of Confluence to version 2.6 or later from September 2009.

ClusterManager for Non-Clustered Licenses causes Critical Errors

The ClusterManager in the non-clustered (i.e. you do not have a clustered license) implementations of Confluence 3.0.1 and Confluence 3.0.2 does not perform any locking (e.g. for job synchronisation). As a result, your Confluence instance may crash due to certain jobs being executed concurrently when they shouldn't be.

We **strongly recommend** that you upgrade to Confluence 3.1 or later to avoid this issue, if possible. If you wish to install or upgrade to a non-clustered implementation of Confluence 3.0.1 or 3.0.2, you must apply the patch attached to CONF-17136 after upgrading.

Introduction of 'Standard' and 'Clustered' editions

Oracle Coherence Licensing Change:

- Due to a license agreement change between Atlassian and Oracle over the Coherence technology, from September 2009, Confluence will be made available in two editions:
    - **Standard Edition** — Confluence with Ehcache's caching technology (available to customers with non-clustered Confluence licenses).
      If you are currently running a clustered installation of Confluence, please do not upgrade it with a standard edition of Confluence.
    - **Clustered Edition** — Confluence with Oracle's Coherence clustering and distributed caching technology (available to customers with Confluence clustered licenses only).
- For more information, please refer to the Coherence License Changes document.
- In standard editions of Confluence, Ehcache replaces the local caching functionality previously provided by the Coherence technology.

Confluence 3.0.1 is the first Confluence version to be released in two editions 'standard' and 'clustered'.

Also, from September 2009:

- Standard editions will be released for each previous major version of Confluence back to 2.6. These will be:
    - Available as Confluence versions 2.10.4, 2.9.3, 2.8.3, 2.7.4 and 2.6.3.
    - Available to customers with non-clustered Confluence licenses.
- The remaining previous versions of Confluence currently available from our download page (from 2.6 to 3.0 inclusively), will be re-released as clustered editions and will only be available to customers with Confluence clustered licenses.
- The installation files for all versions of Confluence prior to 2.6 (which are no longer supported) will be removed from the Atlassian web site and will no longer be available for download and installation.

Confluence Caching Layer Changes

Due to the caching layer changes in the 'standard editions' of Confluence, you will need to reapply any cache customisations made to your cache sizes and/or cache configuration file, if *all three* of the following points are applicable:

- You have implemented cache customisations to your Confluence installation's cache sizes (either via the Administration Console or cache configuration file).
- You have an existing installation of Confluence 3.0 or earlier (excluding the 'standard edition' versions) and will be upgrading to Confluence 3.0.1 or later (or one of the earlier standard editions).
- Your Confluence installation is using a non-clustered Confluence license.

If you customised your Confluence instance's cache settings via the Administration Console, please refer to the Reapplying Cache Size Modifications via the Administration Console section below.

If you customised your Confluence instance's cache settings by modifying the cache configuration file, please refer to the Reapplying Cache Configuration File Modifications section below.

Reapplying Cache Size Modifications via the Administration Console

To reapply your cache size modifications via the administration console:

1. Before you upgrade, use the procedure described on the Cache Statistics page to open the 'Cache Statistics' section of the Administration Console in the 'Advanced' view.
2. Print out this view or save a copy of the web page for later reference. (This contains your existing individual cache settings.)
3. After upgrading Confluence, view the 'Cache Statistics' section of your upgraded Confluence installation in the 'Advanced' view.
4. Use the same procedure describe on the Cache Statistics page to re-adjust the size of each cache based on the previous settings that you had printed out (or saved).

Reapplying Cache Configuration File Modifications

To maintain your existing cache configuration file settings, you will need to transfer any cache customisations you have implemented in the Coherence cache configuration file (`confluence-coherence-cache-config.xml`) to the relevant entries in the Ehcache cache configuration file (`ehcache.xml`).

Each cache has a `cache-mapping` element in the Coherence file (of which there is an equivalent `cache` element in the `ehcache.xml` file). Unfortunately, copying across your customisations is not quite a straightforward process because the Coherence file defines several 'caching schemes' to store the actual cache values, which in turn are referenced by the `cache-mapping` elements. In contrast, the `ehcache.xml` file does not support caching schemes and a cache's values are expressed explicitly in separate parameters of a `cache` element.

To convert your Coherence cache configuration file customisations across to the equivalent Ehcache file:

1. Open both the `confluence-coherence-cache-config.xml` and `ehcache.xml` files in a text editor. These files are located in the `<confluence-home>/config` directory.
   If you implemented your customisations in a version of Confluence prior to 3.0, you will most likely find the `confluence-coherence-cache-config.xml` file in the `<confluence-install>/confluence/WEB-INF/classes` directory.
2. In the customised `confluence-coherence-cache-config.xml` file:
   a. Identify the caching schemes that were customised in this file and make a note of the values of all its child elements.
      Typically, each caching scheme is located inside a `local-scheme` element and all of these are enclosed within the `cache-schemes` element, which appears towards the end of this file.
   b. Note each customised caching scheme by the content of its `scheme-name` element.
   c. For each `cache-mapping` element (which typically appears towards the top of this file), identify if it has a `scheme-name` element whose content matches one noted in the previous step and if so, make a note of its associated `cache-name` element.
3. In the `ehcache.xml` file:
   a. Identify each `cache` element whose 'name' parameter matches the `cache-name` elements noted in step '2c'.
   b. Using the mappings table below, apply the values noted in step '2a' to the appropriate parameters of the `cache` elements identified in the previous step ('3a').

Mappings table showing how elements of the Coherence cache configuration file map to parameters of the equivalent Ehcache file.

Coherence Element	Ehcache Attribute
`high-units`	`maxElementsInMemory`
`expiry-delay` > 0s	`timeToIdleSeconds` - Use this attribute for expiry delays greater than 0s along with the `eternal` attribute set to 'false'
`expiry-delay` = 0s	`eternal` - For expiry delays of 0s, set this attribute to 'true'.

Upgrade Procedure

If you are already running a version of Confluence, please follow these instructions to upgrade to the latest version:

1. Before you upgrade, we strongly recommend that you **back up your `confluence.home` directory and database**.
2. If your version of Confluence is earlier than 3.0.1, read the release notes and upgrade guides for **all releases** between your version and the latest version. In particular:
   - Please read the 3.0.1 release notes.
   - If you are upgrading from 2.1 or earlier, please also read the 2.2 release notes.
3. Download the latest version of Confluence.
4. Follow the instructions in the Upgrade Guide.

Confluence 3.0 Release Notes

Confluence 3.0 fixes some security flaws. Please refer to the security advisory for details of the security vulnerabilities, risk assessment and mitigation strategies.

1 June 2009

With great pleasure, Atlassian presents Confluence 3.0.

Confluence 3.0 is a major release which presents a number of new features and enhancements. With Confluence 3.0, we introduce the new **Macro Browser** feature, which provides a simple, point-and-click interface for discovering and inserting any of Confluence's 80+ bundled macros into a page. Macros greatly enhance the functionality of Confluence pages by allowing users to add tools like task lists, RSS feeds and multi-media content. The macro browser makes it easy to select the macro you need and preview its output with your chosen settings before finally adding it to your page. It also makes it easy to edit and modify the parameters of existing macros on a page.

Confluence 3.0 also provides a range of social features that allow you to discover and connect with other users in your Confluence site. **User Profiles** have been redesigned with a more intuitive layout and provide additional fields for users to share information about themselves like their IM handle or personal website. The new **Network** feature lets you follow other users inside your Confluence site and view an aggregated stream of all the activity such as page edits, blogs posts and status updates they undertake. The new **Hover Profile** feature displays a user's summary profile information whenever you hover your mouse pointer over their name anywhere in Confluence. Hover profiles provide easy access to the user's email address and profile details and allow you to add them to your network. Finally, you can let others know what you are working on, share a link or simply broadcast a short message to your team by setting your **User Status**. Once your current status message has been updated, it will appear in various activity streams, your profile views and your profile sidebar.

The new **PDF Export** feature addresses one of the most highly voted requests for Confluence. You have full control over your PDF exports using Cascading Style Sheet (CSS) modifications. You also have the ability to customise page and margin sizes, add a title page or table of contents list to your exported document, add customisable headers and footers and easily select a subsection of documentation for export.

The **Rich Text Editor** now includes a **right-click context menu** that makes it easier to insert links, images, macros and tables into your Confluence page. It is now also possible to copy/cut and paste rows within tables.

Confluence's **performance** has been substantially improved in both standalone and clustered installations with dramatic improvements in response times. Furthermore, **scalability** has been enhanced, with significantly improved CPU utilisation compared to previous Confluence versions.

Highlights of this Release:

- Introducing the Macro Browser
- Enhanced User Profiles
- Introducing Your Network
- New User Status
- New Hover Profile Feature
- Customisable Enhanced PDF Exports
- Improved Rich Text Editor
- Performance Improvements
- Engine Room and Developer Community
- Administration Improvements
- More than 240 Fixes and Improvements

Responding to your Feedback:

680+ votes satisfied



- Thank you for all your issues and votes. Keep logging, to help us keep improving!
- Below is a list of the highlights in this release.
- Attached is the full list of issues resolved in this release.

Upgrading from a previous version of Confluence

- Upgrading Confluence should be fairly straightforward. ***We strongly recommend that you back up your Confluence Home directory and your database before upgrading.***
- Please refer to the Confluence 3.0 Upgrade Notes for further essential information about **plugins** and other factors affecting your upgrade.

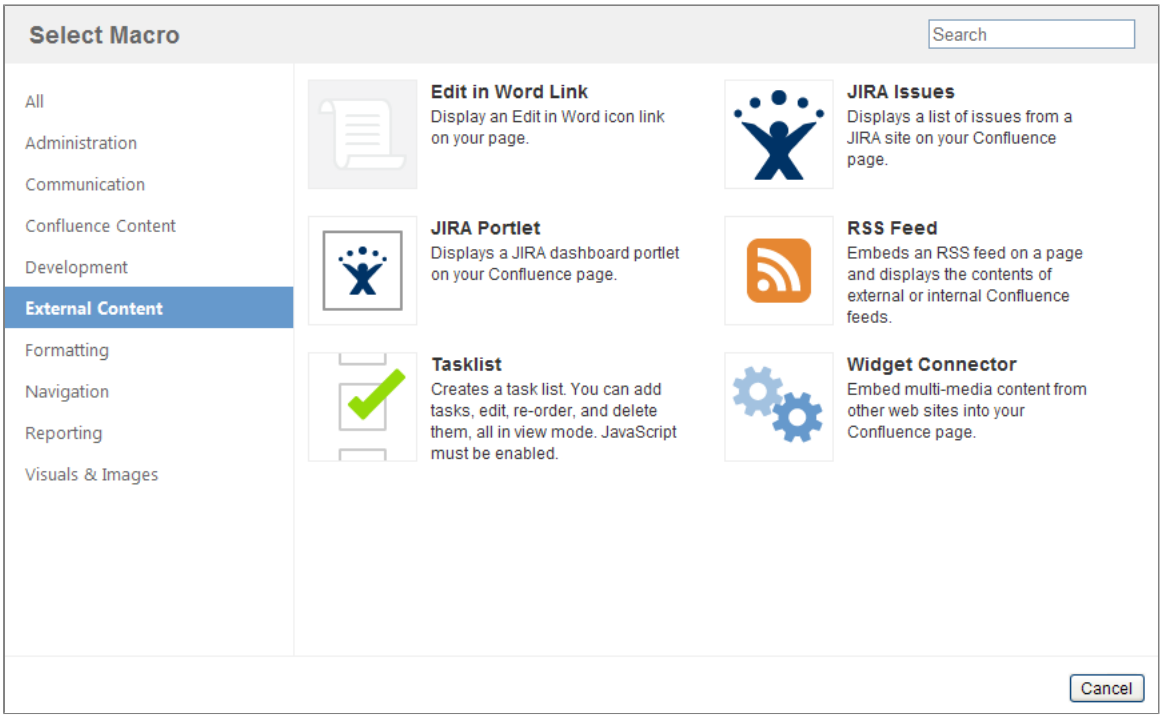
Highlights of Confluence 3.0



Introducing the Macro Browser

Now you can choose from Confluence's plugin-based macros and implement them with ease, from a single point.

- Using the macro browser's fast filtering capabilities, you can quickly find any bundled macro, including additional plugin-based macros installed on your Confluence system.
- Select your macro and modify and preview its parameters before adding it to your Confluence page or blog post.
- Take a look at our documentation for more details on the new macro browser feature.



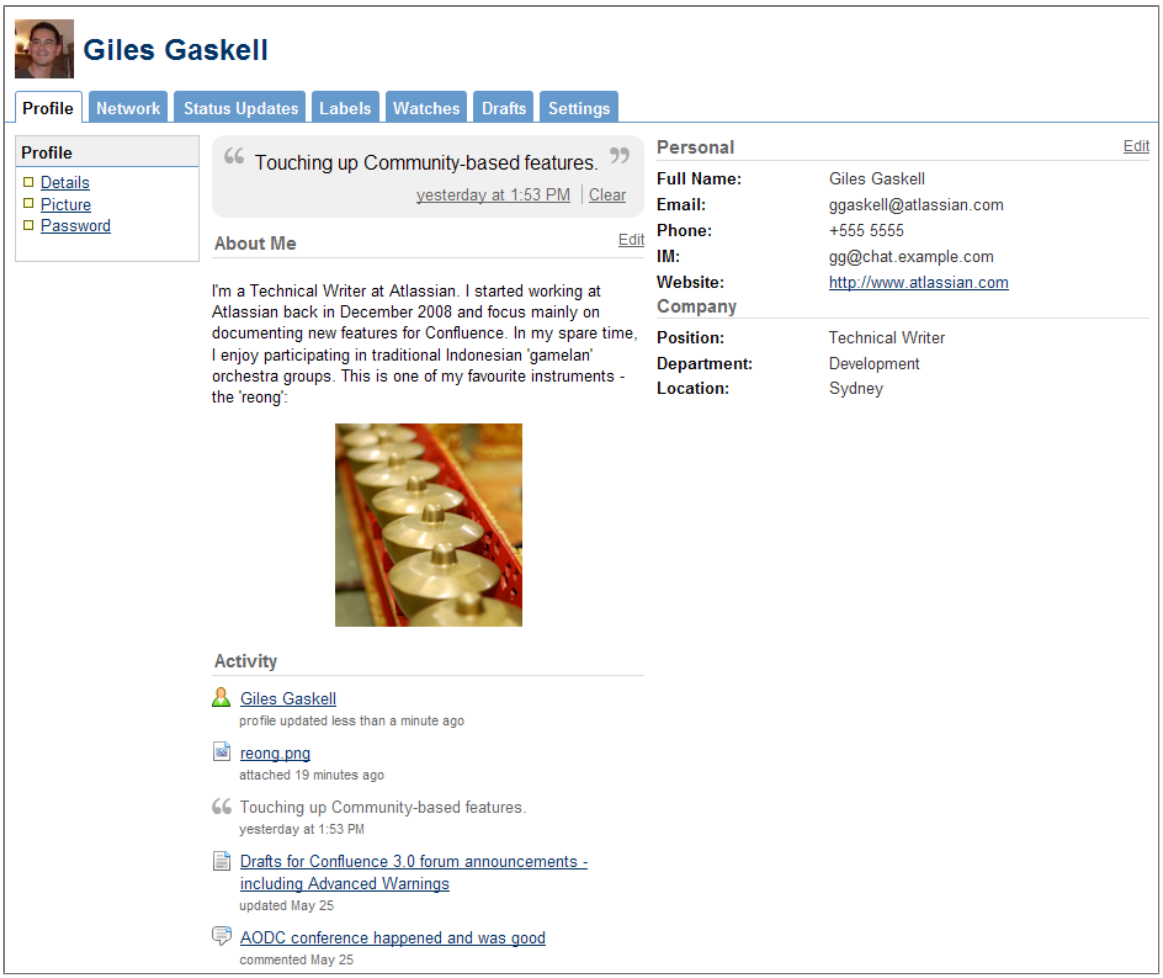


Enhanced User Profiles

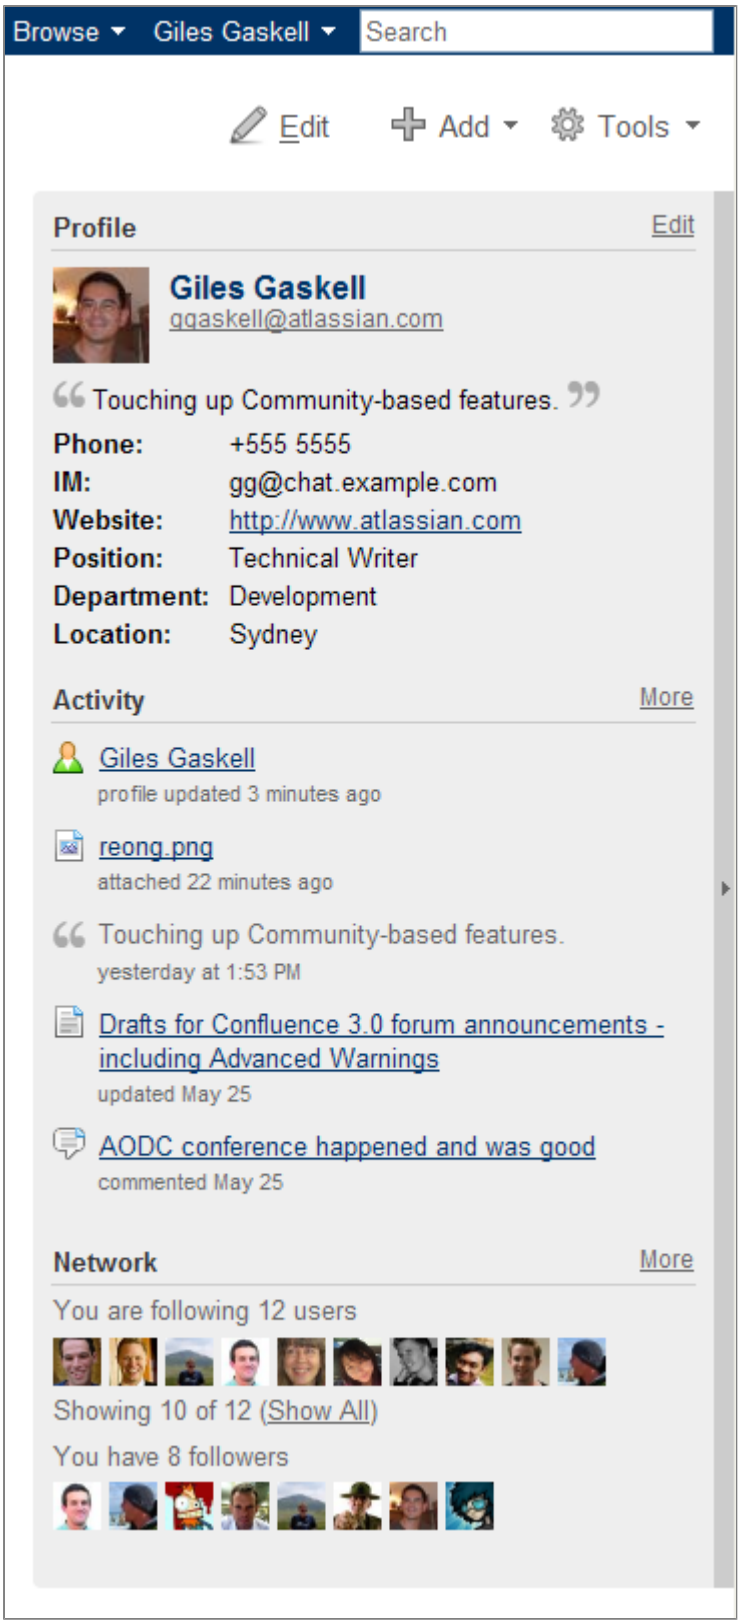
Confluence's enhanced user profiles area has been augmented to incorporate Confluence 3.0's new community-based features and to improve the overall user experience.

- User profiles now allow users to enter 'structured' information about themselves, which can be used by Confluence's community-based features.
- Each user's profile view shows a list of their own recent activities, such as page or blog post updates, changes to their profile information and status updates (described below).
- A summary of each user's profile information is displayed in a 'profile sidebar' on the right-hand side of pages within their personal space and their blog posts too. The profile sidebar appears on pages based on the default Confluence theme.
- Take a look at our documentation for more details on the enhanced user profiles feature.

The new user profile view



The new profile sidebar which appears on your blog posts or pages within your personal space



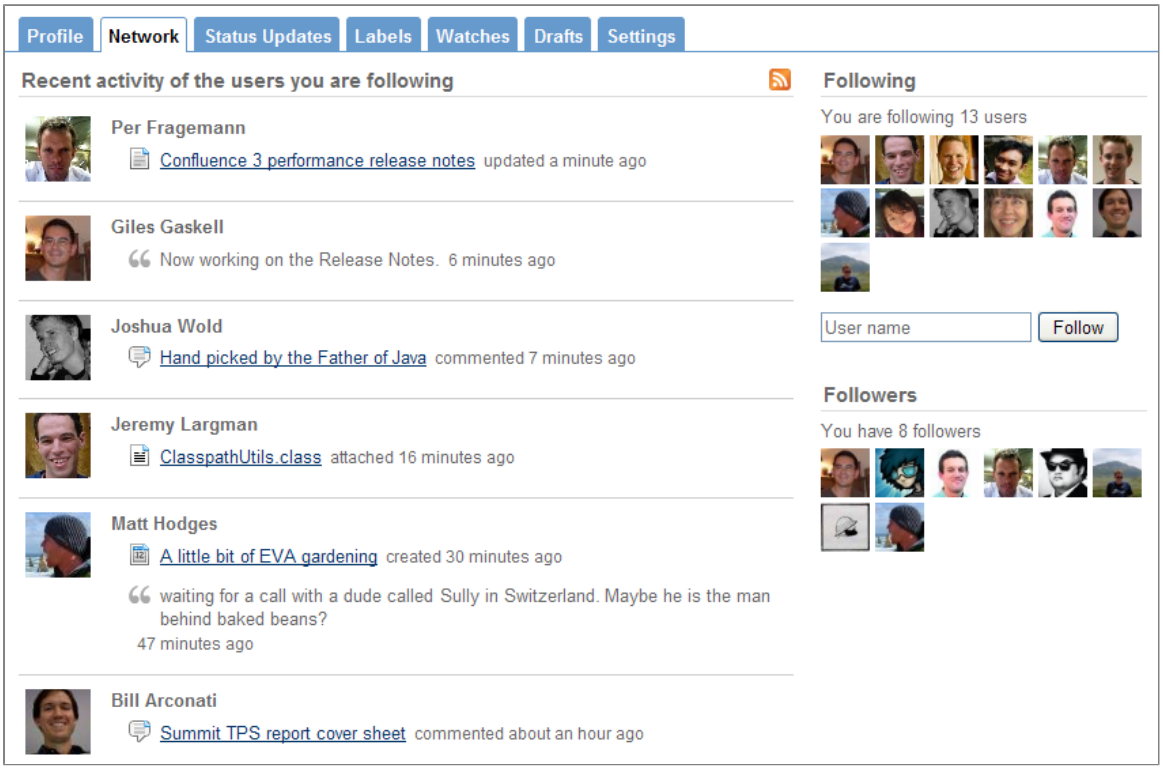
3

Introducing Your Network

The new network feature helps you keep track of what other users are doing throughout your Confluence site, by allowing you to 'follow' their recent activities.

- The activities tracked by the network feature include:
  - Additions or edits to pages or blog posts, including comments of users you are following.
  - Updates to statuses or profile details of users you are following.

- From the network view, you can set up an RSS feed which provides notifications on the activities of users you are following.
- Take a look at our documentation for more details on the network feature.

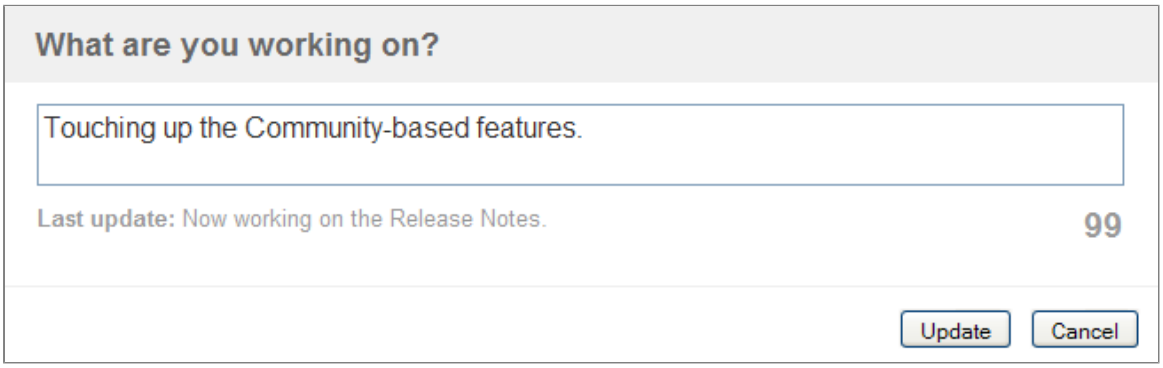


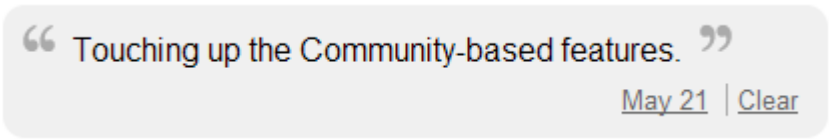


New User Status

User status allows you to broadcast a short message of up to 140 characters rapidly for others to see.

- Your messages could include anything from what you are currently working on to a message or a hyperlink you want to share immediately with other users.
- Other users can see your status messages on various activity streams throughout Confluence and on your profile views and personal space pages.
- Take a look at our documentation for more details on the new user status feature and your status updates page.







New Hover Profile Feature

Hover profile is a convenient tool that provides quick access to key information about other Confluence users, their user profile features and their network functions. Whenever you hover your mouse pointer over a Confluence user's name, key details about them appears in a popup balloon, such as their name, profile picture, email address and their current status.

- From a user's hover profile popup balloon, you can access the following functions:
  - Follow the user to track their recent activities via your network (or stop following them).
  - Directly send the user an email message via your email client.
- You can also access the following features of their user profile via their hover profile popup:
  - The user's personal space.
  - The user's profile, network or status updates views.
- Take a look at our documentation for more details on the new user hover profile feature.



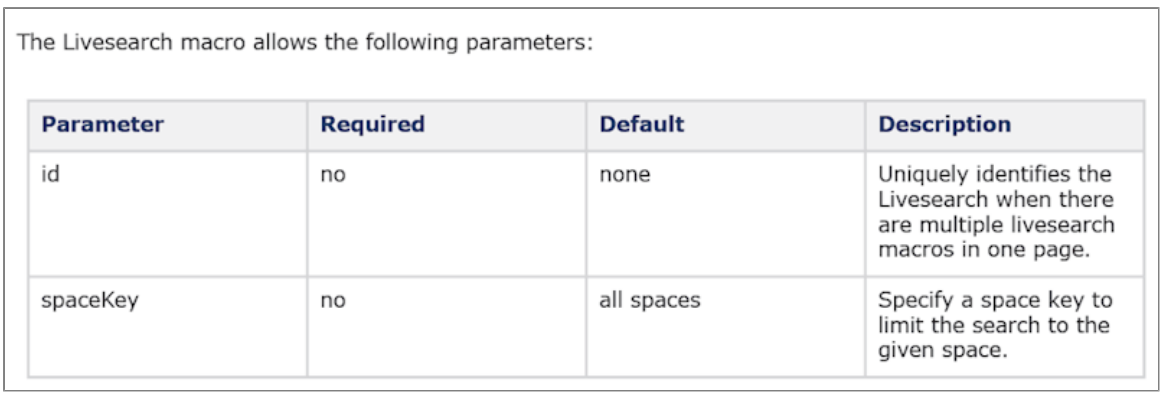


Customisable Enhanced PDF Exports

The enhanced PDF export feature in Confluence 3.0 has been rebuilt from the ground up and provides full customisation of your PDF exports with CSS modifications, to suit your particular requirements.

- Key enhancements to PDF export customisations include the ability to:
  - Customise page and margin sizes.
  - Add a table of contents or add headers and footers with customisable content.
  - Add a title page to your document.
  - Select a subsection of a space (for example, a chapter or section) to export more easily.
- PDF exports are now up to four times faster on large spaces.
- This feature addresses some of the most highly voted Confluence issues. Important bugs in the old PDF export feature have been fixed in this new version, including:
  - Fixed width columns — Table columns were consistently presented with fixed widths using the old PDF export function. However, the new PDF export function presents table columns with variable widths, as they appear on screen.
  - Ability to handle landscape page exports.
- Take a look at our documentation for more details on the enhanced PDF export and PDF stylesheet features.

The old PDF export function only generated fixed-width table columns


The Livesearch macro allows the following parameters:

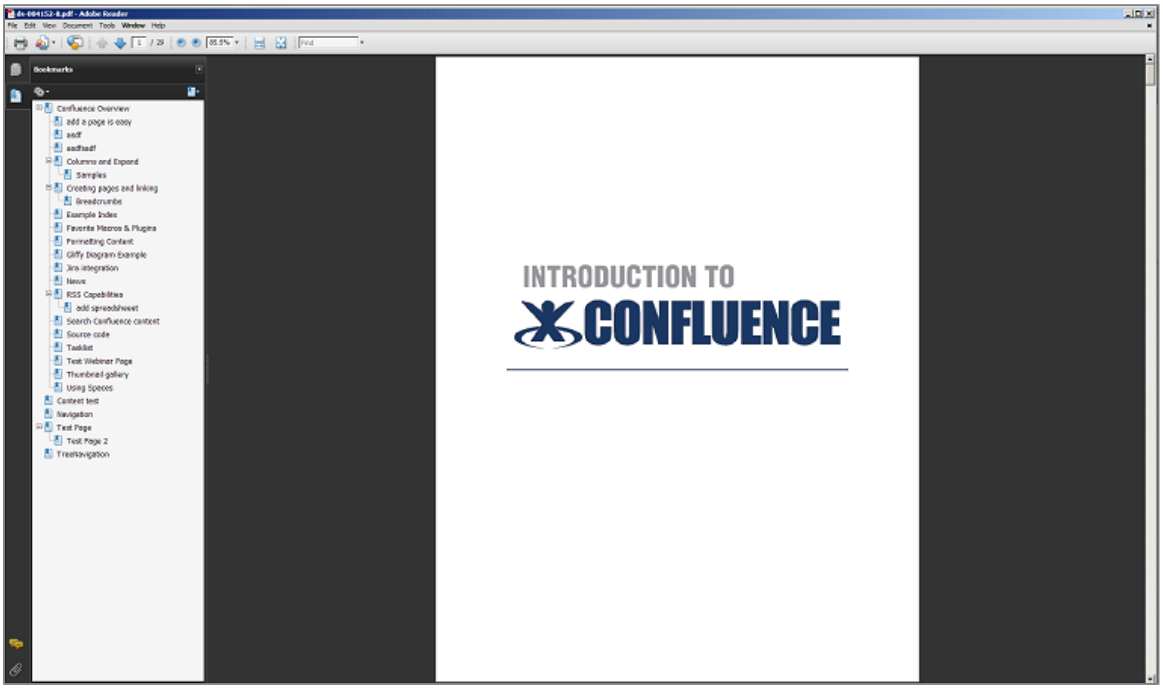
Parameter	Required	Default	Description
id	no	none	Uniquely identifies the Livesearch when there are multiple livesearch macros in one page.
spaceKey	no	all spaces	Specify a space key to limit the search to the given space.

The new PDF export feature now generates variable-width table columns

The Livesearch macro allows the following parameters:

Parameter	Required	Default	Description
id	no	none	Uniquely identifies the Livesearch when there are multiple livesearch macros in one page.
spaceKey	no	all spaces	Specify a space key to limit the search to the given space.

Easily add a title page when you export a space as a PDF



A table of contents is generated in the PDF by default when you export a space

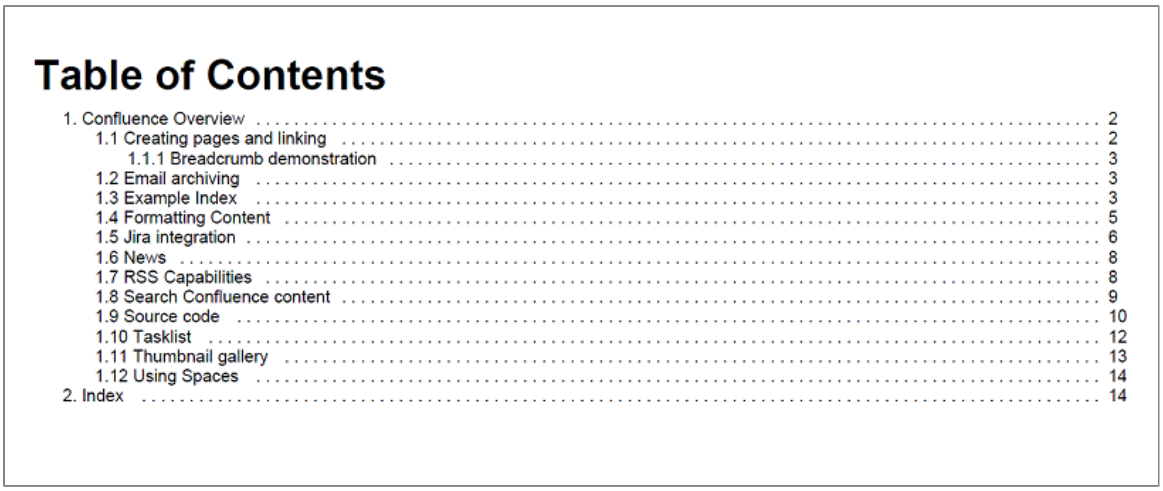
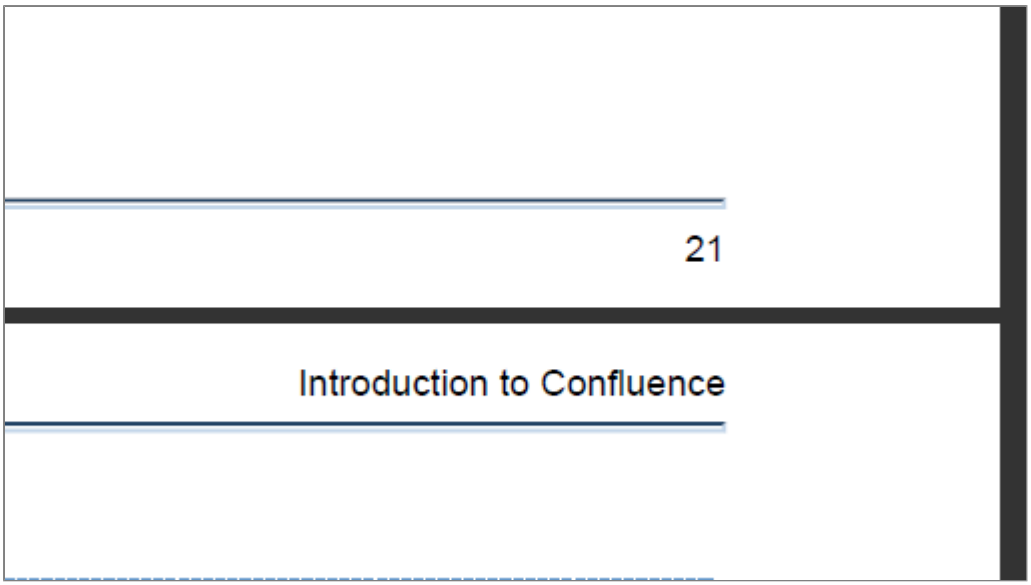


Table of Contents

1. Confluence Overview . . . 2
1.1 Creating pages and linking . . . 2
1.1.1 Breadcrumb demonstration . . . 3
1.2 Email archiving . . . 3
1.3 Example Index . . . 3
1.4 Formatting Content . . . 5
1.5 Jira integration . . . 6
1.6 News . . . 8
1.7 RSS Capabilities . . . 8
1.8 Search Confluence content . . . 9
1.9 Source code . . . 10
1.10 Tasklist . . . 12
1.11 Thumbnail gallery . . . 13
1.12 Using Spaces . . . 14
2. Index . . . 14

It's easy to add custom headers and footers



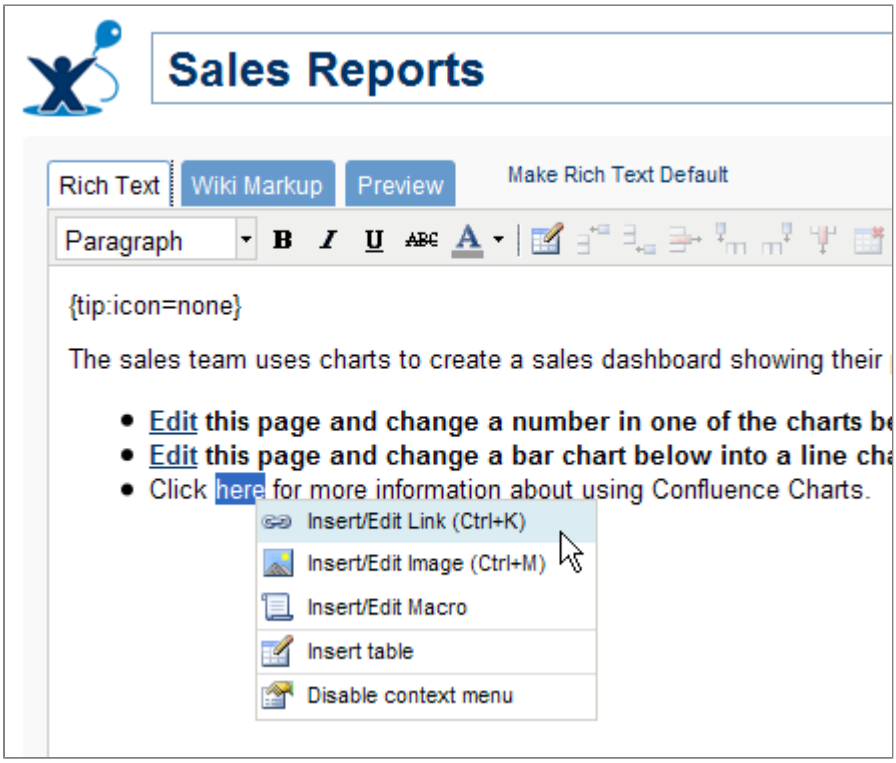


Improved Rich Text Editor

The rich text editor contains several interface enhancements and bug fixes that improve its overall user experience.

- The rich text editor includes a right-click context menu, making it easier to insert links, images, macros and tables into your Confluence page or blog post.
- It is now possible to cut/copy and paste rows.
- Take a look at our documentation for more details on the rich text editor overview to see more information about its enhancements and redesigned interface.

Right-click context menu



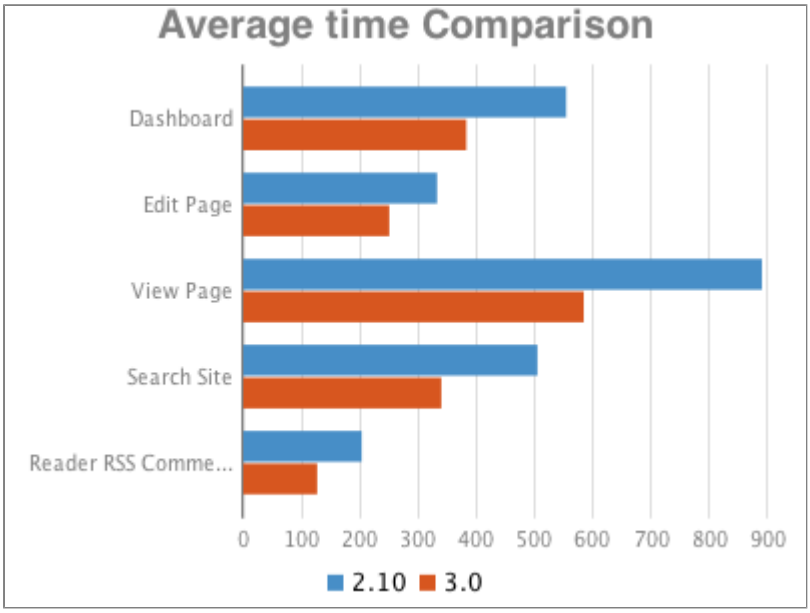


Performance Improvements

A number of performance improvements have been implemented.

- Confluence is now able to use more available resources under high load conditions.
- For Confluence Standalone installations, actions are up to 2 times faster.
- For Confluence Clustered installations:
  - Viewing pages is up to 2.5 times faster under medium, high and peak load.
  - Other actions are up to 5 times faster under medium load and 10 times faster under high and peak load.
- For more information, please refer to the Confluence 3.0 performance improvements documentation.

Confluence Standalone installations - Medium Load



9

Engine Room and Developer Community

- Confluence 3.0 now implements the Atlassian Plugin Framework 2.2 and the new Shared Access Layer (SAL).
- For more information, please refer to the Plugin Framework Documentation.

10

Administration Improvements

Confluence contains a number of improvements to its administrative features, many of which are easily accessible from the Administration Console.

- Security improvements:
  - A new form token authentication mechanism provides Confluence with the means to validate the origin and intent of requested actions, thus adding an additional level of security against phenomena such as cross-site request forgery. This feature also provides a mechanism for Confluence plugin developers to protect their plugins. For more information, please refer to the form token handling documentation.
  - Anti-XSS mode is now enabled by default and the remaining encoding bugs have been fixed.
- You can now generate a thread dump from the Administration Console. See our documentation on generating a thread dump.
- It is now possible to adjust the size of Confluence's internal caches, allowing administrators to fine tune Confluence's cache handling and performance at runtime without the need to restart Confluence. For more information on using this this feature, refer to our page on cache statistics.
- The Office connector contains additional configuration options and provides simplified handling of the {viewfile} macro for the new Macro Browser.

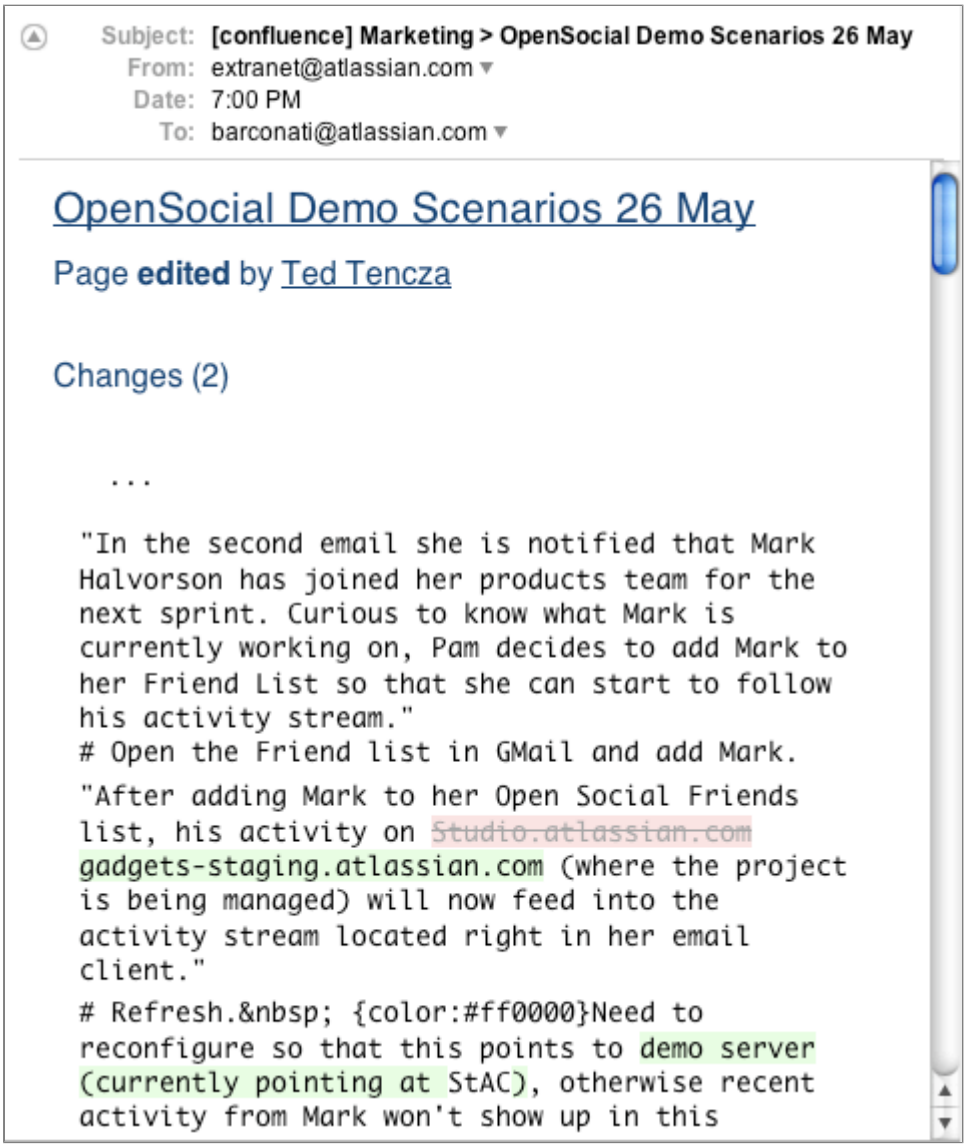
11

More than 240 Fixes and Improvements

Page comparisons have been improved such that deletions or additions of single words (and short phrases) within a single line are highlighted red or green, respectively. Furthermore, large sections of unchanged text are compacted to reduce page length, but their content and context can easily be revealed at a click.

- Page comparisons now appear in email notifications, whenever a user edits a page or blog post.
- As well as the enhanced PDF export feature, you can now choose a subsection of a space (such as a document chapter) more easily, to export to XML or HTML.
- Using the new network, profile and user status list macros, you can incorporate components of Confluence 3.0's new community-based features directly into your page or blog post.
- Atlassian now provides support for recently added features to the widget macro, which include Widgetbox, Yahoo Video, Dailymotion, Wufoo HTML Form Builder, DabbleDB, Google Calendar and BackType micro-blogging.
- The Activity Macros now incorporate improved handling of the 'author' parameter.
- Take a look at the complete list of issues resolved in Confluence 3.0.

Page comparisons in email notifications



Known Issues in this Release

We have an enthusiastic and dedicated group of testers and customers who jump in there, try out the new Confluence release, and report any problems so that we can fix them quickly.

We would like to highlight a known issue affecting the rich text editor, in which text cannot be added to the end of line that already ends with a link. Refer to CONF-15053 for more information about this issue. For more information on other known issues associated with the release of Confluence 3.0, please refer to Confluence 3.0 Known Issues.

A big **thank you** to everyone who helps us ensure that Confluence keeps getting better and better.

The Confluence 3.0 Team

Development

Bugfixing, Security, PDF Export, Office Connector
Andrew Lynch
Charles Miller
Paul Curren
Ryan Ackley

Engine Room and Performance
Matt Ryall
Chris Kiehl
Anatoli Kazatchkov

Macro Browser and Rich Text Editor
Agnes Ro
David Taylor
Dmitry Baranovskiy
Don Willis

Community-based features
Matthew Jensen
David Loeng
Brian Nguyen
Chris Broadfoot

Plugin Updates
David Chui

Team Lead
Per Fragemann

Support

Kuala Lumpur
Arie Murdianto
Azwandi Mohd Aris
Ming Giet Chong
Zed Yap

San Francisco
Jeremy Largman
Maleko Taylor
Tim Wong
Vincent Chang
Peter White

Sydney
Gurleen Anand
Ivan Benko
James Fleming
Michael Seager
Roy Hartono

Others

Design
Stephen Russell
Jason Taylor

Performance Engineering
George Barnett

Product Management
Adnan Chowdhury

Product Marketing Management
Bill Arconati

Quality Assurance
Mark Hrynczak
Peter de Zwart

Technical Writing
Giles Gaskell
Edwin Dawson

Confluence 3.0 Known Issues

We have an enthusiastic and dedicated group of testers and customers who jump in there, try out the new Confluence release, and report any problems so that we can fix them quickly. Below is a list of known issues. We're working on them, and will have a point release out as soon as possible.

A big **thank you** to everyone who helps us ensure that Confluence keeps getting better and better.

While you're waiting, take a look at the great new features in Confluence 3.0 Release Notes.

You can also browse the Confluence project in our issue tracker to see what's fixed and what's not, for each release.

Issues to be Fixed

Issue	Key
In the rich text editor, text cannot be added to the end of line that already ends with a link.	CONF-15053

Other Issues



Early 3.0 Build Issue

Note that an early build of 3.0 (Build number 1626) had a problem when moving a page from one space to another. If the default filesystem attachment storage was configured then any attachments on the page would fail to be moved if the page was moved to a new space. The attachment would only be accessible again if the page was moved back to its original space.

The current 3.0 release (3.0.0_01) does not have this problem. See CONF-15986 for more details.

Confluence on MySQL databases

Some customers who run Confluence on a MySQL database may find that when they upgrade to Confluence 3.0 or later, their Confluence 3.0 upgrade fails, with the Confluence logs revealing a "Specified key was too long" error. This issue is known to occur when MySQL's MyISAM storage engine and UTF-8 character set is used with Confluence.

If this is the case, please refer to the Upgrade to Confluence 3.0 with MySQL database fails with messages like "specified key was too long" knowledge base article on how to resolve this upgrade issue.

Confluence 3.0 Upgrade Notes

Below are some essential notes on upgrading to **Confluence 3.0**. For details of the new features and improvements in this release, please read the Confluence 3.0 Release Notes.

On this page:

- Upgrade Notes
  - Anonymous Access to User Profiles
  - Plugins
  - The Usage Tracking Plugin is disabled by default
- Upgrade Procedure

Upgrade Notes

Anonymous Access to User Profiles

Confluence 3.0 introduces a new Anonymous Access Global Permission, **View User Profiles**, which prevents individuals who have not logged in to Confluence from viewing other users profiles. This permission, which is accessible from the Administration Console was introduced to protect the identity and security of Confluence user accounts and details associated with the Confluence 3.0 Community-based features.

In all new installations of Confluence 3.0, this permission is disabled for Anonymous users by default, such that an individual will not be able to view another user's profile details until they have logged in to Confluence.

Plugins

If you are having trouble with any screens in Confluence 3.0, which could be related to potentially incompatible or unsupported third-party plugins, try using the Plugin Support Mode to disable all unsupported plugins.

The Usage Tracking Plugin is disabled by default

Due to performance issues in high load environments, the Confluence Usage Tracking Plugin is disabled by default in Confluence 3.0. While upgrading to Confluence 3, the plugin will be turned off even if it was active before. We do still deliver the plugin as part of Confluence (and have even improved its performance since 2.10), so if you don't have a high load environment and are happy with the performance of the plugin, you can easily reactivate this plugin through the Administration Console.

Upgrade Procedure



Upgrade a test environment first
As always, please test your upgrades in your test environment before rolling into production.

If you are already running a version of Confluence, please follow these instructions to upgrade to the latest version:

1. Before you upgrade, we strongly recommend that you back up your Confluence Home Directory and database. See the documentation on backing up your Confluence site. If you are using an external database, perform a database backup.

2. If your version of Confluence is earlier than 2.10.x, read the release notes and upgrade guides for all releases between your version and the latest version. In particular:
    - Please read the 2.10 upgrade notes.
    - If you are upgrading from 2.1 or earlier, please read the 2.2 release notes.
    - If you are upgrading from 2.2 or earlier, you will need to upgrade to Confluence 2.7.x first, confirm the upgrade was successful, then upgrade again from version 2.7.x to the latest. For more details, please refer to CONF-11767.

3. Download the latest version of Confluence.

4. Follow the instructions in the Upgrade Guide.

5. If you encounter a problem during the upgrade, please create a support ticket and one of our support engineers will assist you through the process.

RELATED TOPICS

Confluence 3.0 Release Notes

Confluence 3 Performance improvements

Confluence 3.0 has significant performance improvements over Confluence 2.10 and earlier versions. This page explains the performance characteristics of Confluence 3.0 and shows the improvements that were made when compared to its predecessor Confluence 2.10. In brief, compared to version 2.10, Confluence 3.0 response times in Standalone mode are down by 30% to 40%, and response times in a cluster are down by 50%. In other words, the clustered version of Confluence 3.0 is now twice as fast as before. Confluence also scales a lot better than before: More or faster CPU's are better utilised with Confluence 3.0 than they were with 2.10 and earlier versions.

1 Specific performance improvements

We have fixed a few bottlenecks that were so specific that you might have encountered them while analysing logfiles and thread dumps. Even if you did not see them, they might have been slowing your system down, depending on your use-case.

- **Rebuilding the search index is significantly faster**, up to factor 2. In our performance testing, a sample set of 20,000 pages that took 30 minutes using Confluence 2.10 now just takes 16 minutes in Confluence 3.0.
- **Improved Database Queries**
    - CONF-14488 : Added composite database index to SpacePermissions table. This will speed up installations with many page or space restrictions
    - CONF-14422 : Now, only the most recent version of an attachment is loaded when retrieving it for the first time. This had slowed down pages (and the dashboard) when an attachment had hundreds or thousands of revisions
    - CONF-14273 : Reduced overall DB load when rendering pages, which can help overall performance in case the database server was under high load already
- **Caching Enhancements**
    - CONF-12894 : Improved resource caching to improve HTTPS/SSL speed. This will make the screen render faster when you are using HTTPS.
    - CONF-8034 : Now serving caching headers for attachments to improve user interface responsiveness: Attachments (this includes user avatars) will not be downloaded again by the browser, leading to faster page loads
- **Others:**
    - Viewing PDF files through the office connector will use less memory and therefore significantly reduce garbage collection, which has caused some systems to perform a lot slower than needed
    - USGTRK-37 A bug was fixed in the usage statistics plugin that enables it to work smarter. It is still a plugin that is not made for high load so we suggest you disable it on high load scenarios

2 General performance improvements

We have been improving Confluence response time and scalability by implementing small improvements across the board. They are too many and too small to be presented individually, so we will use the results of our general performance-test to demonstrate the effect.

In our general Confluence performance tests, we execute a standardised set of commonly-used functions that simulates the activity of concurrent users. We base this profile on the actual usage patterns of our public Confluence installation, a rather large and active instance. To cater for irregular usage spike, we increase the load by factor 10. On average, this load test performs 10 to 15 Confluence requests per second. Most customer installations do not even get close to these numbers during normal operation. Under normal (low) load, the response times are actually a lot better than what we present here. But we prefer to use this medium load scenario because it simulates cases which *may* occur infrequently, and in which Confluence still needs to perform reasonably well. In addition to this scenario, we defined two additional, more extreme scenarios that perform the same requests, but at 20 to 35 requests per second to simulate an even higher load.



How to read and understand the statistics

Please note that we use the term "**request**" for anything that requests or posts data to Confluence. So viewing a Confluence page is a request, performing a search is a request, posting a comment is a request, and also using the quick navigation drop-down performs requests.

The data table

Each row in the table represents one use case. All use cases are run in parallel for 30 minutes, with a 5 minute ramp-up period.

- **Samplers**: The first column is the name of the requests performed in this scenario, like reading pages, commenting pages, or performing searches.
- **95% Percentile:** This is the time (in milliseconds) by which 95% of *all* requests of this scenario have completed. This is *not* an average value, you rather can think of it as a "how long the slowest requests (except the very worst 5% cases) take" - scenario.
- **Average**: the third column shows the average response time of the requests in this scenario - the lower the better.

The most important use-cases are the following:

- **View Page:** This loads one out of hundreds of different Confluence pages. Some are short, others are long. Some have many images, others have many comments. Some have many macros, others do not. The pages are accessed through their full URL, as if someone had clicked a link within the application or a bookmark.
- **Search Site:** A search across the whole system.
- **Quick Nav:** This simulates typing a character into the search field and getting back suggestions in real time. This is one of the most popular and time-critical operations. Therefore, this operation needs to be very fast.
- **Dashboard**: Simulates visiting the Confluence dashboard.
- **Edit Page**: This saves a page back to Confluence, and notifies all people who are watching this page.

The graph

The chart shows how many concurrent requests per second are being processed. The blue line indicates the moving average per second, and they green lines indicate variation. The blue line is not constant, since the pages and operations requested are extremely different in their CPU usage: A short page with no comments will render faster than a long page with many macros and comments, which in turn, will render faster than a page-edit that triggers many notifications. These differences in requests result in different CPU loads over time.

The more stable the blue average line is, the more consistent the user experience. The higher the line is, the more users can access and use Confluence simultaneously.

Applying the numbers to your company's usage patterns

The notes on this page geared at showing the performance differences between 2.10 and 3.0, using the same tests we used to test Confluence 2.10.

Hardware specification

All tests were conducted on two to four servers, each of which had the following specifications:

Name	Value
Server model	Dell 2950
CPU type	Intel(R) Xeon(R) CPU E5405 @ 2.00GHz (4 Cores)
CPUs per server	1 or 2, depending on test. See test details
RAM	32Gb (but just 2Gb are used for the JVMs, and the database uses 3Gb)
Disks	2 x 15K, 72Gb SAS
Network	1Gbps
Webserver	Tomcat 6, Java 6
Database	Postgres 8.2.4

When testing the Confluence Standalone version, one server acts as the application server and one as the database server, which is the setup we recommend to customers to enable high performance. A third server is used to generate the load, using JMeter. In the cluster, we use two application servers and one database server. In the cluster configuration we use the Pound load balancer, which runs on the same (fourth) server as the load generator JMeter. We do not use any webserver or caching proxy for our tests, and we cannot make any recommendations about which one to use. We want to measure the raw performance of the application server and suggest that you use the webservers/proxies with which you are most familiar.

Software and Settings

The JVM settings we used were -XX:MaxPermSize=192m -Xmx2000m -XX:+PrintGCTimeStamps -verbosegc -XX:+PrintGCDetails -XX:+PrintTenuringDistribution -XX:NewSize=384m -XX:SurvivorRatio=2 -XX:+UseParallelGC -XX:+UseParallelOldGC.

The usage tracking plugin was disabled during these tests because it is known to have performance issues and we recommend that it be turned off in high load deployments.

Confluence standalone

Confluence is most frequently installed on one physical machine. Unless you know you are using (or are planning to use) a cluster, then this section is for you.

Confluence 3.0 Standalone has significantly better performance characteristics than Confluence 2.10 Standalone. We compare three load scenarios and give the details below.

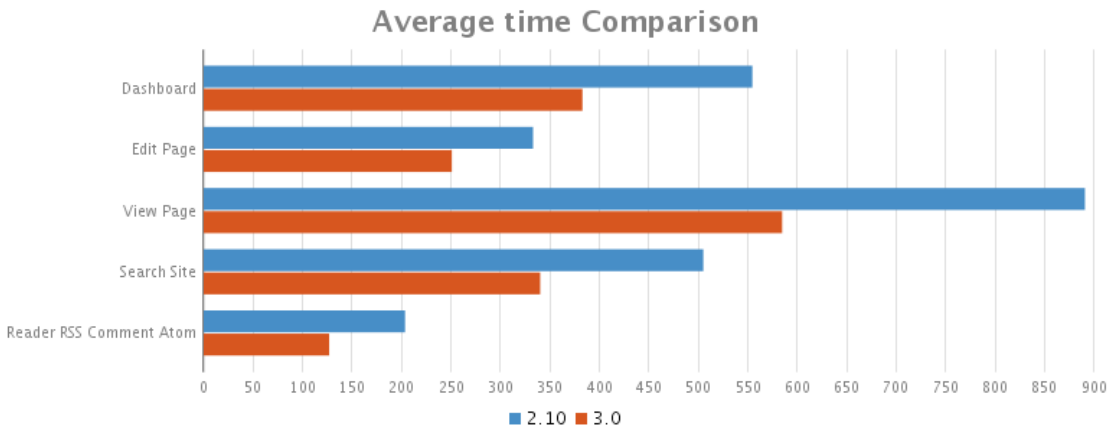
Medium Load scenario, Standalone, 1 CPU

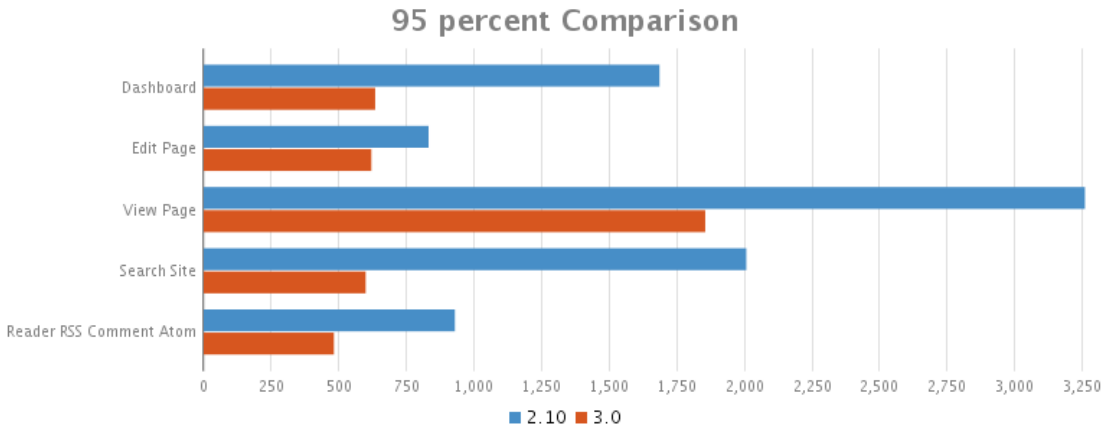
We define Medium Load as requesting roughly 15 requests per second from the loadtest. Most customers with smaller user bases never get even close to this usage, so they will experience a lot faster response times than what you can see below. But occasionally even customers with less than 1000 active users might experience spikes in usage, so we chose 15 requests per second as our medium load scenario.

We are using modest hardware (see above) with just one Xeon CPU with 4 cores, since we assume this is what a medium sized company would be using.

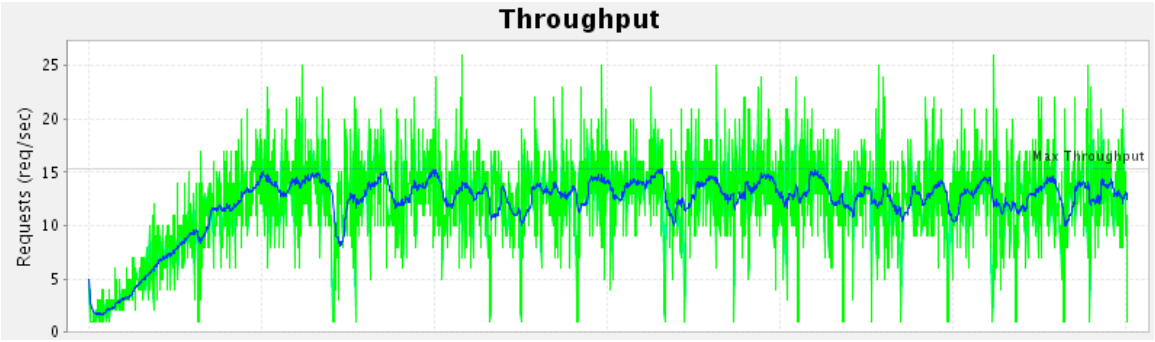
Confluence 2.10 vs Confluence 3.0 response times

Scenario	Average time in 2.10	Average time in 3.0	Improvement	95 percent in 2.10	95 percent in 3.0	Improvement
Browse Label	1619ms	1129ms	43%	3979ms	2387ms	66%
Commentor submit comment	306ms	338ms	-9%	805ms	794ms	1%
Commentor view commented page	737ms	628ms	17%	3386ms	1783ms	89%
Commentor view page	989ms	707ms	39%	4133ms	2168ms	90%
Creator add page	402ms	211ms	90%	765ms	391ms	95%
Creator submit new page	525ms	387ms	35%	1161ms	882ms	31%
Creator view page	256ms	233ms	9%	704ms	501ms	40%
Dashboard	554ms	382ms	44%	1685ms	634ms	165%
Editor display page	520ms	417ms	24%	1881ms	1065ms	76%
Edit Page	332ms	250ms	32%	831ms	620ms	34%
Editor submit edit	1199ms	961ms	24%	3949ms	3459ms	14%
Go to log in page	274ms	456ms	-39%	486ms	2795ms	-82%
Log In	342ms	333ms	2%	774ms	480ms	61%
Quick Navigation Search	134ms	57ms	133%	597ms	110ms	439%
Reader Not Found	551ms	369ms	49%	1266ms	615ms	105%
Reader RSS Blogpost Atom	170ms	59ms	184%	637ms	67ms	838%
Reader RSS Blogpost RSS2	206ms	95ms	116%	754ms	97ms	675%
Reader RSS Comment Atom	203ms	126ms	60%	929ms	481ms	92%
Reader RSS Comment RSS2	369ms	151ms	143%	1602ms	477ms	235%
Reader RSS Page Atom	628ms	513ms	22%	2725ms	2147ms	26%
Reader RSS Page RSS2	800ms	547ms	46%	3381ms	2196ms	53%
View Page	890ms	584ms	52%	3259ms	1854ms	75%
Reader for Space Page	904ms	677ms	33%	2219ms	1566ms	41%
Search Site	505ms	340ms	48%	2006ms	598ms	235%

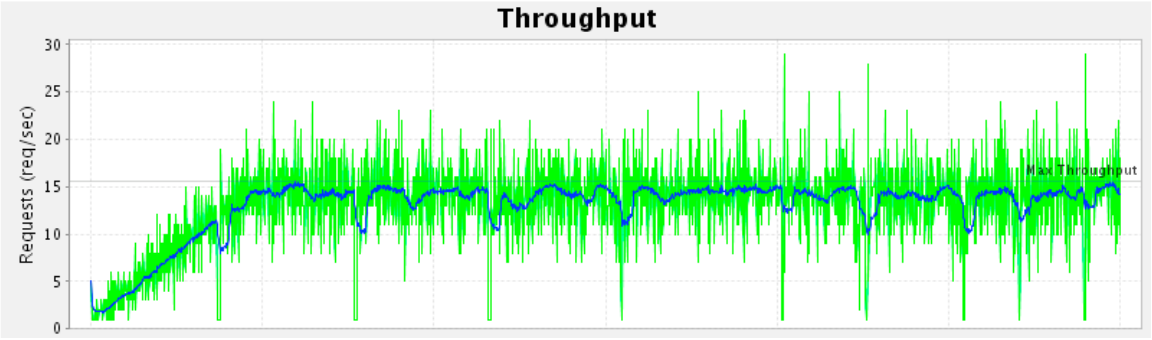




Confluence 2.10 throughput



Confuence 3.0 throughput



Medium Load comparison between 2.10 and 3.0 in standalone mode

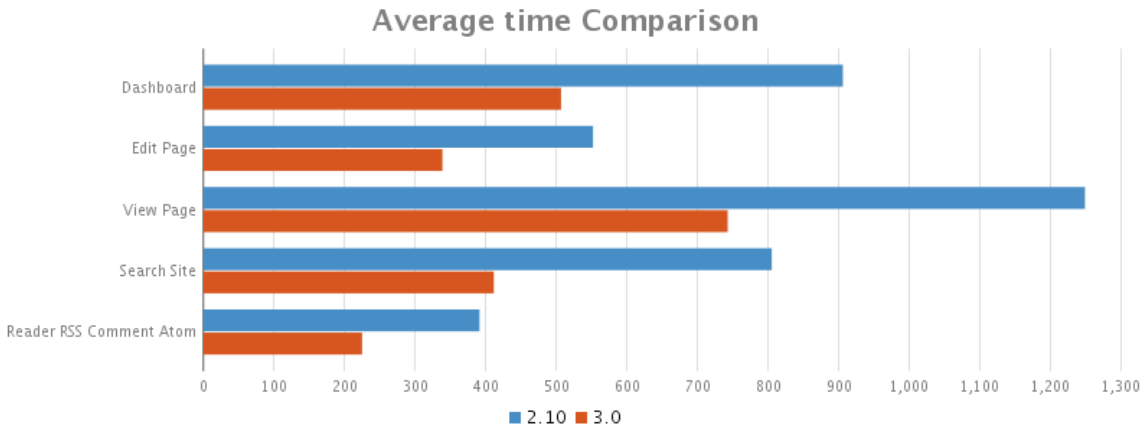
The most important scenario ("View Page") used to take about 900ms in Confluence 2.10, but in 3.0 it is down to 600ms, which is a performance improvement of about 50%. Almost all other scenarios have improved as well, some even by more than 100% (e.g. more than twice as fast). The throughput in this scenario has only changed from approximately 13/s to 14/s. However, this is because the test itself is not making more requests. The main improvement here is that the throughput has less variations (ups/downs) for example when rendering very complicated or large pages. You can improve the smoothness of the line even further by using a different garbage collector, as explained on our tuning page.

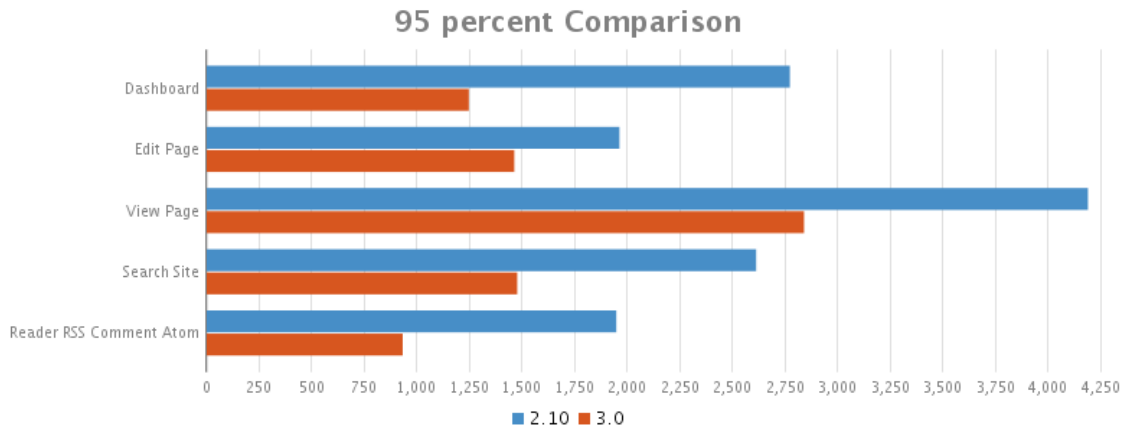
High Load Scenario, Standalone, 2 CPUs

We define a High Load Scenario as one in which the load generation equates to approximately 25 requests per second. In this test, we are using the same hardware as above, but with 2 CPUs. We assume that any company which expects 20 or more requests per second, even if this occurs during a short time frame, will have greater hardware resources (of equivalent cost) than to what is used in this test.

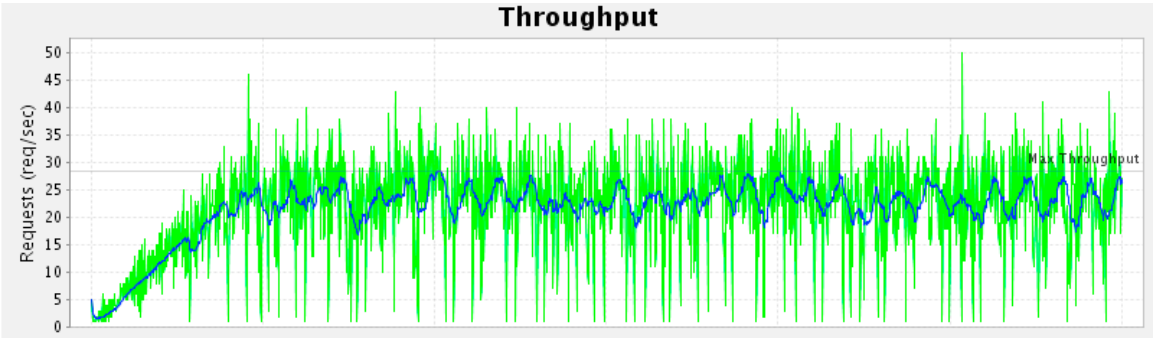
Confluence 2.10 vs Confluence 3.0 response times

Scenario	Average time in 2.10	Average time in 3.0	Improvement	95 percent in 2.10	95 percent in 3.0	Improvement
Browse Label	2389ms	1531ms	56%	6196ms	4195ms	47%
Commentor submit comment	424ms	397ms	6%	1779ms	1603ms	10%
Commentor view commented page	1211ms	815ms	48%	4729ms	2863ms	65%
Commentor view page	1402ms	912ms	53%	5284ms	3094ms	70%
Creator add page	558ms	297ms	87%	1962ms	1543ms	27%
Creator submit new page	783ms	522ms	49%	2567ms	1545ms	66%
Creator view page	443ms	284ms	55%	1845ms	843ms	118%
Dashboard	905ms	506ms	78%	2771ms	1245ms	122%
Editor display page	807ms	504ms	59%	2650ms	2165ms	22%
Edit Page	551ms	338ms	63%	1961ms	1461ms	34%
Editor submit edit	1524ms	1180ms	29%	5115ms	4189ms	22%
Go to log in page	409ms	419ms	-2%	1171ms	982ms	19%
Log In	520ms	346ms	50%	2124ms	700ms	203%
Quick Navigation Search	318ms	124ms	155%	1895ms	369ms	413%
Reader Not Found	866ms	492ms	76%	2439ms	1579ms	54%
Reader RSS Blogpost Atom	300ms	105ms	186%	1549ms	191ms	709%
Reader RSS Blogpost RSS2	299ms	98ms	203%	1954ms	183ms	965%
Reader RSS Comment Atom	390ms	224ms	74%	1946ms	931ms	108%
Reader RSS Comment RSS2	362ms	254ms	42%	1824ms	1196ms	52%
Reader RSS Page Atom	1098ms	777ms	41%	4804ms	2848ms	68%
Reader RSS Page RSS2	1126ms	807ms	39%	4532ms	3406ms	33%
View Page	1248ms	742ms	68%	4188ms	2839ms	47%
Reader for Space Page	1410ms	914ms	54%	3749ms	2487ms	50%
Search Site	804ms	411ms	95%	2611ms	1475ms	76%

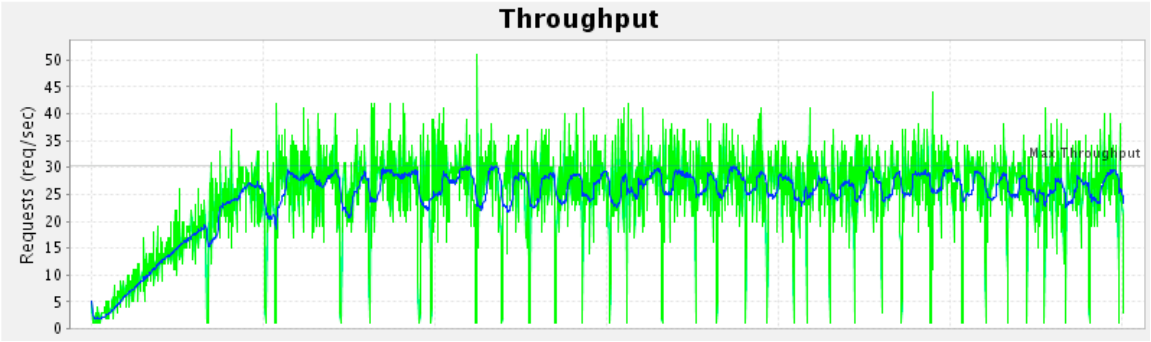




Confluence 2.10 throughput



Confluence 3.0 throughput



High Load comparison between 2.10 and 3.0 in standalone mode

This scenario shows the performance improvements between Confluence 2.10 and 3.0 best. Confluence 2.10 managed about 22 requests per second, Confluence 3.0 about 27 requests per second. Although this is a significant improvement, those in response times are even more impressive. If you have times when there are 20 requests per second, Confluence will respond a lot better and end users will notice the difference.

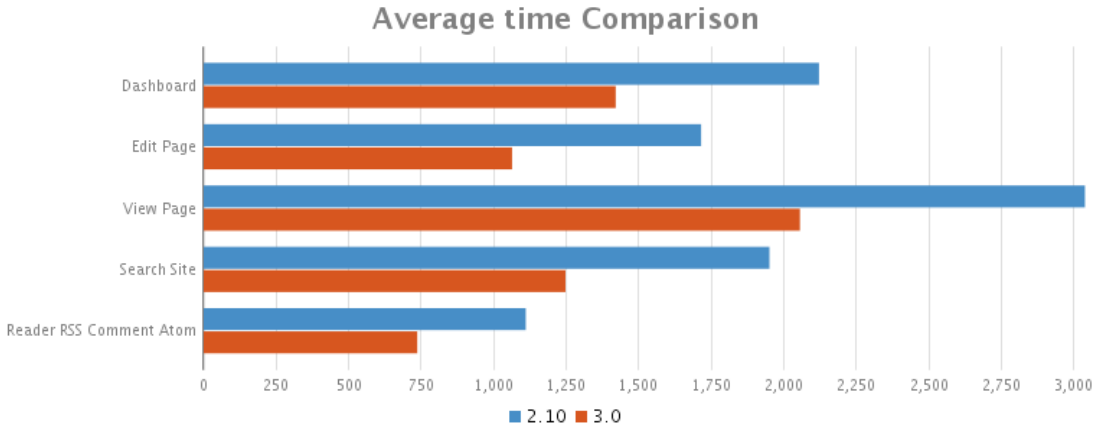
Peak Load Scenario, Standalone, 2 CPUs

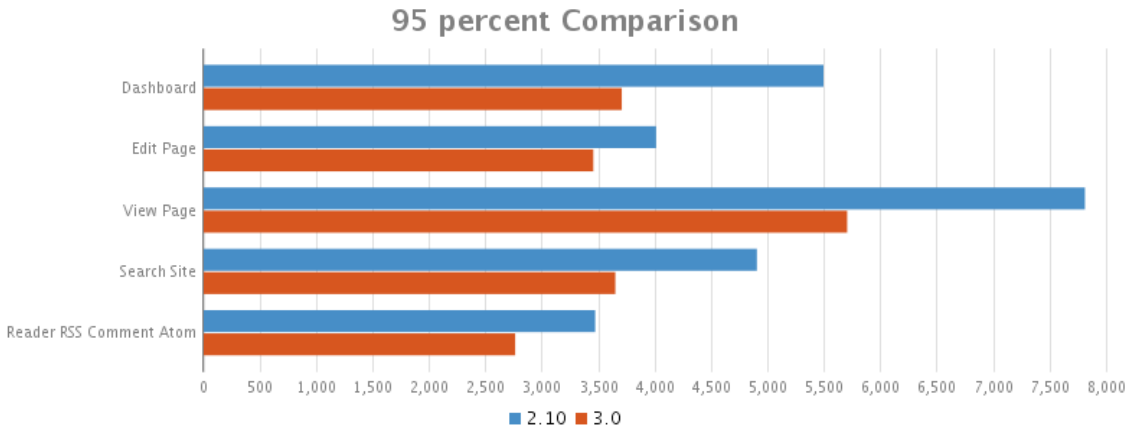
We define a Peak Load Scenario as one in which approximately 35 requests per second from the load generator. Very few of our customers ever reach these high levels of requests per second, but if you do have 100.000 users and many of them view pages at the same time, then the peak load scenario may occasionally be reached. Again, these tests are run on a 2CPU hardware.

Confluence 3.0 vs Confluence 2.10 response times

Scenario	Average time in 2.10	Average time in 3.0	Improvement	95 percent in 2.10	95 percent in 3.0	Improvement
Browse Label	4747ms	3207ms	47%	10951ms	7575ms	44%
Commentor submit comment	1517ms	1146ms	32%	4521ms	3611ms	25%
Commentor view commented page	3148ms	2173ms	44%	9222ms	6184ms	49%
Commentor view page	3302ms	2317ms	42%	9891ms	6410ms	54%
Creator add page	1693ms	934ms	81%	3904ms	3170ms	23%

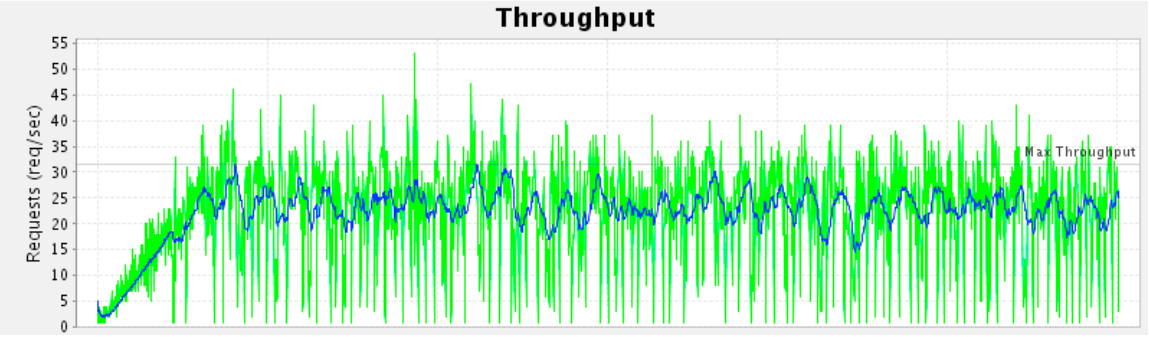
Creator submit new page	2777ms	1959ms	41%	5812ms	5170ms	12%
Creator view page	1589ms	1065ms	49%	3523ms	3358ms	4%
Dashboard	2121ms	1420ms	49%	5492ms	3704ms	48%
Editor display page	2216ms	1502ms	47%	5081ms	4233ms	20%
Edit Page	1714ms	1062ms	61%	4008ms	3452ms	16%
Editor submit edit	3945ms	3205ms	23%	10523ms	9467ms	11%
Go to log in page	934ms	818ms	14%	4544ms	4091ms	11%
Log In	807ms	913ms	-11%	2879ms	3531ms	-18%
Quick Navigation Search	1121ms	568ms	97%	4288ms	2704ms	58%
Reader Not Found	2159ms	1222ms	76%	4265ms	3472ms	22%
Reader RSS Blogpost Atom	864ms	531ms	62%	2796ms	2511ms	11%
Reader RSS Blogpost RSS2	1099ms	527ms	108%	4307ms	2691ms	60%
Reader RSS Comment Atom	1110ms	736ms	50%	3469ms	2760ms	25%
Reader RSS Comment RSS2	1159ms	863ms	34%	3959ms	3130ms	26%
Reader RSS Page Atom	2395ms	2300ms	4%	7588ms	7342ms	3%
Reader RSS Page RSS2	2661ms	2258ms	17%	8295ms	7922ms	4%
View Page	3038ms	2055ms	47%	7809ms	5702ms	36%
Reader for Space Page	3005ms	1990ms	51%	6298ms	4850ms	29%
Search Site	1950ms	1247ms	56%	4902ms	3647ms	34%



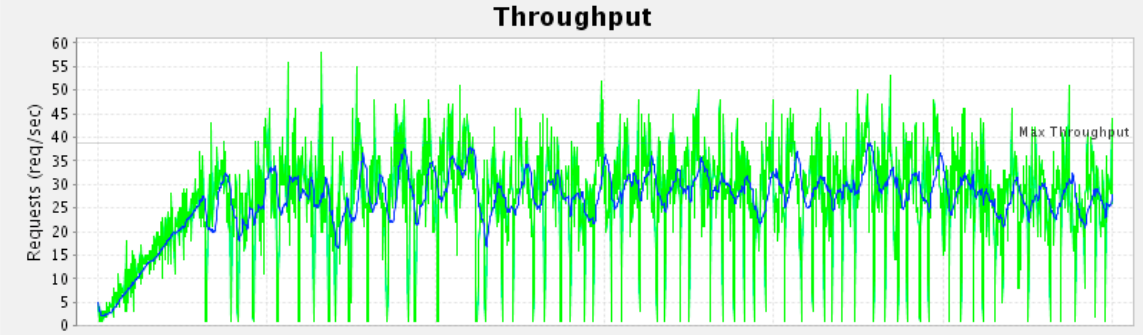


Confluence 2.10

throughput



Confluence 3.0 throughput:



Please note that this test is slightly skewed by the Load generator sitting one the same machine. The actual results will look a bit better.

Peak Load comparison between 2.10 and 3.0 in standalone mode

Confluence 2.10 is able to deliver about 22 requests per second, but response times are not so good. Rendering a page takes 3s and rendering the dashboard takes 2s on average. Confluence 3.0 delivers improved throughput of about 28 requests per second and response times are significantly better than 2.10 (rendering a page is down to 2s, and rendering the dashboard is down to 1.4s). However, response times under Peak Load in 3.0 are still not ideal. Even with 2 CPUs Confluence 3.0 starts reaching its limits here. While standalone is able to deliver results, what we really recommend for this peak load scenario is a clustered solution. Read on for more details.

Confluence Clustered

When rolling out Confluence to a larger amount of users, Clustering becomes important to balance spikes in load. The most commonly used deployment is a 2-node cluster, running on three physical machines (two application servers connected to one database server).

Clustering does not make a single request faster in low load scenarios, but it helps the system dealing with a larger number of requests in parallel, without degrading in performance.

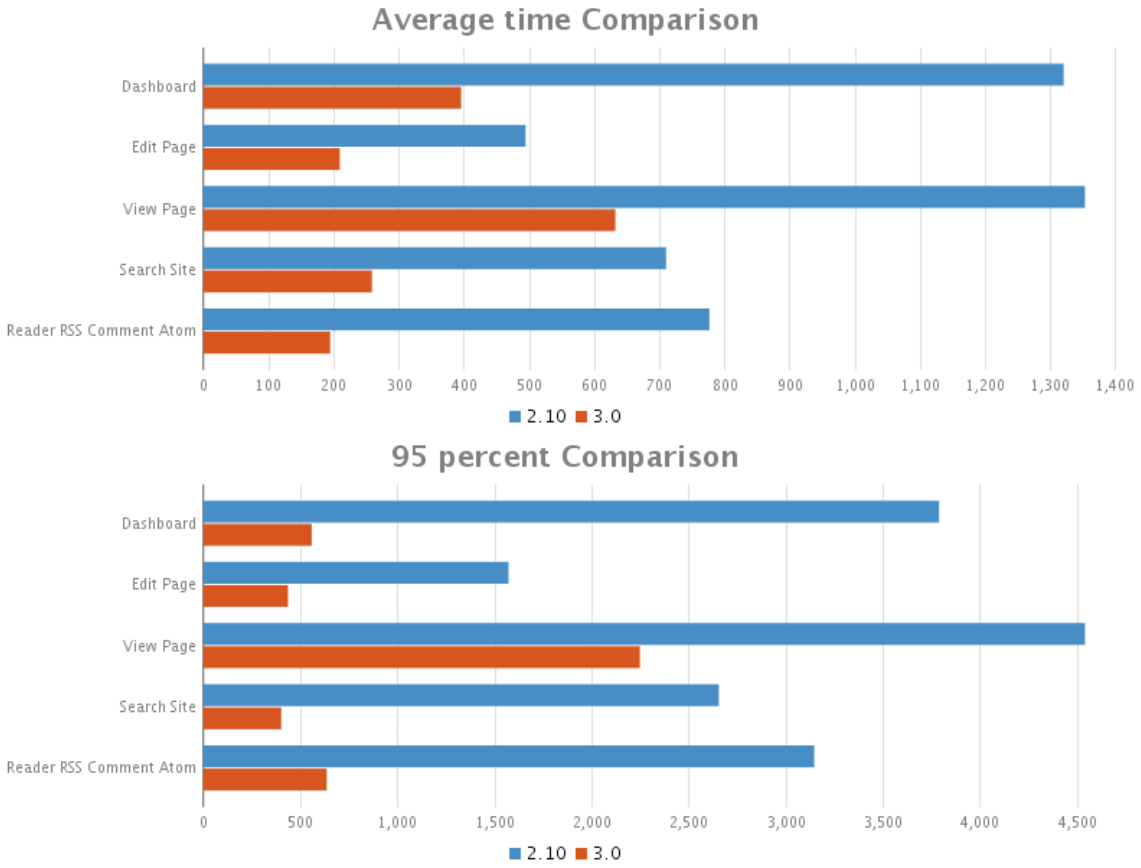
Medium Load Scenario, Clustered, 2 nodes, 1 CPU per node

As above, we define the Medium Load scenario as making 15 requests per second. This test uses just 1 CPU per machine.

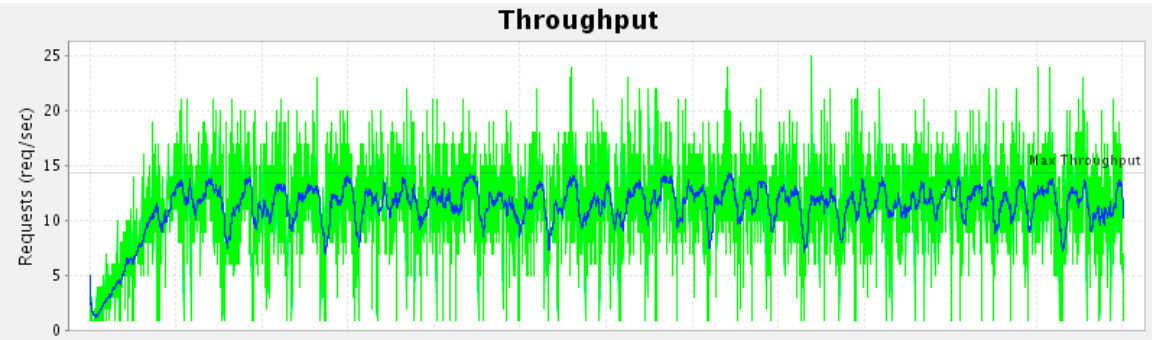
Confluence 3.0 vs Confluence 2.10 response times

Scenario	Average time in 2.10	Average time in 3.0	Improvement	95 percent in 2.10	95 percent in 3.0	Improvement
Browse Label	2477ms	919ms	169%	6365ms	2143ms	196%
Commentor submit comment	410ms	380ms	7%	1127ms	856ms	31%
Commentor view commented page	1029ms	595ms	72%	4193ms	1826ms	129%
Commentor view page	1367ms	786ms	73%	5264ms	2557ms	105%
Creator add page	611ms	214ms	184%	1463ms	414ms	253%
Creator submit new page	692ms	422ms	63%	1596ms	938ms	70%
Creator view page	329ms	131ms	150%	1034ms	205ms	404%
Dashboard	1319ms	395ms	234%	3787ms	556ms	581%
Editor display page	750ms	387ms	93%	2592ms	1420ms	82%
Edit Page	493ms	208ms	136%	1569ms	433ms	261%

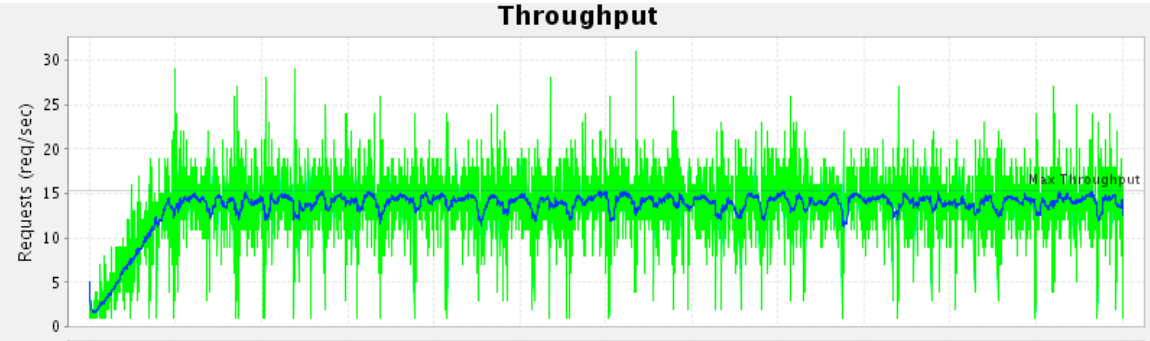
Editor submit edit	1500ms	980ms	53%	4903ms	3435ms	42%
Go to log in page	854ms	357ms	138%	2898ms	547ms	429%
Log In	961ms	422ms	127%	3320ms	597ms	455%
Quick Navigation Search	183ms	56ms	226%	1178ms	107ms	1001%
Reader Not Found	663ms	351ms	89%	1562ms	554ms	181%
Reader RSS Blogpost Atom	499ms	73ms	577%	2674ms	118ms	2166%
Reader RSS Blogpost RSS2	453ms	65ms	589%	1863ms	122ms	1420%
Reader RSS Comment Atom	776ms	194ms	300%	3143ms	633ms	396%
Reader RSS Comment RSS2	742ms	186ms	297%	2930ms	576ms	408%
Reader RSS Page Atom	1378ms	618ms	122%	4693ms	2109ms	122%
Reader RSS Page RSS2	1497ms	584ms	156%	4767ms	1899ms	150%
View Page	1352ms	631ms	114%	4538ms	2246ms	102%
Reader for Space Page	1251ms	793ms	57%	3124ms	1736ms	79%
Search Site	709ms	258ms	174%	2652ms	399ms	564%



Confluence 2.10 throughput



Confluence 3.0 throughput



Medium Load comparison between 2.10 and 3.0 in clustered mode

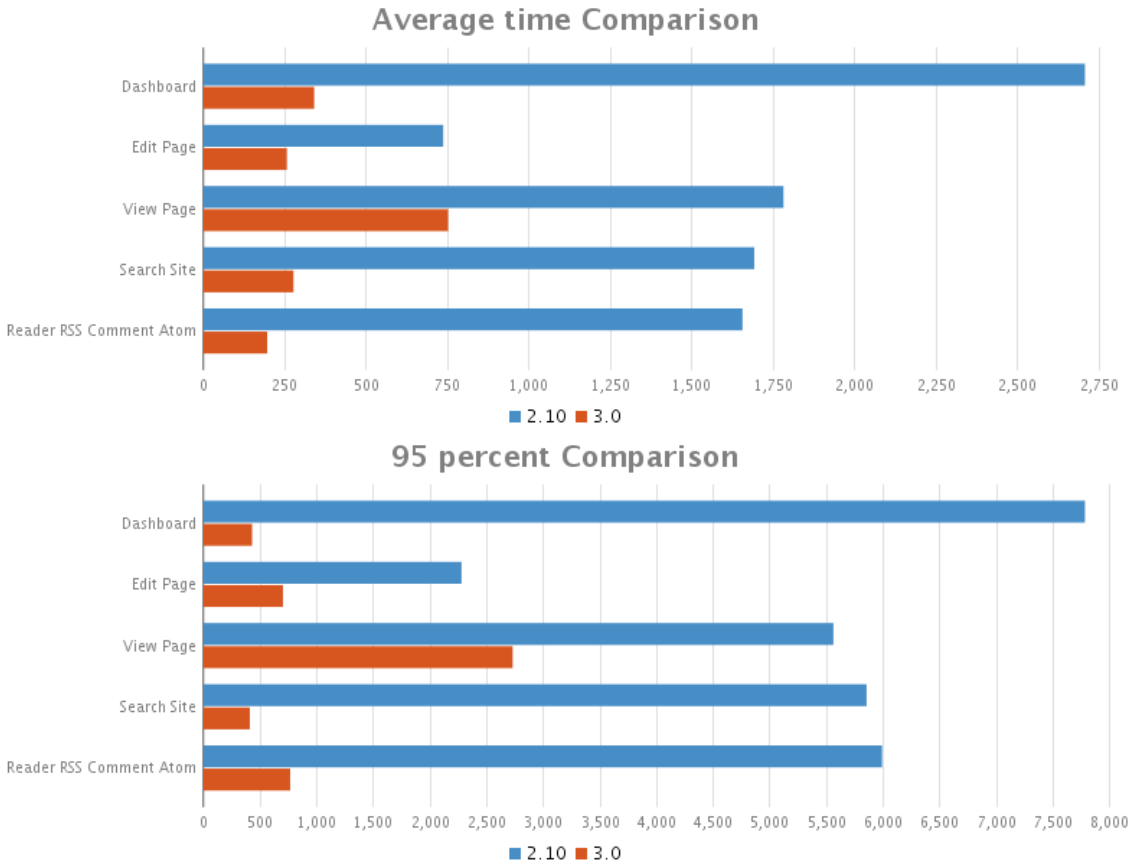
As you can see, the response time of each request is a lot better in Confluence 3.0. On average the performance has doubled, leading to response times that are just 50% of what they used to be. This means that a clustered installation provides the same responsiveness as a standalone installation, while still being much better at scaling, which will be shown below. In this example the load was so low that throughput did not increase very much.

High Load Scenario, Clustered, 2 nodes, 2 CPUs per node

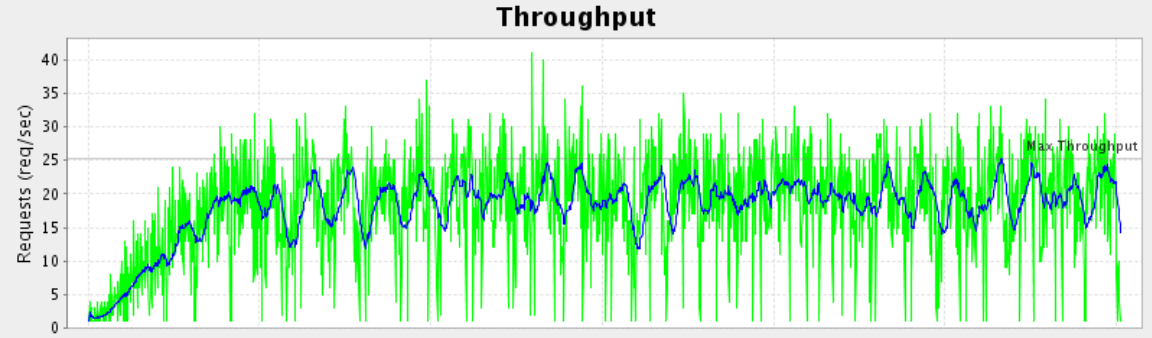
As above, we define the High Load scenario as making 25 requests per second. Few customers will reach these levels of requests per second, but if you have several ten thousand users these levels can be reached during peak business hours. This test is run on servers with 2 CPUs per machine.

Confluence 3.0 vs Confluence 2.10 response times

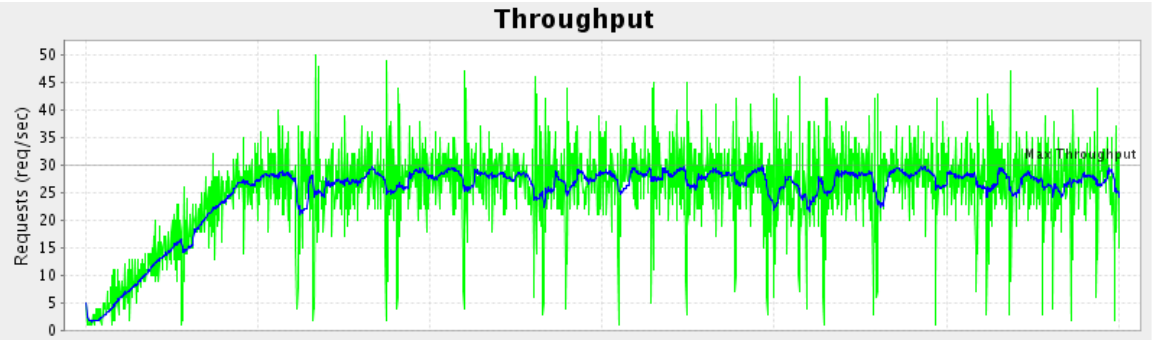
Scenario	Average time in 2.10	Average time in 3.0	Improvement	95 percent in 2.10	95 percent in 3.0	Improvement
Browse Label	4674ms	1447ms	222%	12831ms	3822ms	235%
Commentor submit comment	584ms	442ms	31%	1948ms	1340ms	45%
Commentor view commented page	1728ms	680ms	154%	5943ms	2164ms	174%
Commentor view page	2048ms	893ms	129%	7111ms	2868ms	147%
Creator add page	838ms	245ms	241%	2333ms	562ms	314%
Creator submit new page	896ms	466ms	92%	2308ms	1181ms	95%
Creator view page	443ms	155ms	186%	1300ms	234ms	454%
Dashboard	2707ms	339ms	697%	7781ms	427ms	1722%
Editor display page	960ms	446ms	115%	2909ms	1633ms	78%
Edit Page	735ms	255ms	188%	2276ms	699ms	225%
Editor submit edit	1888ms	1108ms	70%	6513ms	4060ms	60%
Go to log in page	1525ms	256ms	494%	4650ms	524ms	786%
Log In	1278ms	406ms	214%	3712ms	598ms	520%
Quick Navigation Search	540ms	49ms	1000%	4292ms	95ms	4418%
Reader Not Found	882ms	413ms	113%	2151ms	813ms	164%
Reader RSS Blogpost Atom	1165ms	49ms	2245%	5052ms	72ms	6888%
Reader RSS Blogpost RSS2	1494ms	51ms	2825%	5565ms	69ms	7872%
Reader RSS Comment Atom	1655ms	195ms	748%	5990ms	763ms	684%
Reader RSS Comment RSS2	1497ms	197ms	656%	5892ms	822ms	616%
Reader RSS Page Atom	2440ms	630ms	287%	8300ms	2263ms	266%
Reader RSS Page RSS2	2562ms	685ms	273%	8027ms	2530ms	217%
View Page	1780ms	750ms	137%	5560ms	2728ms	103%
Reader for Space Page	1668ms	835ms	99%	4177ms	1976ms	111%
Search Site	1691ms	275ms	513%	5853ms	407ms	1338%



Confluence 2.10 throughput



Confluence 3.0 throughput



High Load comparison between 2.10 and 3.0 in clustered mode

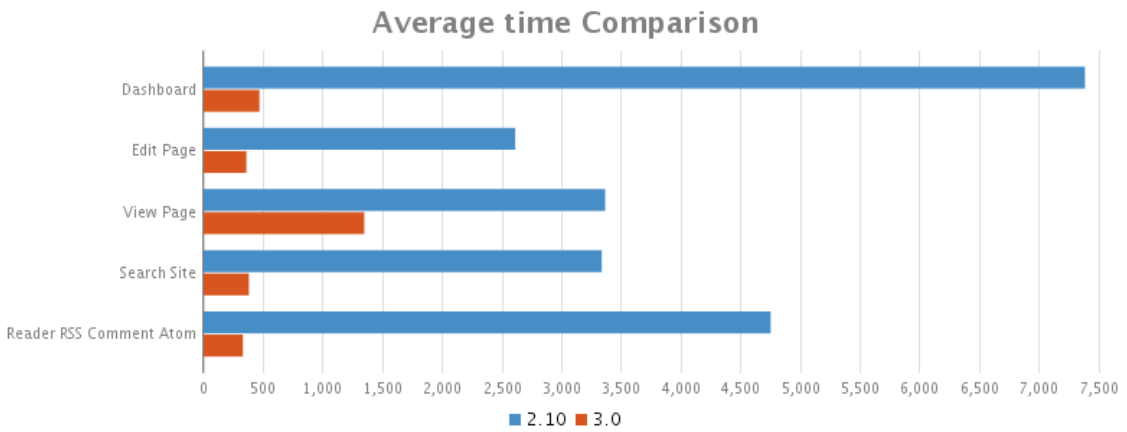
In this test we show how using a Cluster for high load instances can increase throughput and reduce response time. Confluence 3.0 has many improvements which benefit the clustered version. In the case of the test above, we can see that as the load is increased, Confluence is able to use more of the available CPU power on the 8 core machines to scale up and handle the higher load with an very good response time. This is where clustering makes a lot of sense now.

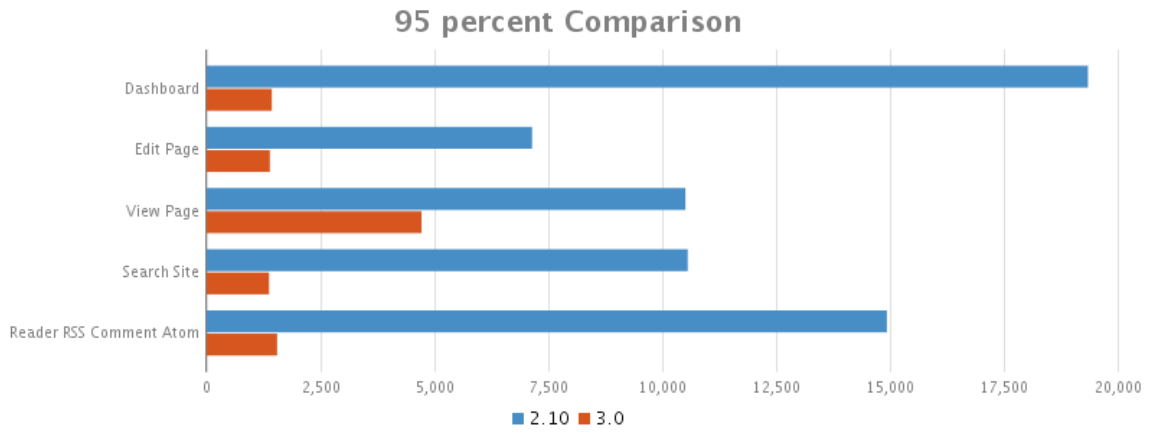
Peak Load Scenario, Clustered, 2 nodes, 2 CPUs per node

As above, we define peak load as the load generator making around 35 requests per second. During this test we used 2 CPUs per machine.

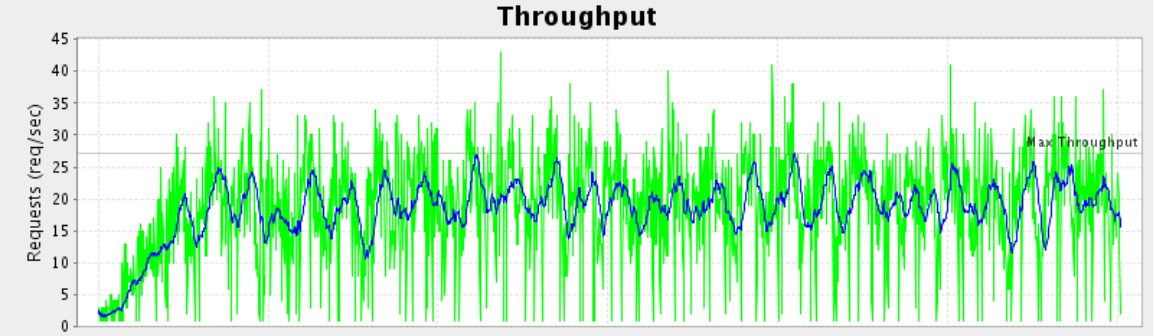
Confluence 3.0 vs Confluence 2.10 response times

Scenario	Average time in 2.10	Average time in 3.0	Improvement	95 percent in 2.10	95 percent in 3.0	Improvement
Browse Label	9179ms	2150ms	326%	22987ms	5695ms	303%
Commentor submit comment	999ms	859ms	16%	3175ms	2712ms	17%
Commentor view commented page	2760ms	1213ms	127%	8739ms	3908ms	123%
Commentor view page	3225ms	1672ms	92%	10941ms	5323ms	105%
Creator add page	2396ms	379ms	532%	7285ms	1487ms	389%
Creator submit new page	1913ms	858ms	122%	4850ms	2548ms	90%
Creator view page	1070ms	270ms	296%	2925ms	1130ms	158%
Dashboard	7383ms	466ms	1481%	19349ms	1429ms	1254%
Editor display page	2460ms	761ms	223%	7388ms	2737ms	169%
Edit Page	2609ms	357ms	630%	7143ms	1385ms	415%
Editor submit edit	4064ms	1821ms	123%	12599ms	6287ms	100%
Go to log in page	3622ms	280ms	1191%	13728ms	497ms	2657%
Log In	3603ms	497ms	625%	12477ms	1045ms	1093%
Quick Navigation Search	1268ms	92ms	1273%	9515ms	435ms	2084%
Reader Not Found	1591ms	539ms	195%	3616ms	1622ms	122%
Reader RSS Blogpost Atom	4514ms	78ms	5688%	14570ms	136ms	10605%
Reader RSS Blogpost RSS2	4666ms	84ms	5416%	14554ms	134ms	10689%
Reader RSS Comment Atom	4750ms	328ms	1346%	14934ms	1545ms	866%
Reader RSS Comment RSS2	4723ms	302ms	1460%	16526ms	1412ms	1070%
Reader RSS Page Atom	6443ms	1012ms	536%	20556ms	4005ms	413%
Reader RSS Page RSS2	6287ms	1109ms	466%	17762ms	4175ms	325%
View Page	3363ms	1345ms	150%	10510ms	4717ms	122%
Reader for Space Page	3069ms	1161ms	164%	9475ms	3023ms	213%
Search Site	3334ms	378ms	780%	10560ms	1368ms	671%

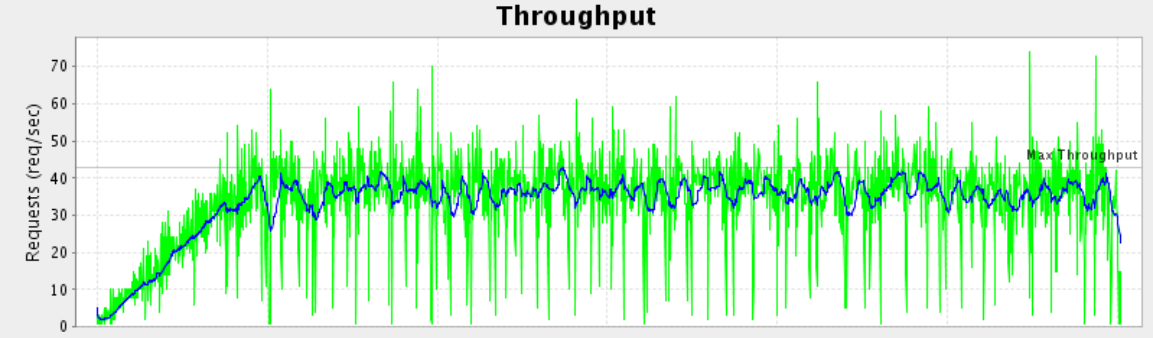




Confluence 2.10 throughput



Confluence 3.0 throughput



Peak Load comparison between 2.10 and 3.0 in clustered mode

This test highlights how well Confluence 3.0 can now scale. Response times remain low as the load is increased. Confluence 3.0 is able to make far better use of more powerful hardware than Confluence 2.10 which is shown by the improved response times for key scenarios like Page view and Dashboard.

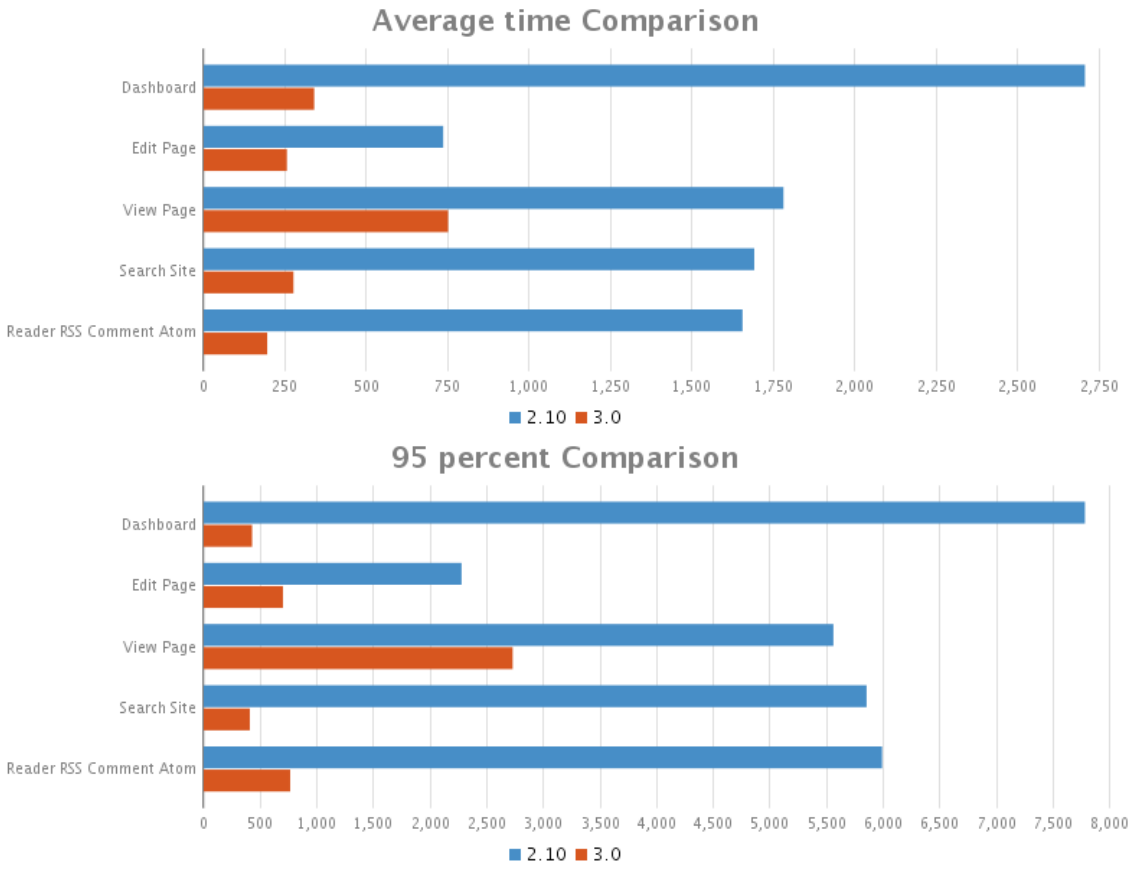
Feedback welcome

We welcome your feedback! Is this document understandable, does it cover the areas that you are most interested about? Tell us and leave comments on this page!

High Load Cluster

Scenario	Average time in 2.10	Average time in 3.0	Improvement	95 percent in 2.10	95 percent in 3.0	Improvement
Browse Label	4674ms	1447ms	222%	12831ms	3822ms	235%
Commentor submit comment	584ms	442ms	31%	1948ms	1340ms	45%
Commentor view commented page	1728ms	680ms	154%	5943ms	2164ms	174%
Commentor view page	2048ms	893ms	129%	7111ms	2868ms	147%
Creator add page	838ms	245ms	241%	2333ms	562ms	314%
Creator submit new page	896ms	466ms	92%	2308ms	1181ms	95%

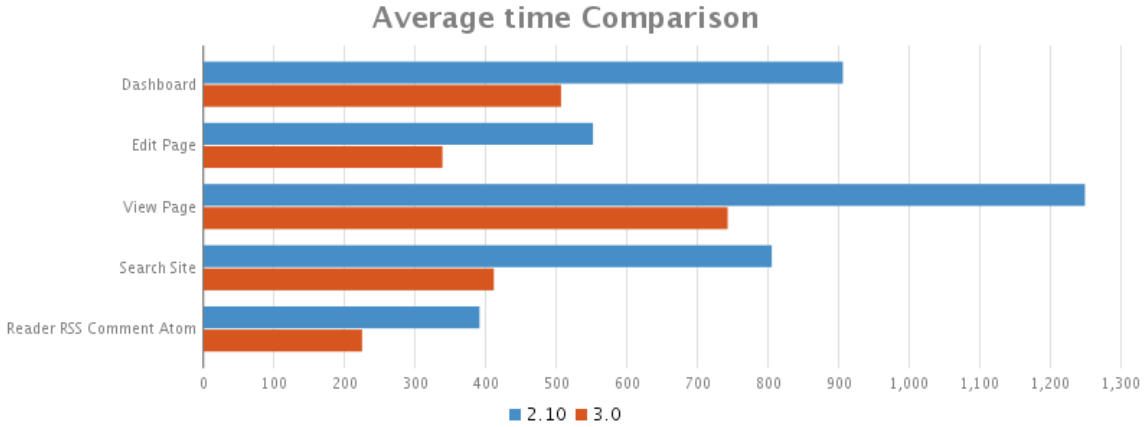
Creator view page	443ms	155ms	186%	1300ms	234ms	454%
Dashboard	2707ms	339ms	697%	7781ms	427ms	1722%
Editor display page	960ms	446ms	115%	2909ms	1633ms	78%
Edit Page	735ms	255ms	188%	2276ms	699ms	225%
Editor submit edit	1888ms	1108ms	70%	6513ms	4060ms	60%
Go to log in page	1525ms	256ms	494%	4650ms	524ms	786%
Log In	1278ms	406ms	214%	3712ms	598ms	520%
Quick Navigation Search	540ms	49ms	1000%	4292ms	95ms	4418%
Reader Not Found	882ms	413ms	113%	2151ms	813ms	164%
Reader RSS Blogpost Atom	1165ms	49ms	2245%	5052ms	72ms	6888%
Reader RSS Blogpost RSS2	1494ms	51ms	2825%	5565ms	69ms	7872%
Reader RSS Comment Atom	1655ms	195ms	748%	5990ms	763ms	684%
Reader RSS Comment RSS2	1497ms	197ms	656%	5892ms	822ms	616%
Reader RSS Page Atom	2440ms	630ms	287%	8300ms	2263ms	266%
Reader RSS Page RSS2	2562ms	685ms	273%	8027ms	2530ms	217%
View Page	1780ms	750ms	137%	5560ms	2728ms	103%
Reader for Space Page	1668ms	835ms	99%	4177ms	1976ms	111%
Search Site	1691ms	275ms	513%	5853ms	407ms	1338%

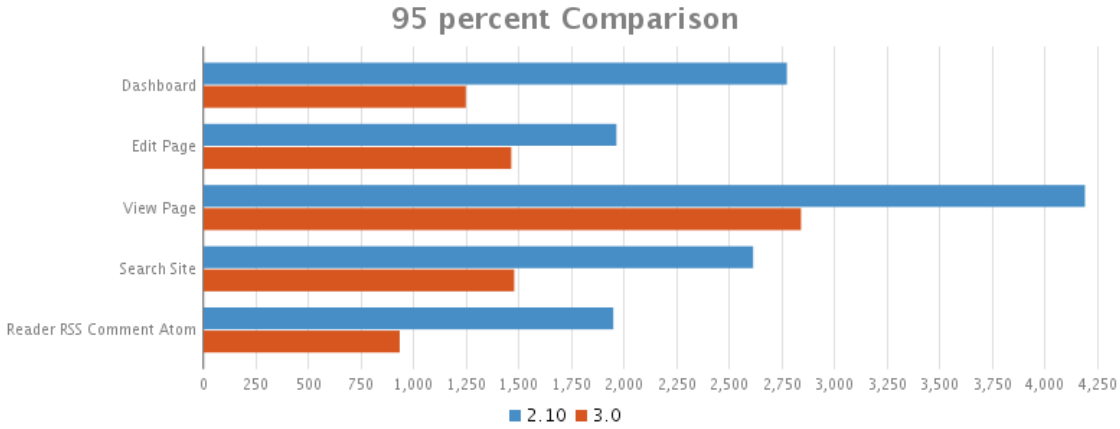


High Load Single Node

Scenario	Average time in 2.10	Average time in 3.0	Improvement	95 percent in 2.10	95 percent in 3.0	Improvement
Browse Label	2389ms	1531ms	56%	6196ms	4195ms	47%
Commentor submit comment	424ms	397ms	6%	1779ms	1603ms	10%

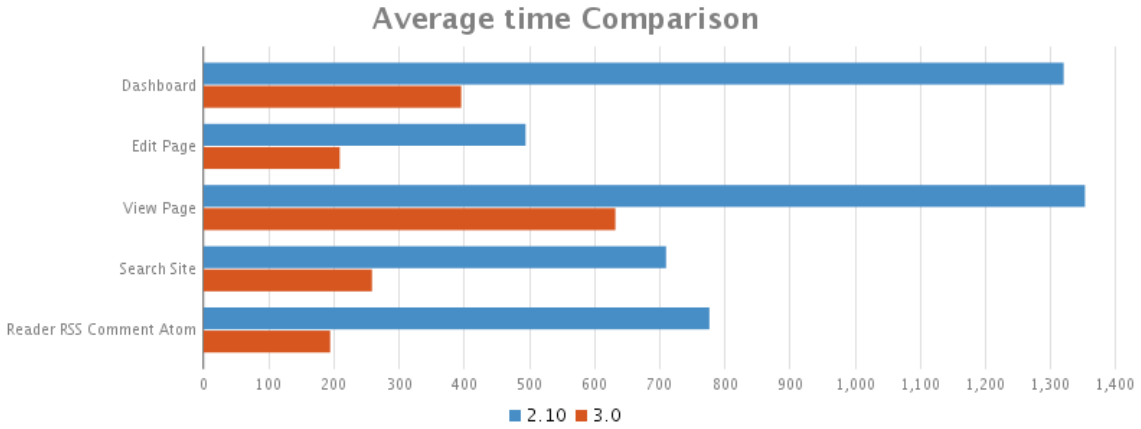
Commentor view commented page	1211ms	815ms	48%	4729ms	2863ms	65%
Commentor view page	1402ms	912ms	53%	5284ms	3094ms	70%
Creator add page	558ms	297ms	87%	1962ms	1543ms	27%
Creator submit new page	783ms	522ms	49%	2567ms	1545ms	66%
Creator view page	443ms	284ms	55%	1845ms	843ms	118%
Dashboard	905ms	506ms	78%	2771ms	1245ms	122%
Editor display page	807ms	504ms	59%	2650ms	2165ms	22%
Edit Page	551ms	338ms	63%	1961ms	1461ms	34%
Editor submit edit	1524ms	1180ms	29%	5115ms	4189ms	22%
Go to log in page	409ms	419ms	-2%	1171ms	982ms	19%
Log In	520ms	346ms	50%	2124ms	700ms	203%
Quick Navigation Search	318ms	124ms	155%	1895ms	369ms	413%
Reader Not Found	866ms	492ms	76%	2439ms	1579ms	54%
Reader RSS Blogpost Atom	300ms	105ms	186%	1549ms	191ms	709%
Reader RSS Blogpost RSS2	299ms	98ms	203%	1954ms	183ms	965%
Reader RSS Comment Atom	390ms	224ms	74%	1946ms	931ms	108%
Reader RSS Comment RSS2	362ms	254ms	42%	1824ms	1196ms	52%
Reader RSS Page Atom	1098ms	777ms	41%	4804ms	2848ms	68%
Reader RSS Page RSS2	1126ms	807ms	39%	4532ms	3406ms	33%
View Page	1248ms	742ms	68%	4188ms	2839ms	47%
Reader for Space Page	1410ms	914ms	54%	3749ms	2487ms	50%
Search Site	804ms	411ms	95%	2611ms	1475ms	76%

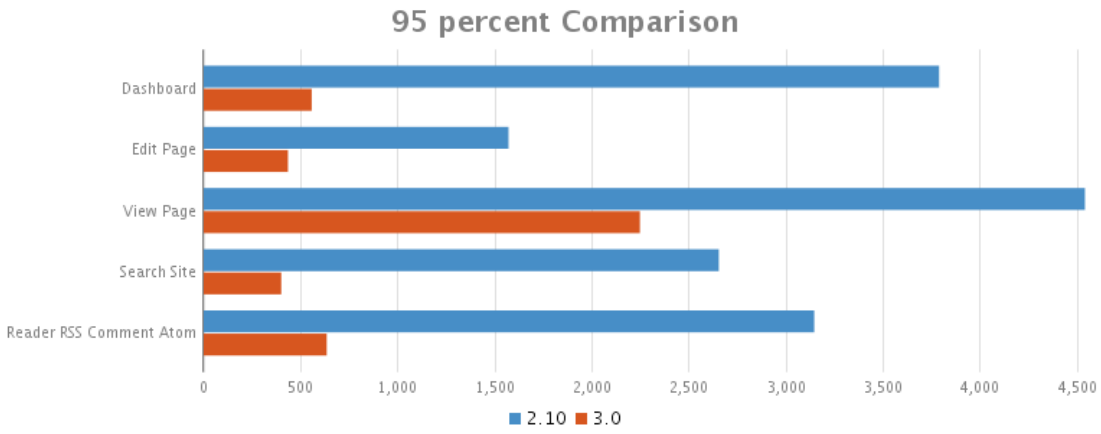




Medium Load Cluster

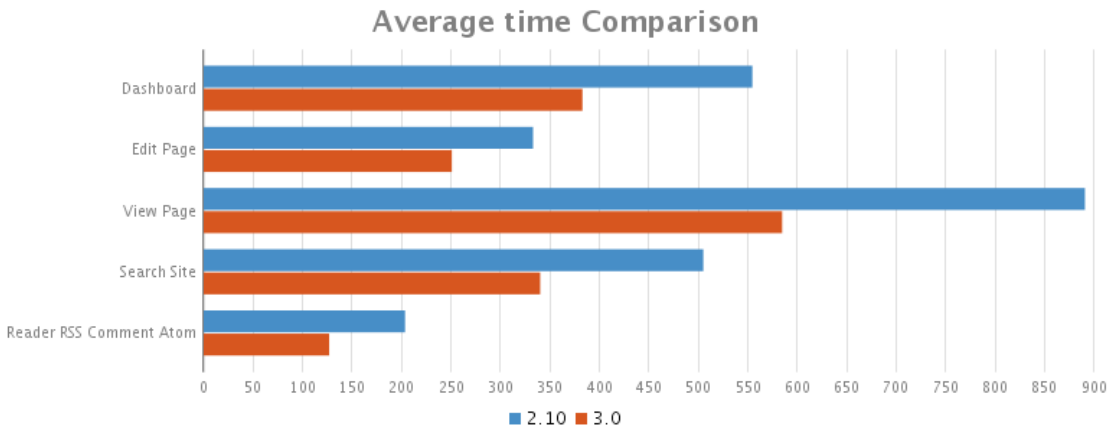
Scenario	Average time in 2.10	Average time in 3.0	Improvement	95 percent in 2.10	95 percent in 3.0	Improvement
Browse Label	2477ms	919ms	169%	6365ms	2143ms	196%
Commentor submit comment	410ms	380ms	7%	1127ms	856ms	31%
Commentor view commented page	1029ms	595ms	72%	4193ms	1826ms	129%
Commentor view page	1367ms	786ms	73%	5264ms	2557ms	105%
Creator add page	611ms	214ms	184%	1463ms	414ms	253%
Creator submit new page	692ms	422ms	63%	1596ms	938ms	70%
Creator view page	329ms	131ms	150%	1034ms	205ms	404%
Dashboard	1319ms	395ms	234%	3787ms	556ms	581%
Editor display page	750ms	387ms	93%	2592ms	1420ms	82%
Edit Page	493ms	208ms	136%	1569ms	433ms	261%
Editor submit edit	1500ms	980ms	53%	4903ms	3435ms	42%
Go to log in page	854ms	357ms	138%	2898ms	547ms	429%
Log In	961ms	422ms	127%	3320ms	597ms	455%
Quick Navigation Search	183ms	56ms	226%	1178ms	107ms	1001%
Reader Not Found	663ms	351ms	89%	1562ms	554ms	181%
Reader RSS Blogpost Atom	499ms	73ms	577%	2674ms	118ms	2166%
Reader RSS Blogpost RSS2	453ms	65ms	589%	1863ms	122ms	1420%
Reader RSS Comment Atom	776ms	194ms	300%	3143ms	633ms	396%
Reader RSS Comment RSS2	742ms	186ms	297%	2930ms	576ms	408%
Reader RSS Page Atom	1378ms	618ms	122%	4693ms	2109ms	122%
Reader RSS Page RSS2	1497ms	584ms	156%	4767ms	1899ms	150%
View Page	1352ms	631ms	114%	4538ms	2246ms	102%
Reader for Space Page	1251ms	793ms	57%	3124ms	1736ms	79%
Search Site	709ms	258ms	174%	2652ms	399ms	564%

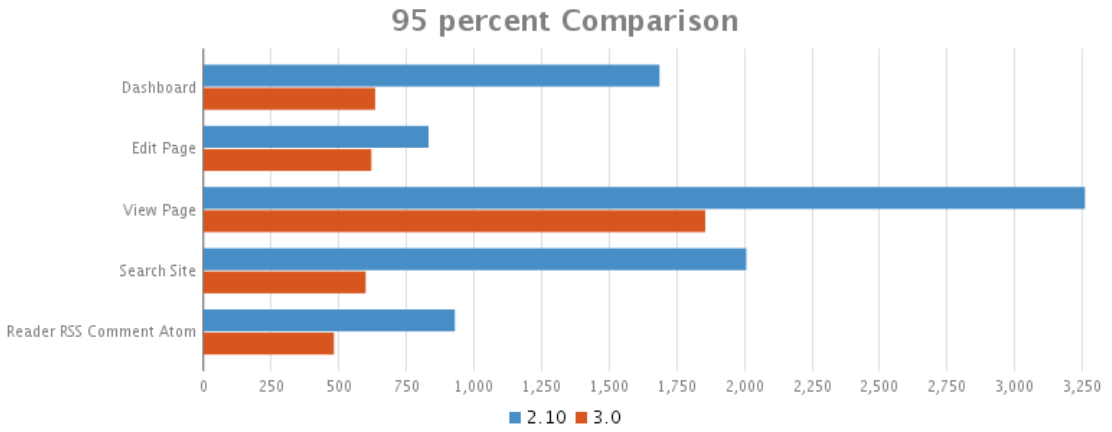




Medium Load Single Node

Scenario	Average time in 2.10	Average time in 3.0	Improvement	95 percent in 2.10	95 percent in 3.0	Improvement
Browse Label	1619ms	1129ms	43%	3979ms	2387ms	66%
Commentor submit comment	306ms	338ms	-9%	805ms	794ms	1%
Commentor view commented page	737ms	628ms	17%	3386ms	1783ms	89%
Commentor view page	989ms	707ms	39%	4133ms	2168ms	90%
Creator add page	402ms	211ms	90%	765ms	391ms	95%
Creator submit new page	525ms	387ms	35%	1161ms	882ms	31%
Creator view page	256ms	233ms	9%	704ms	501ms	40%
Dashboard	554ms	382ms	44%	1685ms	634ms	165%
Editor display page	520ms	417ms	24%	1881ms	1065ms	76%
Edit Page	332ms	250ms	32%	831ms	620ms	34%
Editor submit edit	1199ms	961ms	24%	3949ms	3459ms	14%
Go to log in page	274ms	456ms	-39%	486ms	2795ms	-82%
Log In	342ms	333ms	2%	774ms	480ms	61%
Quick Navigation Search	134ms	57ms	133%	597ms	110ms	439%
Reader Not Found	551ms	369ms	49%	1266ms	615ms	105%
Reader RSS Blogpost Atom	170ms	59ms	184%	637ms	67ms	838%
Reader RSS Blogpost RSS2	206ms	95ms	116%	754ms	97ms	675%
Reader RSS Comment Atom	203ms	126ms	60%	929ms	481ms	92%
Reader RSS Comment RSS2	369ms	151ms	143%	1602ms	477ms	235%
Reader RSS Page Atom	628ms	513ms	22%	2725ms	2147ms	26%
Reader RSS Page RSS2	800ms	547ms	46%	3381ms	2196ms	53%
View Page	890ms	584ms	52%	3259ms	1854ms	75%
Reader for Space Page	904ms	677ms	33%	2219ms	1566ms	41%
Search Site	505ms	340ms	48%	2006ms	598ms	235%

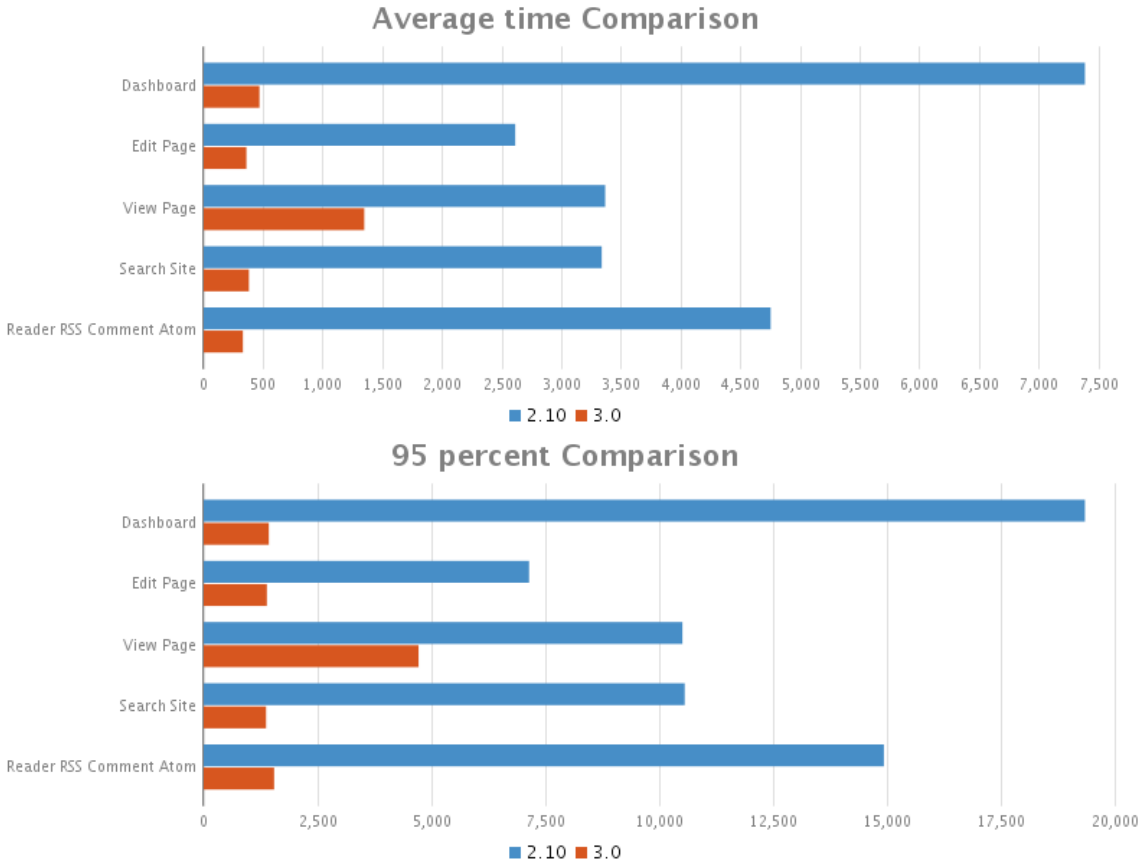




Peak Load Cluster

Scenario	Average time in 2.10	Average time in 3.0	Improvement	95 percent in 2.10	95 percent in 3.0	Improvement
Browse Label	9179ms	2150ms	326%	22987ms	5695ms	303%
Commentor submit comment	999ms	859ms	16%	3175ms	2712ms	17%
Commentor view commented page	2760ms	1213ms	127%	8739ms	3908ms	123%
Commentor view page	3225ms	1672ms	92%	10941ms	5323ms	105%
Creator add page	2396ms	379ms	532%	7285ms	1487ms	389%
Creator submit new page	1913ms	858ms	122%	4850ms	2548ms	90%
Creator view page	1070ms	270ms	296%	2925ms	1130ms	158%
Dashboard	7383ms	466ms	1481%	19349ms	1429ms	1254%
Editor display page	2460ms	761ms	223%	7388ms	2737ms	169%
Edit Page	2609ms	357ms	630%	7143ms	1385ms	415%
Editor submit edit	4064ms	1821ms	123%	12599ms	6287ms	100%
Go to log in page	3622ms	280ms	1191%	13728ms	497ms	2657%
Log In	3603ms	497ms	625%	12477ms	1045ms	1093%
Quick Navigation Search	1268ms	92ms	1273%	9515ms	435ms	2084%
Reader Not Found	1591ms	539ms	195%	3616ms	1622ms	122%
Reader RSS Blogpost Atom	4514ms	78ms	5688%	14570ms	136ms	10605%
Reader RSS Blogpost RSS2	4666ms	84ms	5416%	14554ms	134ms	10689%
Reader RSS Comment Atom	4750ms	328ms	1346%	14934ms	1545ms	866%
Reader RSS Comment RSS2	4723ms	302ms	1460%	16526ms	1412ms	1070%
Reader RSS Page Atom	6443ms	1012ms	536%	20556ms	4005ms	413%

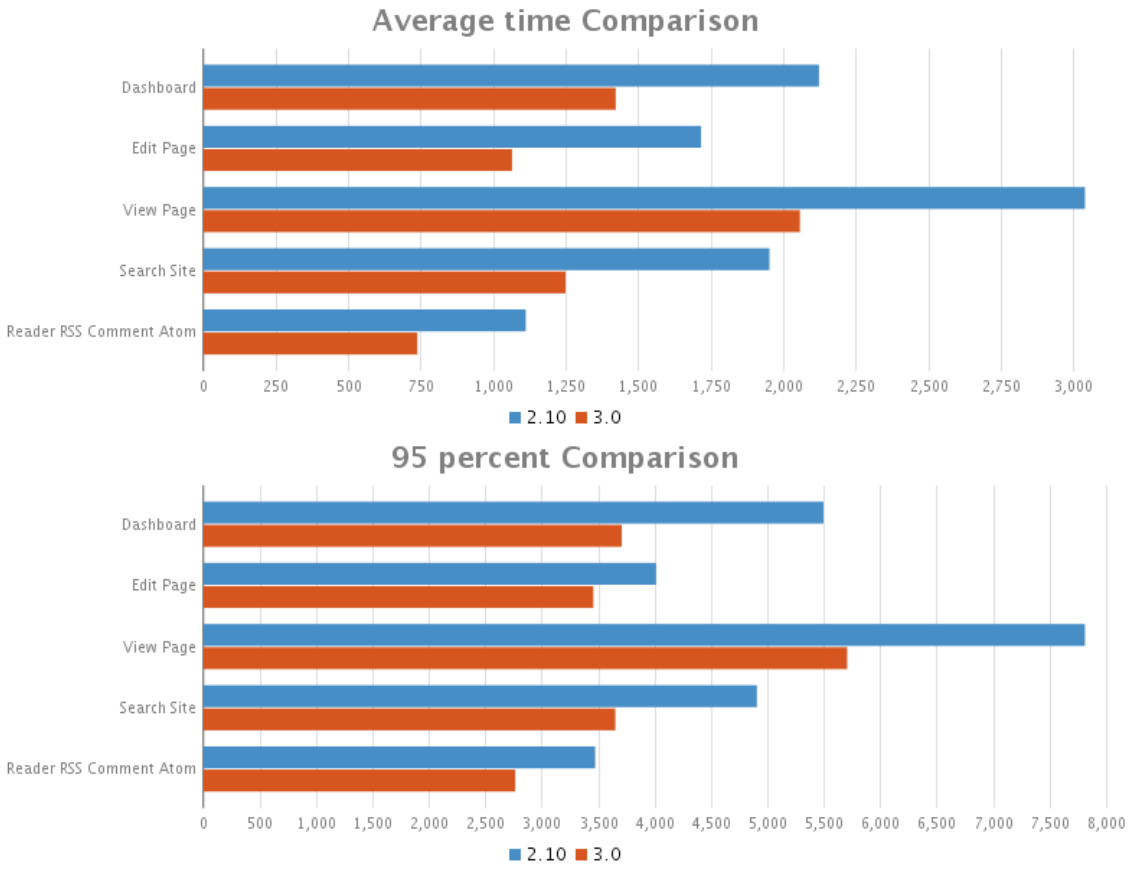
Reader RSS Page RSS2	6287ms	1109ms	466%	17762ms	4175ms	325%
View Page	3363ms	1345ms	150%	10510ms	4717ms	122%
Reader for Space Page	3069ms	1161ms	164%	9475ms	3023ms	213%
Search Site	3334ms	378ms	780%	10560ms	1368ms	671%



Peak Load Single Node

Scenario	Average time in 2.10	Average time in 3.0	Improvement	95 percent in 2.10	95 percent in 3.0	Improvement
Browse Label	4747ms	3207ms	47%	10951ms	7575ms	44%
Commentor submit comment	1517ms	1146ms	32%	4521ms	3611ms	25%
Commentor view commented page	3148ms	2173ms	44%	9222ms	6184ms	49%
Commentor view page	3302ms	2317ms	42%	9891ms	6410ms	54%
Creator add page	1693ms	934ms	81%	3904ms	3170ms	23%
Creator submit new page	2777ms	1959ms	41%	5812ms	5170ms	12%
Creator view page	1589ms	1065ms	49%	3523ms	3358ms	4%
Dashboard	2121ms	1420ms	49%	5492ms	3704ms	48%
Editor display page	2216ms	1502ms	47%	5081ms	4233ms	20%
Edit Page	1714ms	1062ms	61%	4008ms	3452ms	16%
Editor submit edit	3945ms	3205ms	23%	10523ms	9467ms	11%
Go to log in page	934ms	818ms	14%	4544ms	4091ms	11%
Log In	807ms	913ms	-11%	2879ms	3531ms	-18%
Quick Navigation Search	1121ms	568ms	97%	4288ms	2704ms	58%
Reader Not Found	2159ms	1222ms	76%	4265ms	3472ms	22%

Reader RSS Blogpost Atom	864ms	531ms	62%	2796ms	2511ms	11%
Reader RSS Blogpost RSS2	1099ms	527ms	108%	4307ms	2691ms	60%
Reader RSS Comment Atom	1110ms	736ms	50%	3469ms	2760ms	25%
Reader RSS Comment RSS2	1159ms	863ms	34%	3959ms	3130ms	26%
Reader RSS Page Atom	2395ms	2300ms	4%	7588ms	7342ms	3%
Reader RSS Page RSS2	2661ms	2258ms	17%	8295ms	7922ms	4%
View Page	3038ms	2055ms	47%	7809ms	5702ms	36%
Reader for Space Page	3005ms	1990ms	51%	6298ms	4850ms	29%
Search Site	1950ms	1247ms	56%	4902ms	3647ms	34%



Issues Resolved in Confluence 3.0

Below are the issues resolved in Confluence 3.0, ordered by number of votes. For the full details of the fixes, improvements and new features, please take a look at our issue tracker. You can also take a look at the Confluence 3.0 Release Notes.

JIRA Issues (200 issues)

Type	Key	Summary	Status	Resolution	Votes
	CONF-13276	Restrict anonymous users from viewing user profiles.	Resolved	Fixed	7
	CONF-8298	Reduce the number of directories created under ${confluence.home}/attachments	Resolved	Fixed	5
	CONF-14156	Code Macro - Rich Text Editor removes spaces	Resolved	Fixed	4
	CONF-9928	WIki markup containing <tag> elements is erased when switching to the RTE	Resolved	Fixed	3
	CONF-11172	Line-break markup (\\) uses clear=all style, breaking blog post layout in 2.8	Resolved	Fixed	2
	CONF-13823	Editing a page with columns in RTE overwrites columns width parameter and sets		Fixed	2

		to 100%	Resolved		
	CONF-14798	pdf {viewfile} slide renderer creates vast amounts of garbage which lead to poor performance	Resolved	Fixed	1
	CONF-14248	Exception given when Doc Import tries to create pages with duplicate names	Resolved	Fixed	1
	CONF-15986	Page's attachments become inaccessible when page is moved to a different space	Resolved	Fixed	0
	CONF-16184	Lock contentions in didyoumean code	Resolved	Fixed	0
	CONF-14539	The XStream serializer only uses the plugins class loader	Resolved	Fixed	0
	CONF-15376	Gallery Macro XSS	Resolved	Fixed	0
	CONF-15644	Did-you-mean search IncrementalIndexBuilder uses large amount of PermGen memory to store words	Resolved	Fixed	0
	CONF-15146	Saving a page or Switching to wiki markup fails with java.lang.AbstractMethodError: getTextContent error in Weblogic	Resolved	Fixed	0
	CONF-15754	Jiraissues add icon mapping configuration is susceptible to XSS	Resolved	Fixed	0
	CONF-15809	Viewfile macros do not respect page restrictions	Resolved	Fixed	0
	CONF-2079	More control over PDF exporting	Resolved	Fixed	232
	CONF-5853	Tables in PDF Exports have the same width for every column. Width should be appropriate to the column content.	Resolved	Fixed	96
	CONF-12836	Add ability to tune caches from the UI	Resolved	Fixed	36
	CONF-8436	Add support for page breaks in PDF page exports	Resolved	Fixed	25
	CONF-8302	Allow recently updated macro to filter results by users	Resolved	Fixed	21
	CONF-8945	Watch emails (change notifications) should contain html diffs	Resolved	Fixed	21
	CONF-1790	Export Page Hierarchy	Resolved	Fixed	19
	CONF-8071	Generate PDF with specific cover page or custom layout	Resolved	Fixed	17
	CONF-7910	PDF Export of wide tables does not resize the table properly to fit the 'paper' size.	Resolved	Fixed	15
	CONF-1104	Allow configuration of page size and orientation in PDF Export	Resolved	Fixed	14
	CONF-8273	Copying a page to a new space should not prepend 'Copy Of' to the copied page	Resolved	Fixed	12
	CONF-7708	html macro loses formatting wiki markup and rich text	Resolved	Fixed	10
	CONF-14058	Whitespace added between {section} and {column} every wysiwyg round-trip	Resolved	Fixed	9
	CONF-13524	Left navigation theme displays "news operations" instead of "page operations" when viewing the left nav main menu categories	Resolved	Fixed	9
	CONF-6750	Custom layouts for PDF exports	Resolved	Fixed	7
	CONF-12083	Outlook 2007 is unable to read Confluence RSS feeds over HTTPS		Fixed	7

			Resolved		
	CONF-11552	height and width of embedded movies is changed to 32 by Rich Text Editor	Resolved	Fixed	7
	CONF-7225	PDF Export throws an exception when an unterminated macro is used in a heading	Resolved	Fixed	6
	CONF-6756	Improve the Export Space pages selection to allow easy selection of a page and all it's children	Resolved	Fixed	5
	CONF-14698	A number of {{FelixOsgiContainerManager}} messages in the logs since the introduction of Atlassian Plugins 2	Resolved	Fixed	5
	CONF-11878	PDF Customization Console	Resolved	Fixed	5
	CONF-9934	Full screen editor does not autosave to Drafts folder	Resolved	Fixed	5
	CONF-11286	Allow ability to view a users profile after they've created a personal space	Resolved	Fixed	5
	CONF-7027	Wiki markup for graphical emoticons not escaped by rich text editor	Resolved	Fixed	5
	CONF-7491	Escaped (+) and (-) in content revert to emoticons when opening for Edit	Resolved	Fixed	5
	CONF-8034	Serve attachments, including embedded images, with cache headers	Resolved	Fixed	4
	CONF-7398	Make PDF export fonts and styles themeable	Resolved	Fixed	4
	CONF-14175	Confluence Macro Browser	Resolved	Fixed	4
	CONF-12321	Update attachment comment should not update last modification date of all historical versions	Resolved	Fixed	3
	CONF-4128	Can't copy and paste text from Confluence generated PDFs	Resolved	Fixed	3
	CONF-15040	Confluence 2.10.x incompatible with WebSphere - plugins do not load	Resolved	Fixed	3
	CONF-14505	Wildcards in "manage user" searches cause performance problems with large LDAP repositories	Resolved	Fixed	3
	CONF-9234	Label order in popular-labels macro changes when using a space key parameter	Resolved	Fixed	3
	CONF-7623	PDF export is missing the title information from the code macro in the page	Resolved	Fixed	3
	CONF-9901	Export space to PDF or HTML fails if the space has attachments	Resolved	Fixed	3
	CONF-11799	Editing tables with Rich Text - Tab button jumps to location in Firefox 3	Resolved	Fixed	3
	CONF-13188	Code macro breaks PDF export when there is a pattern //<keyword> <keyword>	Resolved	Fixed	2
	CONF-7425	Error in pdf export using content formatting macro	Resolved	Fixed	2
	CONF-15864	Switching between Rich text editor and Wiki markup editor causes code macro add line space	Resolved	Fixed	2
	CONF-6865	Side by side images will expand across the bottom page margin on PDF Export.	Resolved	Fixed	2
	CONF-5990	export to PDF truncates HTML links to external sites	Resolved	Fixed	2
	CONF-14791	Confluence doesn't start up on OpenJDK because it requires		Fixed	2

		com.sun.image.codec.jpeg.JPEGImageEncoder	Resolved		
	CONF-8709	On Space Export screen add the ability to select all children of a page for export	Resolved	Fixed	2
	CONF-12366	Image width and height are to 32 by wysiwyg editor for missing image files.	Resolved	Fixed	2
	CONF-14072	Alphabetical listing of pages does not work when there are too many pages beginning with a letter	Resolved	Fixed	2
	CONF-3904	Hint macro not displaying correctly in page created from PDF Export	Resolved	Fixed	2
	CONF-13933	Add optional diff to email notifications and improve the diffing logic	Resolved	Fixed	2
	CONF-12634	PDF Exporter doesn't handle situations when JTidy fails	Resolved	Fixed	2
	CONF-13094	Export page tree should be use subtree selection	Resolved	Fixed	2
	CONF-7209	Em-dash inserted in Rich Text Edit breaks after saving or changing to WIKI markup mode	Resolved	Fixed	2
	CONF-4334	Escaping smileys in wiki markup doesn't work	Resolved	Fixed	2
	CONF-13663	Links within Charting Portlet displaying results inside IFrame	Resolved	Fixed	1
	CONF-14491	Add a system property to enable logging of Macros used.	Resolved	Fixed	1
	CONF-14128	Reduce blocking in Velocity when accessing cache	Resolved	Fixed	1
	CONF-7412	Text inside code macro displays outside boundaries in exported PDF	Resolved	Fixed	1
	CONF-13769	The full screen view has no scroll bar	Resolved	Fixed	1
	CONF-13795	Index searcher not properly closed after searches	Resolved	Fixed	1
	CONF-8453	Links created via the Rich Text Editor dialog lose the space character immediately after the link	Resolved	Fixed	1
	CONF-13684	Inserting a link in RTE behaves inconsistently across different browsers	Resolved	Fixed	1
	CONF-6860	Blank lines in a table row are lost in a PDF Export of the table	Resolved	Fixed	1
	CONF-14282	Section macro loses parameters when switching between RTE and Wiki Markup	Resolved	Fixed	1
	CONF-7432	PDF export fails when uploaded font contains a space in its path	Resolved	Fixed	1
	CONF-7868	PDF Export of page with user macros fail	Resolved	Fixed	1
	CONF-14783	Encoding test should use a GET submission to mimic quick search	Resolved	Fixed	1
	CONF-8320	Search macro output inconsistent with notation guide	Resolved	Fixed	1
	CONF-13843	Recently-Updated macro throws a BooleanQuery$TooManyClauses exception	Resolved	Fixed	1
	CONF-12058	popular-labels macro doesn't order the labels alphabetically	Resolved	Fixed	1
	CONF-7749	Don't block index readers during index updates and optimization		Fixed	1

			Resolved		
	CONF-14724	Clicking on thumbnail images in IE prompts for download instead of opening popup with large image	Resolved	Fixed	1
	CONF-12612	Menus look bad in IE7	Resolved	Fixed	1
	CONF-9736	Large embedded images are converted to a thumbnail size when exported into the PDF format	Resolved	Fixed	1
	CONF-15120	Can't install Office Connector firefox plugin because 'Signing could not be verified'	Resolved	Fixed	1
	CONF-14733	Exessive logging during export word page	Resolved	Fixed	1
	CONF-8322	PDF export fails under certain unfortunate naming of pages and headings	Resolved	Fixed	1
	CONF-9603	Restore recent activity to a user profile screen	Resolved	Fixed	1
	CONF-11762	Support for DB2 9	Resolved	Fixed	1
	CONF-15450	If a user has no text in their profile, XML-RPC calls throw a null pointer exception (NPE)	Resolved	Fixed	1
	CONF-14222	i18n support for macro browser	Resolved	Fixed	1
	CONF-13814	Error removing space with custom space logo and attachments stored in the database	Resolved	Fixed	1
	CONF-8647	Removing an entire table in Rich Text Editor is not intuitive	Resolved	Fixed	1
	CONF-13825	NullPointerException when thumbnail cannot be generated	Resolved	Fixed	1
	CONF-10166	Inconsistent escaping of emoticons in rich-text edit mode	Resolved	Fixed	1
	CONF-12894	Slow page rendering when using Confluence via HTTPS / SSL due to lack of caching	Resolved	Fixed	1
	CONF-14029	Javascript error on "Add Members" in Manage Groups in IE	Resolved	Fixed	1
	CONF-14462	User Status	Resolved	Fixed	0
	CONF-14094	Confluence Support Request Improvements	Resolved	Fixed	0
	CONF-14387	Formatting problems in PDF export	Resolved	Fixed	0
	CONF-14438	Deprecate com.atlassian.confluence.util.io.IOUtils	Resolved	Fixed	0
	CONF-14495	Remove public logger in VersionedHibernateObjectDao	Resolved	Fixed	0
	CONF-14490	Add the ability to Hibernate to have more than one index per column	Resolved	Fixed	0
	CONF-15399	Contributors Macro noneFoundMessage XSS Vector	Resolved	Fixed	0
	CONF-14516	Only aquire a lock for deferred operations cache if neccessary	Resolved	Fixed	0
	CONF-14517	Look and Feel improvements to User Profile	Resolved	Fixed	0
	CONF-14533	Expose the id in content name search results		Fixed	0

			Resolved		
	CONF-12267	When editing existing page, extra space appears at bottom	Resolved	Fixed	0
	CONF-16133	WebDAV plugin documentation is out dated on CAC	Resolved	Fixed	0
	CONF-16073	{contentbylabel} macro does not filter results based on comments	Resolved	Fixed	0
	CONF-5415	Investigate performance gains of multicolumn and functional indexes in Postgres	Resolved	Fixed	0
	CONF-7484	Pages with image larger than the PDF 'paper' size will cause a failure during PDF export.	Resolved	Fixed	0
	CONF-15935	Improve upgrade guide with some performance tips	Resolved	Fixed	0
	CONF-15397	Instant Messenger Macro XSS Vector	Resolved	Fixed	0
	CONF-14523	Insert macro form for macros that don't provide parameter metadata	Resolved	Fixed	0
	CONF-14531	Add filtering/search box to macro browser	Resolved	Fixed	0
	CONF-12573	Enable automatic HTML encoding by default	Resolved	Fixed	0
	CONF-14559	RTE does not deal with <foo> correctly.	Resolved	Fixed	0
	CONF-16495	User cannot update the page which contains pagetree macro when it is updated using XMLRPC	Resolved	Fixed	0
	CONF-2333	PDF Export created on Linux server cannot be opened in Adobe AcrobatReader.	Resolved	Fixed	0
	CONF-14431	Write Anti-XSS documentation for plugin developers	Resolved	Fixed	0
	CONF-6864	When a note macro is split between pages in a PDF export the second part loses its style	Resolved	Fixed	0
	CONF-14651	Improve cluster performance by avoiding cache puts for already-cached objects	Resolved	Fixed	0
	CONF-14689	Comment edits do not trigger email notification	Resolved	Fixed	0
	CONF-14705	Workaround for authentication login prompt when editing Microsoft Office files	Resolved	Fixed	0
	CONF-11500	Disabling daily Backups on one node in a cluster does not reflect on other nodes in the cluster	Resolved	Fixed	0
	CONF-14707	Better Cache/Queue management for the Office Connector	Resolved	Fixed	0
	CONF-14213	Correct MySQL collation for Confluence database.	Closed	Fixed	0
	CONF-14981	Viewfile macro with a word document throws java.lang.ArrayIndexOutOfBoundsException: -10	Resolved	Fixed	0
	CONF-14073	Deprecate ThumbnailInfoFactory	Resolved	Fixed	0
	CONF-14984	Update green arrows on the blog calendar display	Resolved	Fixed	0
	CONF-14969	Remove "Open Web Folder" link from attachment lists	Resolved	Fixed	0
	CONF-14735	Update Usernames in Notifications table to match LDAP users with different case		Fixed	0

			Resolved		
	CONF-13490	Attachments can be moved to a page which the user cannot access	Resolved	Fixed	0
	CONF-12264	Extra newline appears when converting to H2 in a new page	Resolved	Fixed	0
	CONF-15061	Extract build information responsibilities from GeneralUtil	Resolved	Fixed	0
	CONF-13270	Fix slow start-up caused by AspectJ hitting the plugin classloader too often	Resolved	Fixed	0
	CONF-14294	Check cache sizes and increase if necessary	Resolved	Fixed	0
	CONF-15267	Import Pages is not restricted to system admins	Resolved	Fixed	0
	CONF-13776	Default panel background colour does not apply in PDF export	Resolved	Fixed	0
	CONF-15059	Deprecate com.atlassian.confluence.cache.CacheManagerKeys	Resolved	Fixed	0
	CONF-11749	Uploaded pdf fonts break when the home directory is moved.	Resolved	Fixed	0
	CONF-15073	Add web-items plugin point on Dashboard for navigation items similar to the Page and BlogPost pages	Resolved	Fixed	0
	CONF-14688	Improve performance of HtmlSafe annotation checking	Resolved	Fixed	0
	CONF-9038	PDF Generation fails when bgcolor wiki markup is used	Resolved	Fixed	0
	CONF-15385	Confluence runs out of file handles when running the did you mean index job in our performance build	Resolved	Fixed	0
	CONF-15119	Update to atlassian-plugins 2.2	Resolved	Fixed	0
	CONF-15157	viewfile macro does not work for non printable PDFs	Resolved	Fixed	0
	CONF-15484	Cache personal information objects to improve performance	Resolved	Fixed	0
	CONF-15396	Improve performance of TransactionalCacheFactory	Resolved	Fixed	0
	CONF-5782	Selecting several lines and applying the heading style produces incorrect markup	Resolved	Fixed	0
	CONF-15251	Remove Table Layout from Profile Pages	Resolved	Fixed	0
	CONF-13858	Rendering of insert table dialog poor	Resolved	Fixed	0
	CONF-14015	500page.jsp always reports Confluence User as anonymous	Resolved	Fixed	0
	CONF-15531	BuildNumberCondition would make it easier to control display of web items in plugins	Resolved	Fixed	0
	CONF-15534	Add a new search filter to allow in place pagination of search results	Resolved	Fixed	0
	CONF-15550	Remove dependency on atlassian-bucket	Resolved	Fixed	0
	CONF-14085	Document confluence.jmx.disabled system property	Closed	Fixed	0
	CONF-15566	Upgrade atlassian xwork to include XWORK-16 fix		Fixed	0

			Resolved		
	CONF-14093	Implement Hibernate's ReadWriteCache in a non blocking way	Resolved	Fixed	0
	CONF-14011	ConfluenceVelocityResourceCache causes stack overflow if "A FakeSerializableResource has been serialized"	Resolved	Fixed	0
	CONF-14445	Doc update: Component plugin modules are not treated as StateAware in OSGi plugins	Resolved	Fixed	0
	CONF-14081	Allow disabling of JMX via a system property	Resolved	Fixed	0
	CONF-13865	Admin users should be able to see content in all spaces when using QuickNav	Resolved	Fixed	0
	CONF-13945	History version comparison should never be a POST	Resolved	Fixed	0
	CONF-14098	Remove unncessary sychronization on Hibernates UpdateTimestampsCache	Resolved	Fixed	0
	CONF-14121	Use putAll() instead of put() on the Coherence cache	Resolved	Fixed	0
	CONF-14422	All attachments versions loaded when you query for just one	Resolved	Fixed	0
	CONF-13911	Using heading formats in lists and tables inserts unwanted new lines @ cursor	Resolved	Fixed	0
	CONF-13917	Improve performance by caching Userprofile-Pictures and Space-logos	Resolved	Fixed	0
	CONF-13893	Add context menus for the RTE	Resolved	Fixed	0
	CONF-15703	One pagetree-macro in a pagelayout prevents another pagetree-macro on a page from working correctly	Resolved	Fixed	0
	CONF-13775	Creating headings in the RTE & Firefox causes empty headings also be created	Resolved	Fixed	0
	CONF-15253	Jsonator does not encode string according to the spec	Resolved	Fixed	0
	CONF-14823	Dynamic tasklist, widget macros and plugin repository don't work after upgrade to Confluence 2.10.1 or 2.10.2	Resolved	Fixed	0
	CONF-7633	Error in PDF Export after the first call of a sql-macro producing a table output	Resolved	Fixed	0
	CONF-15643	Sensible default settings for Remote API and Public Sign Up	Resolved	Fixed	0
	CONF-14938	viewfile causing OutOfMemoryError for Powerpoint files that contain embedded PICT file (Mac OS X)	Resolved	Fixed	0
	CONF-13469	Replace LuceneSmartListManager with V2SearchSmartListManager and ensure backwards compatibility	Resolved	Fixed	0
	CONF-12191	Page title font should match the header fonts when printing	Resolved	Fixed	0
	CONF-13215	Cancelling content-creation partially broken and inconsistent	Resolved	Fixed	0
	CONF-14059	synchronization inside ognl.EvaluationPool is a performance bottleneck in high load instances	Resolved	Fixed	0
	CONF-11576	Add a remove link button to the Rich Text editor	Resolved	Fixed	0
	CONF-14434	Display the most recent edit by a user, even if other users edited the same pages afterwards	Resolved	Fixed	0
	CONF-15345	Follow macro does not render the 'Full theme' (large icons) when previewing it		Fixed	0

		either in the Macro Browser or on the page itself in IE7.	Resolved		
	CONF-15720	Include-Excerpt macro doesn't include excerpts until the page has been viewed	Resolved	Fixed	0
	CONF-9615	Cluster information page should tell you which node you're currently on.	Resolved	Fixed	0
	CONF-13004	Attachment directory needs to be hashed to avoid filesystem limitations	Resolved	Fixed	0
	CONF-14272	Remove deprecated ThumbnailManager classes	Resolved	Fixed	0
	CONF-14257	Avatars should be served with proper cache control headers	Resolved	Fixed	0
	CONF-15660	Big dimension of image file crashes Confluence when we tries to export the page to PDF.	Resolved	Fixed	0
	CONF-13846	Wildcard and range queries always converted to lower case	Resolved	Fixed	0
	CONF-13844	Allow plugins to indicate whether they should be subject to automatic HTML encoding	Resolved	Fixed	0
	CONF-14343	Add and display macro parameter descriptions in macro browser	Resolved	Fixed	0
	CONF-14273	Reduce DB load for loading pages	Resolved	Fixed	0
	CONF-11371	ConfluenceEncodingFilter removes Unicode dashes, quotes, bullets and ellipses	Resolved	Fixed	0
	CONF-14298	Last modifier and creator username gets set to null when modifying CEOs from outside a request	Resolved	Fixed	0

Above are the issues resolved in Confluence 3.0, ordered by number of votes. For the full details of the fixes, improvements and new features, please take a look at our issue tracker.

Confluence 2.10.4 Release Notes

Please skip this version and upgrade to Confluence 3.0 or later
Please be advised that there is a known bug in Confluence 2.10.4, whereby the `ehcache.xml` file is completely missing and there is no 'Cache Statistics' link on the Administration console. If the link `http://.../admin/cachestatistics.action` is directly accessed, a blank page is returned with no cache statistics details. Please upgrade to the latest version of **Confluence 3.x**, as indicated in the release notes. If you cannot upgrade to Confluence 3.x, please contact our support team.

Confluence 2.10.4 is the **standard edition** version of Confluence 2.10.x. This version of Confluence is equivalent to Confluence version 2.10.3, but it does not include Oracle's Coherence technology. Due to licensing changes with Oracle over the Coherence clustering and distributed caching technology, Atlassian is no longer able to distribute this technology to customers without a Confluence clustered license.

Therefore, Atlassian has released 'standard editions' for each previous major version of Confluence back to 2.6, which includes Confluence versions 2.10.4, 2.9.3, 2.8.3, 2.7.4 and 2.6.3. Confluence 2.10.4 will be the only Confluence 2.10.x version available to customers with a non-clustered Confluence license.

For its caching functionality, Confluence 2.10.4 utilises Ehcache. This functionality was provided by the Coherence technology in previous Confluence 2.10.x versions.

For more information about the features, updates and fixes in Confluence 2.10.3, please refer to the Confluence 2.10.3 Release Notes.

Upgrading from a Previous Version of Confluence

Upgrading Confluence should be fairly straightforward. Please read the Confluence 2.10.3 Upgrade Notes. However, if you have customised the cache settings in your installation of Confluence (e.g. for performance reasons), then please read the Confluence 3.0.1 Upgrade Notes for important information about transferring your caching layer customisations from Coherence to Ehcache.

We *strongly recommend* that you back up your `confluence.home` directory and database before upgrading.

Confluence 2.10.3 Release Notes

15 April 2009

Confluence 2.10.3 is a recommended upgrade which fixes some security flaws and other issues.

Please refer to the security advisory for details of the security vulnerabilities, risk assessments and mitigation strategies.

General Fixes

A bug was identified, whereby viewing or editing restrictions could not be assigned to a page, whose parent page contained an apostrophe in its title and also possessed existing page restrictions. This bug has now been fixed.

When a user is restricted from viewing a page, Confluence presents them with a more informative **Access Denied** error rather than a general **Page Not Found** error.

When the {gallery} macro is used on a page with no parameters or image attachments, it would render into an error in HTML or PDF exports. This issue has now been fixed.

An issue was identified whereby under certain circumstances, clicking on a page's or blog's thumbnail image to expand it would result in a Runtime Error in Internet Explorer versions 6 and 7. This issue has now been fixed.

Widget Connector Plugin

Several new features have been added to the Widget Connector Plugin packaged with Confluence 2.10.3, including support for new widget, video and micro-blogging sites. Other supported features include Google Calendar and the Wufoo HTML Form Builder. For more information on how to add these features to your Confluence page or blog, refer to Widget Macro.

Episodic made changes to the format of IDs they designate for all new videos, allowing them to be alphanumeric rather than solely numeric. The Widget Connector plugin has been updated to support this new URL format.

Engine Room Fixes

An issue was identified in Confluence's PDF Export feature that could result in memory leaks. These in turn may have affected the performance and stability of Confluence instances. This issue has now been fixed.

A few other issues were identified which under certain or specific circumstances, could affect the stability of Confluence. However, these have now been fixed.

There's a complete list of fixes below. Click a specific issue to see details of the fix.

Don't have Confluence 2.10 yet?

Take a look at the new features and other highlights in the Confluence 2.10 Release Notes.

Download Latest Version

Upgrading from a Previous Version of Confluence

Upgrading Confluence should be fairly straightforward. Please read the Confluence 2.10.3 Upgrade Notes. We *strongly recommend* that you back up your `confluence.home` directory and database before upgrading.

Updates and Fixes in this Release

JIRA Issues (31 issues)

Type	Key	Summary	Priority	Status	Resolution
	CONF-17439	Fix the 2.10 patch for link insertion point in IE RTE		Resolved	Fixed
	CONF-15541	Previewing a page does not display long content when Confluence is embedded in a frame		Resolved	Invalid
	CONF-15538	em dash in the attachment causes Confluence's crash		Resolved	Duplicate
	CONF-14988	SAXParser memory leaks		Resolved	Fixed
	CONF-14922	com.sun.pdfview.font.Type1CFont.readCommand(Type1CFont.java:357) consumes 100% CPU, blocks all other pdf threads		Resolved	Fixed
	CONF-14849	Discarding a draft and refreshing the page results in an error		Resolved	Fixed
	CONF-14775	Add new filter and response wrapper to prevent header injection attacks		Resolved	Fixed
	CONF-14753	XSS vulnerability can be exploited with the Page Index macro		Resolved	Fixed

	CONF-14704	Improper sanitisation of attachment filenames allows header injection		Resolved	Fixed
	CONF-14537	Can not alter the permissons on a page, if the parent page has permissions and the parent has an apostrophe in the page name. Similar to CONF-10717		Resolved	Fixed
	CONF-14510	Fix upgrade tasks that access the database through a connection from the HibernateSession which is later garbage collected		Resolved	Fixed
	CONF-14493	Password is being logged for 500 errors		Resolved	Fixed
	CONF-14386	Empty gallery macro throws error in HTML and PDF export		Resolved	Fixed
	CONF-14337	XSS in the Widget Connector		Resolved	Fixed
	CONF-14326	Site search query box and submit button too small with Left Nav theme and Clickr theme		Resolved	Fixed
	CONF-14310	Studio plugins: Using components not available to plugins		Resolved	Fixed
	CONF-14178	System error when adding users to a group if the group name contains a space		Resolved	Fixed
	CONF-14127	New evaluation expiry message		Resolved	Fixed
	CONF-14102	anti-XSS mode breaks RTE-by-default editing and view page source		Resolved	Fixed
	CONF-14092	AspectJ caused CAC crash 2008-12-31		Resolved	Fixed
	CONF-13785	Errors when exporting demonstration space after migration to 2.10-m8		Resolved	Fixed
	CONF-13771	HTML export results in NPE from requireResource velocimacro when used by plugins		Resolved	Fixed
	CONF-13494	SOAP calls break when trying to write {tasklist} macro in contents		Resolved	Fixed
	CONF-13331	Search box display issue when not using default theme in IE and FireFox		Resolved	Fixed
	CONF-13316	NullPointerException thrown if user cannot move attachments		Resolved	Fixed
	CONF-13063	Confluence Advanced Search Pane search text box is flattened vertically in 2.91 using Left Navigation Theme		Resolved	Fixed
	CONF-12835	Error in Java Script when adding an image thumbnail and clicking on it		Resolved	Fixed
	CONF-12366	Image width and height are to 32 by wysiwyg editor for missing image files.		Resolved	Fixed
	CONF-11552	height and width of embedded movies is changed to 32 by Rich Text Editor		Resolved	Fixed
	CONF-9239	"Page not found" (404) is displayed to users without view permissions, should be "Access denied"		Resolved	Fixed
	CONF-8130	image link breaks if the target is renamed		Resolved	Fixed

Click here to open a report on http://jira.atlassian.com for Resolved or Closed issues in Confluence 2.10.3.

Confluence 2.10.3 Upgrade Notes

Confluence 2.10.3 is a recommended upgrade which fixes some security flaws as well as other bugs. You'll find details of the fixes in the release notes.

Upgrade Procedure

If you are already running a version of Confluence, please follow these instructions to upgrade to the latest version:

1. Before you upgrade, we strongly recommend that you **back up your `confluence.home` directory and database**.
2. If your version of Confluence is earlier than 2.10, read the release notes and upgrade guides for **all releases** between your version and the latest version. In particular:
   - Please read the 2.10 upgrade notes.
   - If you are upgrading from 2.1 or earlier, please also read the 2.2 release notes.
3. Download the latest version of Confluence.
4. Follow the instructions in the Upgrade Guide.

Confluence 2.10.2 Release Notes

18 February 2009

Confluence 2.10.2 is a recommended upgrade which fixes a security flaw and other issues.

Please refer to the security advisory for details of the security vulnerability, risk assessment and mitigation strategies.

Rich Text Editor fixes

A bug was identified in Confluence version 2.10, which affected the Rich Text Editor in Internet Explorer browsers. This bug prevented new content that was entered into a single line break between existing chunks of text from being saved. For instance, if a user placed their cursor on a single blank line (or line break) between two existing sections of text, entered new text (consisting of any number of paragraphs) and saved it, the new text would not be saved. This issue has now been fixed.

Another issue was identified in the Rich Text Editor which made text appearing immediately under an image concatenate with the image's Wiki Markup when the page content was saved. This prevented the image from being displayed. While inserting an additional line break between the image and text provided a workaround, this issue has now been fixed and this workaround is no longer required.

Content by Label Macro fixes

A bug was identified in the Content by Label Macro that prevented its sort parameter from functioning correctly. However, this has now been fixed. Furthermore, the performance of the Content by Label Macro has been improved.

With the release of Confluence 2.10, the default behaviour of the Content by Label Macro's **space** parameter was modified to **@self**. Due to customer feedback and popular demand, however, we reverted this parameter's default behaviour back to **@all**.

Other fixes

A stability issue was identified in Confluence version 2.10 which has now been fixed. However, a minor side effect has been identified which can result in some superfluous non-breaking spaces not being removed from the end of lines, when either saving a page or switching from the Rich Text Editor to Wiki Markup modes.

A bug was identified in Confluence version 2.10 that prevented Confluence from playing SWF files with Flash Player 10 in Internet Explorer. This has now been fixed.

Some issues were identified with the Code Block Macro in Confluence 2.10, which resulted in the removal of white space within a code block when switching from Wiki Markup to Rich Text Editor modes. However, a fix was introduced to mitigate these issues.

An issue was identified when Viewing Pages Alphabetically, which prevented Confluence from listing pages by specific letters of the alphabet when 1,000 or more pages started with any one letter of the alphabet. This issue has now been fixed.

When writing content in Wiki Markup, URLs containing accented characters now render correctly into links. Additionally, Confluence's French and German product interface translations have been improved.

There's a complete list of fixes below. Click a specific issue to see details of the fix.

Don't have Confluence 2.10 yet?

Take a look at the new features and other highlights in the Confluence 2.10 Release Notes.

Download Latest Version

Upgrading from a Previous Version of Confluence

Upgrading Confluence should be fairly straightforward. Please read the Confluence 2.10.2 Upgrade Notes. We *strongly recommend* that you back up your `confluence.home` directory and database before upgrading.

Updates and Fixes in this Release

JIRA Issues (47 issues)					
Type	**Key**	**Summary**	**Priority**	**Status**	**Resolution**

	CONF-14840	Update Java Policy Security page for 2.10.x		Resolved	Fixed
	CONF-14744	Users may have to rebuild indices after performing migration		Resolved	Fixed
	CONF-14451	Attachments on Comments will not be indexed correctly		Resolved	Fixed
	CONF-14449	Attachment Manager allows for storing data against history versions of a page		Resolved	Fixed
	CONF-14429	"java.lang.NullPointerException at java.util.regex.Matcher.getTextLength(Unknown Source)" when importing a .doc file		Resolved	Fixed
	CONF-14424	Append new line character to end of code macro to mitigate loss of spaces in code macro		Resolved	Fixed
	CONF-14333	Internal system information is shown on the GUI for list of restricted pages		Resolved	Fixed
	CONF-14311	Document new Web UI plugin location		Resolved	Fixed
	CONF-14275	HTTP Header Injection vulnerability: os_destination value not properly escaped when used as redirect location		Resolved	Fixed
	CONF-14255	Disallow adding a child page under a page it cannot view in default theme		Resolved	Fixed
	CONF-14152	stylesheet.page.description in Colour Screen help tips		Resolved	Fixed
	CONF-14138	Confluence installed as a service seems to be started but the server start actually fails		Resolved	Fixed
	CONF-14136	Typo in screenshot for attachment migration documentation		Resolved	Fixed
	CONF-14129	Exporting to Word in anti-XSS mode will include html tags		Resolved	Fixed
	CONF-14119	Typo in warning message when changing Attachment Storage Configuration		Resolved	Fixed
	CONF-14101	HTML export skips file attachements with space characters (blanks) in their name		Resolved	Fixed
	CONF-14094	Confluence Support Request Improvements		Resolved	Fixed
	CONF-14069	Remote API rename feature should not allow duplicate titles		Resolved	Fixed
	CONF-14048	Incorrect German Translation for "Sign Up"		Resolved	Fixed
	CONF-14047	Remove contextPath from drop-down.js		Resolved	Fixed
	CONF-14023	Content By Label macro now defaults to showing content from the current space only		Resolved	Fixed
	CONF-14020	Upgrade to latest atlassian-gzipfilter to fix remote APIs		Resolved	Fixed
	CONF-14018	Rich Text mode doesn't save changes between linebreaks in IE		Resolved	Fixed
	CONF-14006	Get 500 when trying to communicate to confluence via trusted apps.		Resolved	Fixed
	CONF-14002	HTML export: Attachments are missing filetypes/extensions		Resolved	Fixed
	CONF-13974	Confluence does not play SWF file in Internet Explorer when using Flash Player 10 (it works when using flash player 9)		Resolved	Fixed

	CONF-13965	CPU Spike: DefaultWysiwygConverter.convertXHtmlToWikiMarkup risks blowing the stack		Resolved	Fixed
	CONF-13937	Incorporate more German and French Translations		Resolved	Fixed
	CONF-13825	NullPointerException when thumbnail cannot be generated		Resolved	Fixed
	CONF-13824	Images and text concatenated by RTE save operation, also causing image markup to be escaped.		Resolved	Fixed
	CONF-13707	Inserting image of size bigger than attachment limit yields "Page Not Found" error		Resolved	Fixed
	CONF-13694	PATCH: Confluence sets incorrect content type for css files bundled in war		Resolved	Fixed
	CONF-13677	Cannot update the name of a space using storeSpace(String in0, RemoteSpace in1)		Resolved	Fixed
	CONF-13588	Users should not have to click Refresh to update their license count after removing a user		Resolved	Fixed
	CONF-13435	mismatched tag. Expected:		Resolved	Fixed
	CONF-13119	Pressing the search button always appends previous search query to proposed new query.		Resolved	Fixed
	CONF-13028	Paging to the ninth result set of any pagination displays "one" twice in pagination choices		Resolved	Fixed
	CONF-12809	Anonymous Users can still see links to the people directory even when disabled		Resolved	Fixed
	CONF-12707	Accented characters (all non-ascii?) cause link rendering to fail		Resolved	Fixed
	CONF-12475	RSS feeds created by Space->Advanced ->Rss Feeds do not respect the base server url		Resolved	Fixed
	CONF-11795	Export With {toc} Causes NullPointerException		Resolved	Fixed
	CONF-10174	Pages created with restrictions show restrictions as created by Anonymous		Resolved	Fixed
	CONF-9715	Confluence crash can result in empty confluence.cfg.xml file.		Resolved	Fixed
	CONF-9635	System Error when accessing via a TinyURL that ends with a dash		Resolved	Fixed
	CONF-9213	PageManager can create Duplicate Pages		Resolved	Fixed
	CONF-8865	Improve content-by-label macro performance		Resolved	Fixed
	CONF-5322	Add sorting options to the contentbylabel macro		Resolved	Fixed

Click here to open a report on http://jira.atlassian.com for Resolved or Closed issues in Confluence 2.10.2.

Confluence 2.10.2 Upgrade Notes

Confluence 2.10.2 is a recommended upgrade which fixes some security flaws as well as other bugs. You'll find details of the fixes in the release notes.

Upgrade Procedure

If you are already running a version of Confluence, please follow these instructions to upgrade to the latest version:

1. Before you upgrade, we strongly recommend that you **back up your `confluence.home` directory and database**.
2. If your version of Confluence is earlier than 2.10, read the release notes and upgrade guides for **all releases** between your version

and the latest version. In particular:
- Please read the 2.10 upgrade notes.
- If you are upgrading from 2.1 or earlier, please also read the 2.2 release notes.

3. Download the latest version of Confluence.
4. Follow the instructions in the Upgrade Guide.

Confluence 2.10.1 Release Notes

7 January 2009

Confluence 2.10.1 is a recommended upgrade which fixes some security flaws and other issues.

Please refer to the security advisory for details of the security vulnerabilities, risk assessment and mitigation strategies.

A bug was identified in Confluence version 2.10.0 that made the label parameter of the Content by Label Macro case-sensitive. This resulted in link breakages where differences existed in case usage between {contentbylabel} label parameter values and the page labels they referenced. However, this issue has now been fixed.

Another issue was identified in Confluence version 2.10.0 that prevented specific Confluence components (for example, those of the Rich Text Editor) from loading correctly when running Confluence behind certain proxy server configurations. For example, this may have become apparent when running Confluence behind an Apache HTTP Server using the mod_proxy connection module. This issue has also been fixed in Confluence version 2.10.1, along with a number of other issues.

There's a complete list of fixes below. Click a specific issue to see details of the fix.

Don't have Confluence 2.10 yet?

Take a look at the new features and other highlights in the Confluence 2.10 Release Notes.

Download Latest Version

Upgrading from a Previous Version of Confluence

Upgrading Confluence should be fairly straightforward. Please read the Confluence 2.10.1 Upgrade Notes. We *strongly recommend* that you back up your `confluence.home` directory and database before upgrading.

Updates and Fixes in this Release

JIRA Issues (23 issues)

Type	Key	Summary	Priority	Status	Resolution
	CONF-14329	viewfile macro does not work in case of umlauts in attachment name		Resolved	Fixed
	CONF-14108	Recently Updated Macro Documentation Needs update		Resolved	Fixed
	CONF-14049	Clarify source code shipped with product		Closed	Fixed
	CONF-14033	Documentation on support for platforms not x86 is unclear		Resolved	Fixed
	CONF-14032	Office Connector plugin shows vendor as unknown in the Plugin Manager		Resolved	Fixed
	CONF-14026	Upgrade to Atlassian Plugins 2.1.3		Resolved	Fixed
	CONF-14015	500page.jsp always reports Confluence User as anonymous		Resolved	Fixed
	CONF-14014	Word import with Office Connector can overwrite existing content without permission		Resolved	Fixed
	CONF-13998	In SharedMode the SpaceGroupFilter will filter out all search results if the user is not an administrator		Resolved	Fixed
	CONF-13987	Confluence 2.10 Standalone Default SSL Implementation Not Compatible With Tomcat 6		Resolved	Fixed
	CONF-13969	Content By Label macro now uses case sensitive label names		Resolved	Fixed

	CONF-13942	Rich Text editor does not display when Confluence is running behind Apache		Resolved	Fixed
	CONF-13934	The "purchasing it" link at the footer of a Confluence page (under an Evaluation Licence) is incorrect.		Resolved	Fixed
	CONF-13930	confluence-2.10-std README file needs to correct the jdk version.		Resolved	Fixed
	CONF-13900	Plugin being disabled after being upgraded or re-enabled		Resolved	Fixed
	CONF-13892	Plugin state conflict when plugin installed at same time as plugin state is checked.		Resolved	Fixed
	CONF-13865	Admin users should be able to see content in all spaces when using QuickNav		Resolved	Fixed
	CONF-13846	Wildcard and range queries always converted to lower case		Resolved	Fixed
	CONF-13702	Session must not be invalidated on logout		Resolved	Fixed
	CONF-13524	Left navigation theme displays "news operations" instead of "page operations" when viewing the left nav main menu categories		Resolved	Fixed
	CONF-9693	Exporting a space that is associated to a SpaceGroup causes an error (lazy init exception)		Resolved	Fixed
	CONF-9033	change-history macro produces a IndexOutOfBoundsException when previewed on create page/news		Resolved	Fixed
	CONF-8643	Rationalise FAQ		Resolved	Fixed

Confluence 2.10.1 Upgrade Notes

Confluence 2.10.1 is a recommended upgrade which fixes some security flaws as well as other bugs. You'll find details of the fixes in the release notes.

Upgrade Procedure

If you are already running a version of Confluence, please follow these instructions to upgrade to the latest version:

1. Before you upgrade, we strongly recommend that you **back up your `confluence.home` directory and database**.
2. If your version of Confluence is earlier than 2.10, read the release notes and upgrade guides for **all releases** between your version and the latest version. In particular:
   - Please read the 2.10 upgrade notes.
   - If you are upgrading from 2.1 or earlier, please also read the 2.2 release notes.
3. Download the latest version of Confluence.
4. Follow the instructions in the Upgrade Guide.

Confluence 2.10 Release Notes

Confluence 2.10 fixes some security flaws. Please refer to the security advisory for details of the security vulnerabilities, risk assessment and mitigation strategies.

3 December 2008
With great pleasure, Atlassian presents **Confluence 2.10**.

Confluence 2.10 is a major release which presents a number of new features and enhancements. With Confluence 2.10 we introduce the new **Widget Connector**, an easy way to embed multi-media content from all over the web directly into your Confluence page. Add Youtube videos, Flickr slide shows and Google Gadgets just to name a few. Now that the improved **Office Connector** is bundled with Confluence 2.10, you can also display the contents of attached documents, spreadsheets, presentations, and PDFs directly on any page. Also, with the new Office Connector, you can now view the contents of attachments using the new 'View' feature from the attachments page, or from a search.

In Confluence 2.10, finding all your content in Confluence is a lot easier. With **Quick Navigation** just start typing in the search box and immediately see suggested results. With **OpenSearch** autodiscovery, you can add Confluence search to your Firefox or IE7 search box with

just one click. And with **'Did You Mean'**, you no longer have to worry about mistyping search terms since Confluence now suggests corrections to misspelled words. Take a look at the **new search results page**, you'll find the right page, attachment or person far more easily.

You can better control your avatar with new **profile picture cropping** and spruce up your personal space with **custom stylesheets.** Page editing has become more reliable through an upgrade to our **rich text editor** which now supports Safari and has a new styles dropdown. JIRA users will be delighted to use our new **JIRA Issues Macro** with drag and drop column sorting, paging and support for custom fields.

Confluence administrators will benefit from improved user management, a broader range of supported wikis for the **Universal Wiki Converter** and some good performance improvements.

Please see our **overview video** to see a demonstration of the new features in Confluence 2.10.

Highlights of this Release:

- Introducing the Widget Connector
- Improved Office Connector Now Bundled
- Introducing Quick Navigation
- 'Did You Mean', OpenSearch and More
- Custom Stylesheets for Confluence Spaces
- Updated JIRA Issues Macro with Custom Fields and Dynamic Display
- Enhanced User and Group Management
- Upgraded Rich Text Editor
- Universal Wiki Converter now with SharePoint Import and More
- Improved Activity Macros
- Plugin Framework 2
- More than 250 Fixes and Improvements

Responding to your Feedback:

720+ votes satisfied

11 new feature requests implemented



- Thank you for all your issues and votes. Keep logging, to help us keep improving!
- Below is a list of the highlights in this release.
- Attached is the full list of issues resolved in this release.

Upgrading from a previous version of Confluence
• Upgrading Confluence should be fairly straightforward. ***We strongly recommend that you back up your Confluence Home directory and your database before upgrading.*** • Please refer to the Confluence 2.10 Upgrade Notes for further essential information about **plugins** and other factors affecting your upgrade.

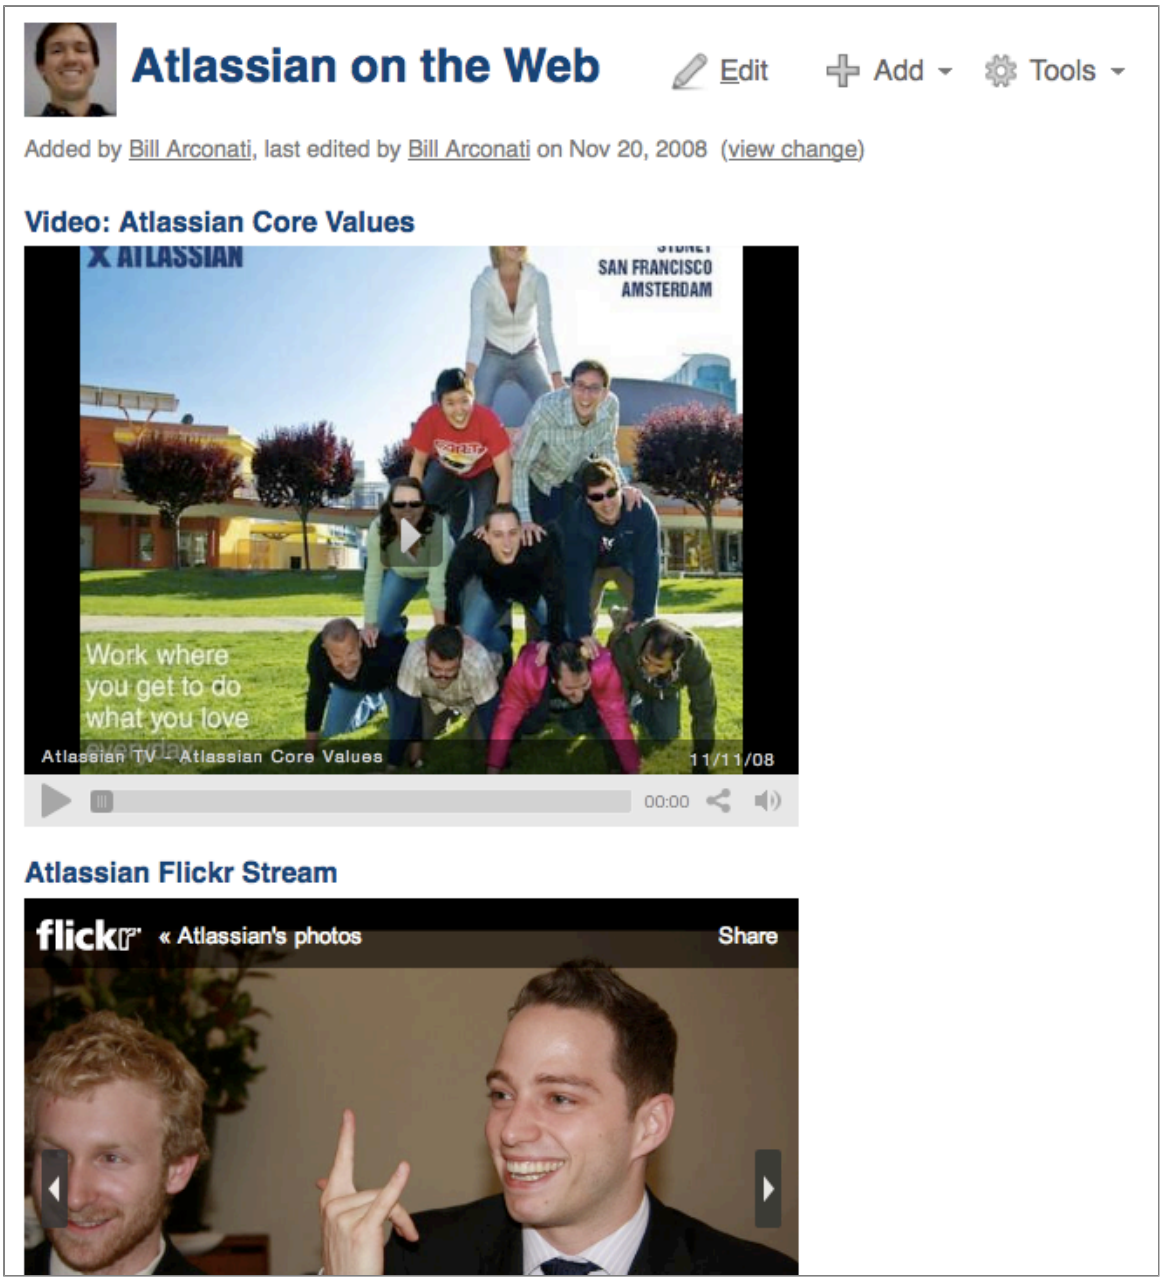
Highlights of Confluence 2.10



Introducing the Widget Connector

Now you can embed multi-media content from other web sites into a Confluence page. Bring your wiki page to life with Google Gadgets, videos, slide shows, Twitter messages and more. Simply type the word '{widget}' and give it the web address of the content you want to embed.

- Gadgets: Google Gadgets.
- Videos: YouTube, MySpace Video, Google Video, Episodic, Vimeo, Metacafe, blip.tv, Viddler.
- Photos and images: Flickr, Skitch.com.
- Micro-blogging: Twitter, FriendFeed.
- Documents and presentations: SlideShare, SlideRocket, Scribd, presentations on Google Docs.
- Our documentation shows you what to do.

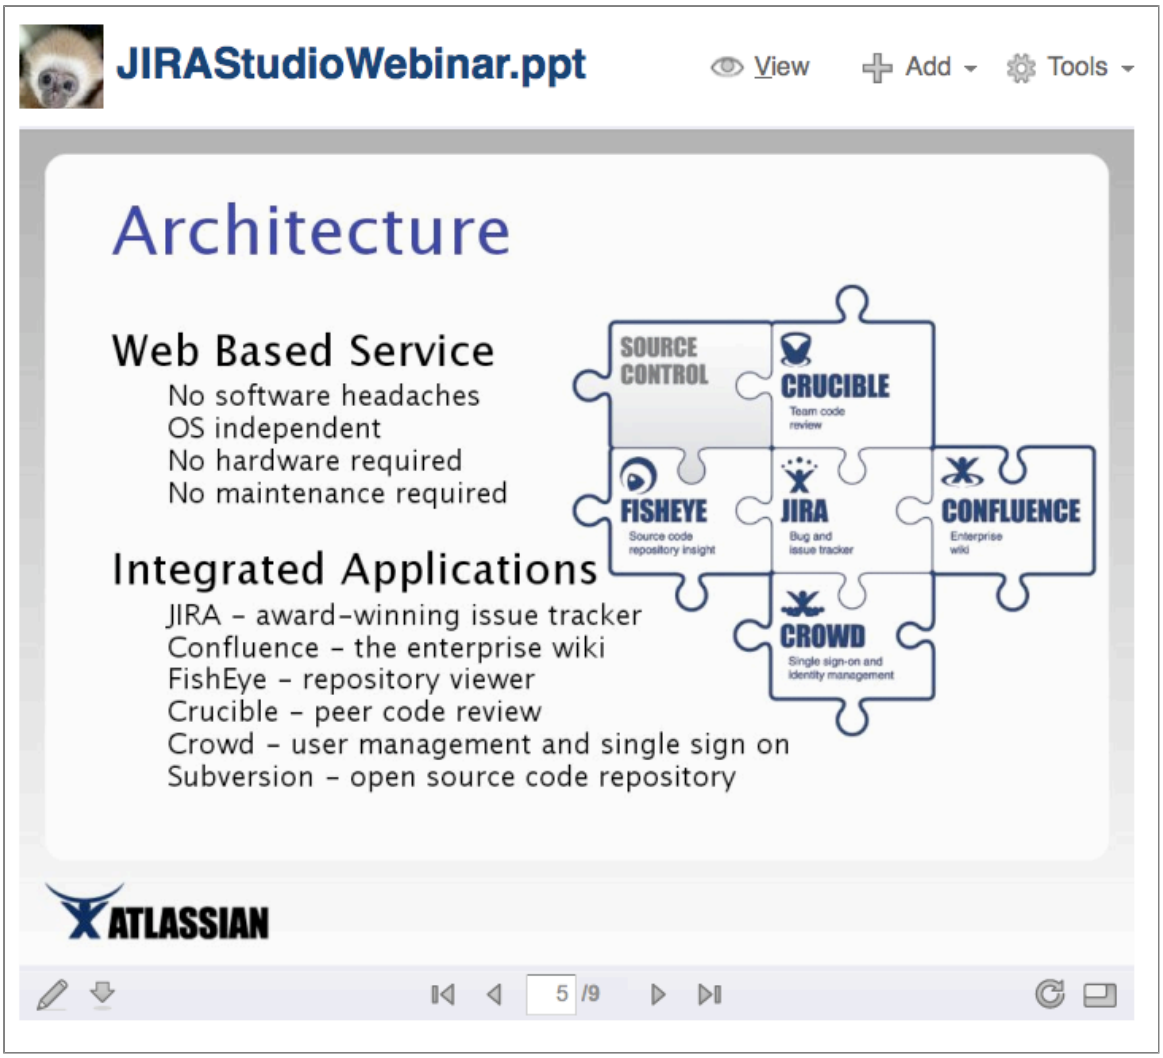


2

Improved Office Connector Now Bundled

The Office Connector is shipped as part of Confluence 2.10. There is no need to install it separately. Use the Office Connector to create and edit rich content for Confluence using Microsoft Office or OpenOffice.

- The new 'View' feature lets you view Office documents from the Search results, from the Attachments page and from a list of attachments displayed by the Attachments macro. You do not need to have Office installed on your machine to view an Office document in Confluence.
- We have also fixed a number of bugs in the Office Connector. See the list of fixes.
- See all the features of the Office Connector in our documentation.

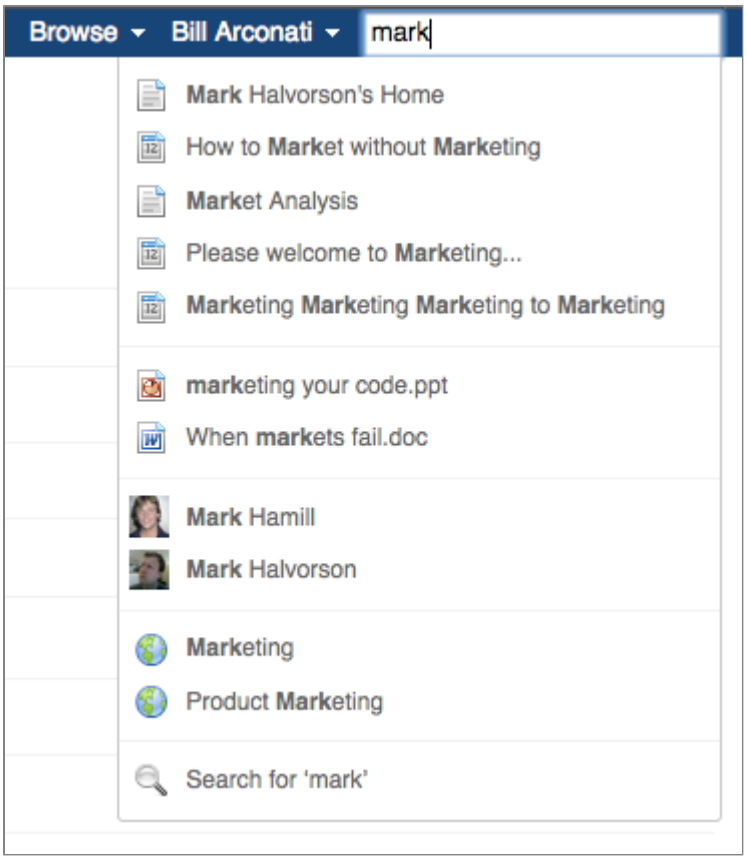


3

Introducing Quick Navigation

Confluence's search box now offers a quick navigation feature, the fastest way to find content in Confluence.

- Start typing your search term. Confluence matches titles as you type, showing a quickly-adjusting list of pages, news items, personal profiles, attachments and so on.
- The matching items are grouped by content type so that you can quickly find the type you want.
- When the matching item is a person's name, their profile picture appears next to their name in the list.
- Still not found what you are looking for? Click the 'Search' option at the bottom of the list to do a full search.

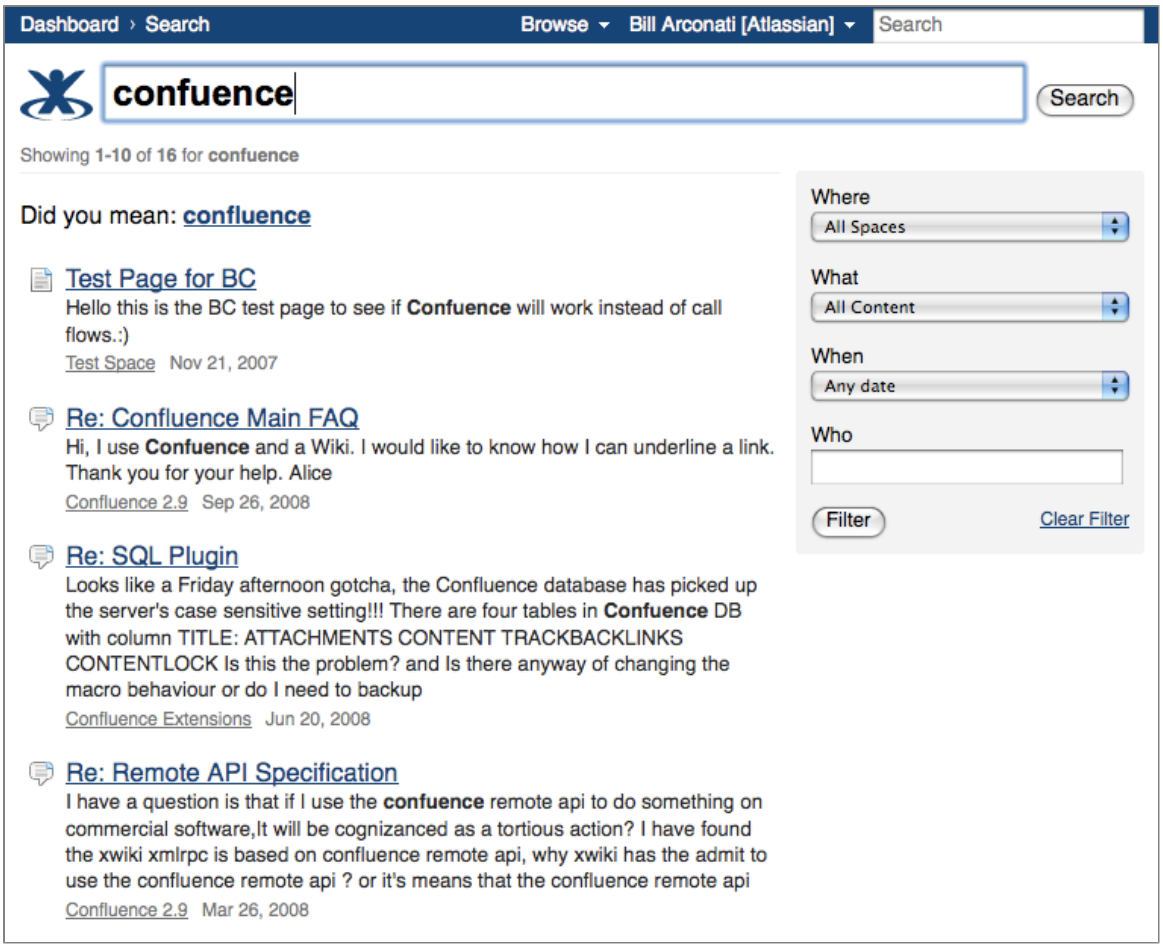




'Did You Mean', OpenSearch and More

Confluence 2.10 includes a number of improvements to the Search functionality.

- Find what you're looking for even if you mistype your search term. Confluence's new 'Did you mean' feature analyses your search term and suggests an alternative spelling to give you more relevant search results.
- You can now search Confluence from the convenience of your browser's search box, if you are using Firefox or IE7. Just add your Confluence site as a search provider, via the dropdown menu next to the browser's search box. This is because Confluence now supports the autodiscovery part of the OpenSearch standard.
- We have also improved the layout of the Search screen so that it is easier to read. The titles now stand out more and a longer extract is shown.
- If the matching item is a person, the profile picture and other profile information appear in the search results.
- If the matching item is an attached Office document, a new 'View' link allows you to view the document online.
- There is a search box at the bottom as well as the top of the screen, so you do not need to scroll up to enter a new search term.
- The 'Clear Filter' link replaces the old 'Clear Search', and now just clears the filter criteria instead of the entire search.
- When ranking the search results, Confluence now gives slightly higher priority to pages created recently.
- Take a look at our documentation for full details of the new features.

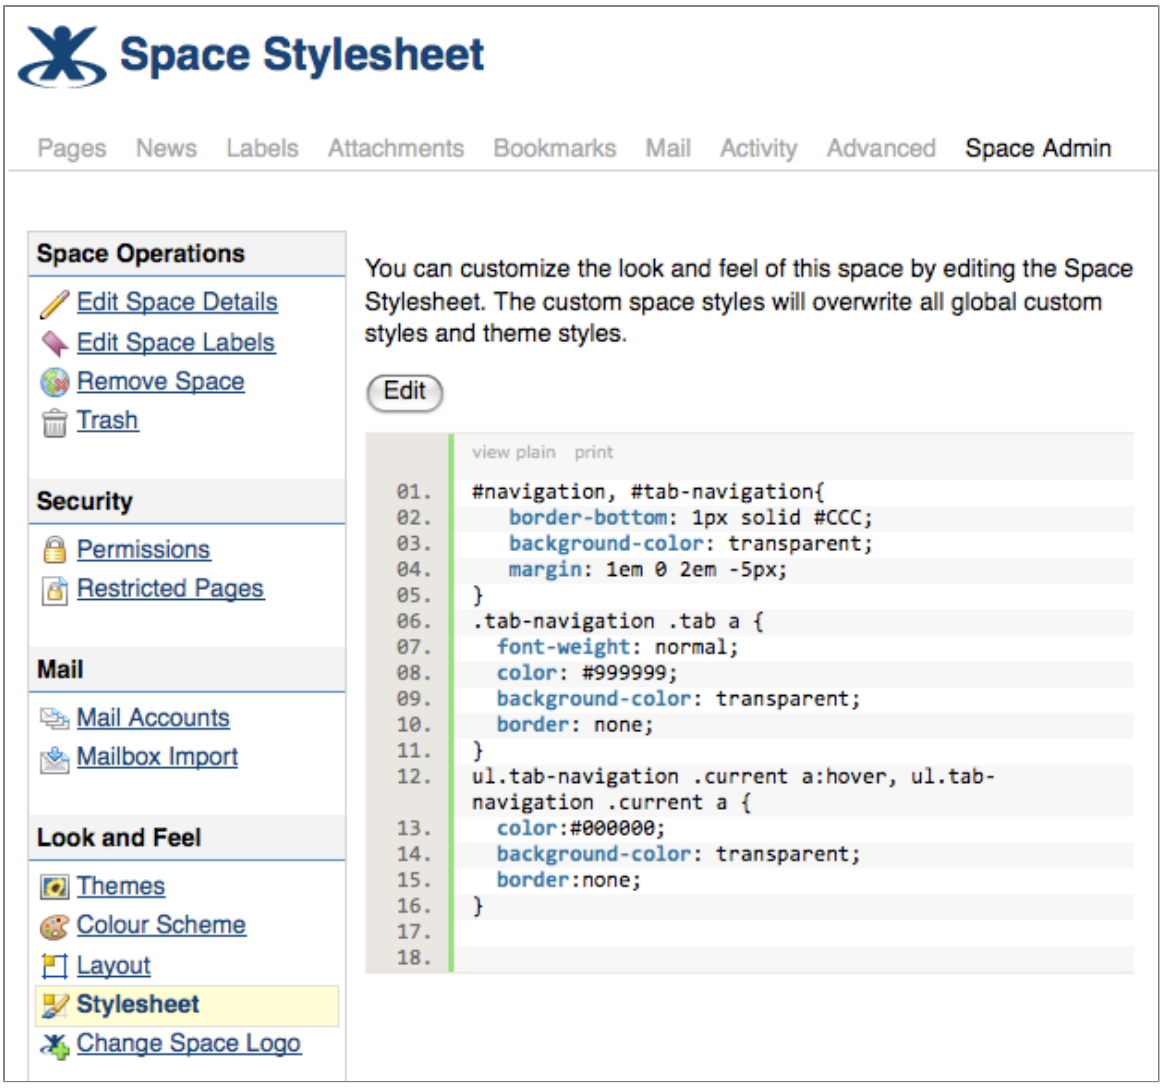




Custom Stylesheets for Confluence Spaces

Change the look of your Confluence space by specifying your own CSS styles. Cascading Style Sheets (CSS) are the standard way of styling web pages.

- For example, you might choose to change the background for the header at the top of each Confluence page. See our tutorial.
- Or you might change the look of the tabs in your Space Admin screens, as described in this example.
- To get started, take a look at our documentation.

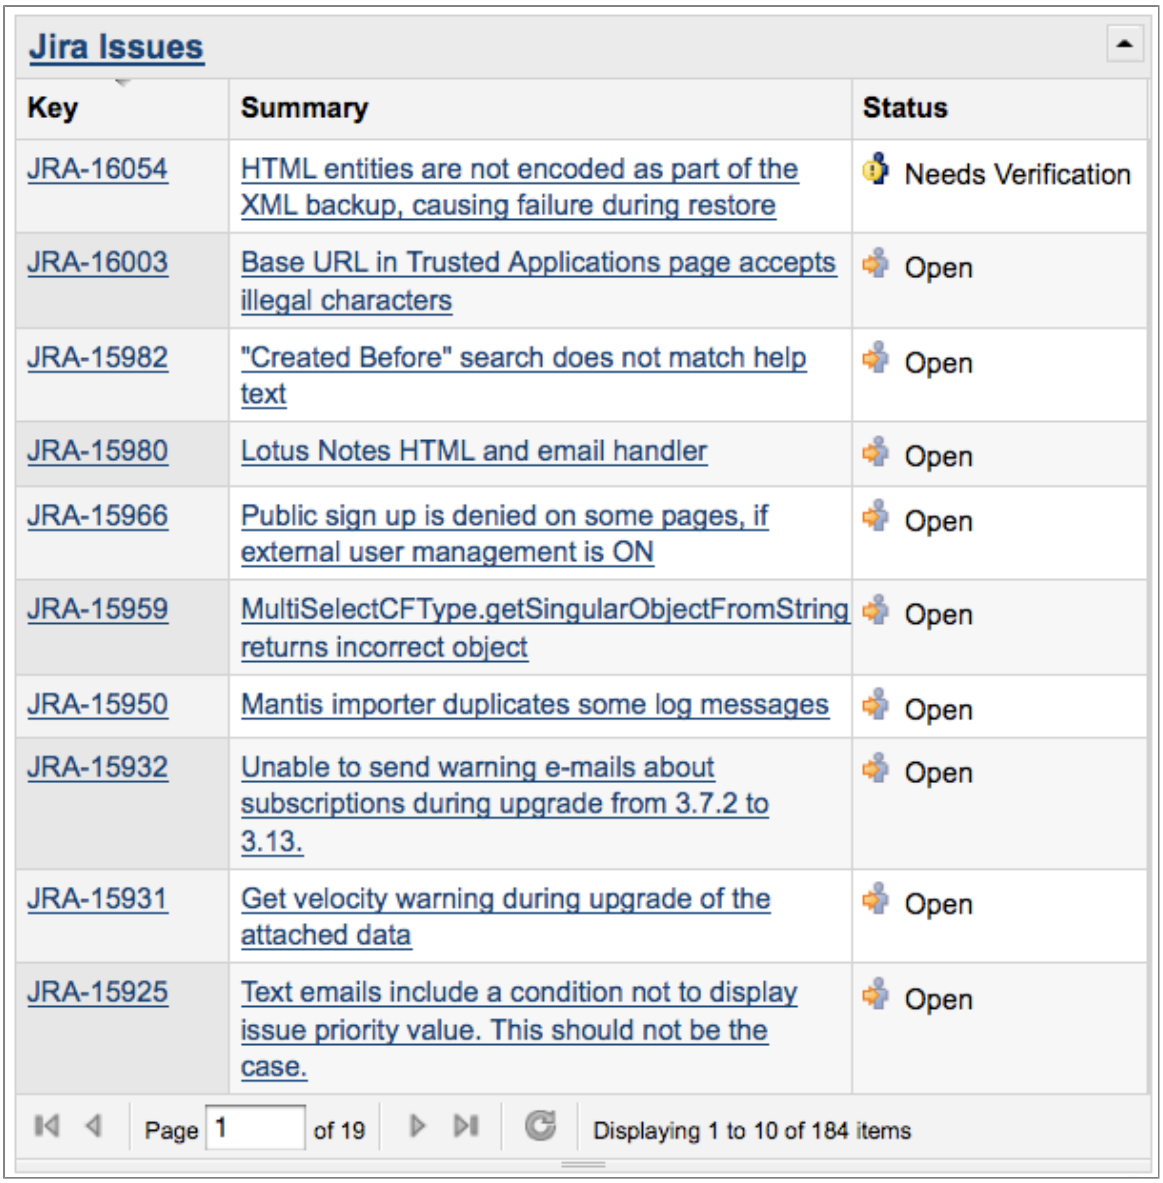




Updated JIRA Issues Macro with Custom Fields and Dynamic Display

The JIRA Issues macro now gives you more control both when viewing the output and when coding the macro.

- Specify any JIRA field as a column for display, including custom fields.
- Click the column headers to sort the output.
- Drag and drop the columns into a different order.
- Temporarily remove a column from the display.
- Click the triangle at top right of the issue table to collapse the table.
- Retrieve a page of issues at a time, rather than a huge list all at once.
- Take advantage of the improved performance of your JIRA site. The JIRA Issues macro now caches the results for anonymous users and for results retrieved using trusted communication. This should improve the performance of your JIRA site if you have set it up for trusted communication, because Confluence no longer needs to access the JIRA site as often.
- Our documentation tells you how to use the macro.



Jira Issues

Key	Summary	Status
JRA-16054	HTML entities are not encoded as part of the XML backup, causing failure during restore	Needs Verification
JRA-16003	Base URL in Trusted Applications page accepts illegal characters	Open
JRA-15982	"Created Before" search does not match help text	Open
JRA-15980	Lotus Notes HTML and email handler	Open
JRA-15966	Public sign up is denied on some pages, if external user management is ON	Open
JRA-15959	MultiSelectCFType.getSingularObjectFromString returns incorrect object	Open
JRA-15950	Mantis importer duplicates some log messages	Open
JRA-15932	Unable to send warning e-mails about subscriptions during upgrade from 3.7.2 to 3.13.	Open
JRA-15931	Get velocity warning during upgrade of the attached data	Open
JRA-15925	Text emails include a condition not to display issue priority value. This should not be the case.	Open

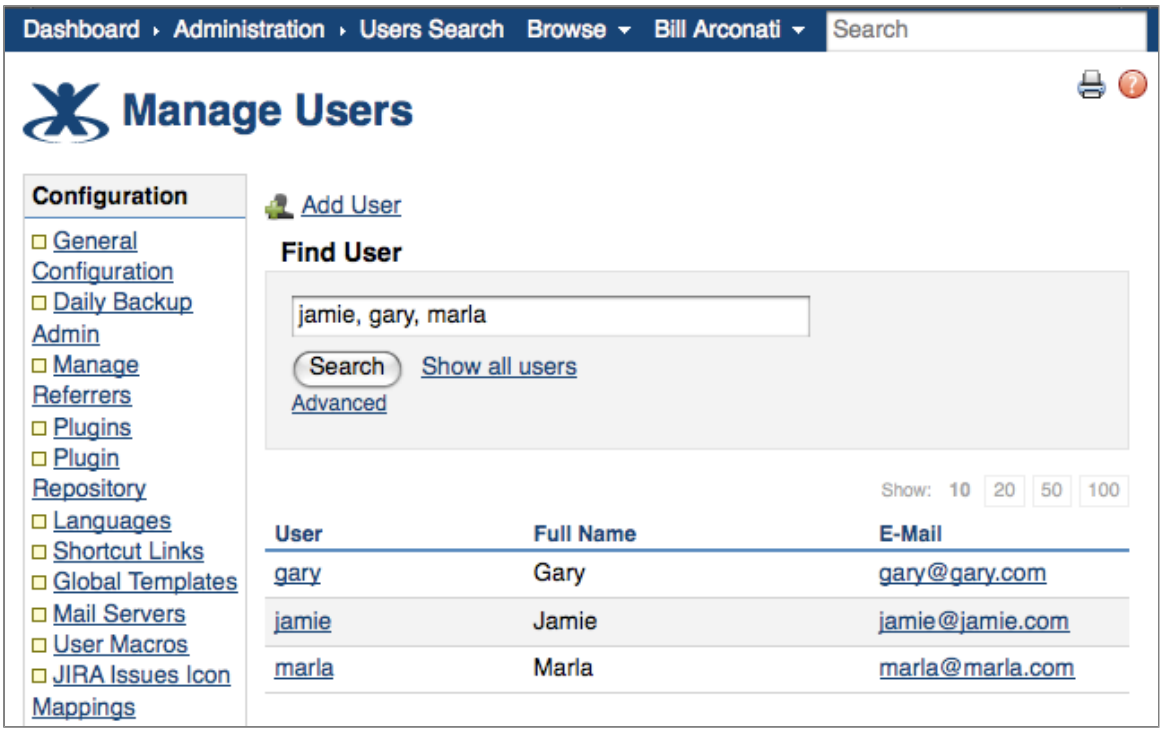
Page 1 of 19 Displaying 1 to 10 of 184 items



Enhanced User and Group Management

Searching for users and adding users to groups is now much easier.

- You can add and remove users directly from the group management screen. This allows you to manage the group membership for a number of users at the same time.
- The new user search offers a simple and an advanced option. With the simple option, just type all or part of the person's name, username or email address. If you want to restrict your search, use the advanced option to specify the field you want searched. Or you can search for users in a specific group.
- A new user picker lets you select the people you need from the list of matching users. This makes things much easier when adding members to a group, or when assigning page permissions and space permissions.





Upgraded Rich Text Editor

Confluence 2.10 brings many improvements to the Rich Text Editor and is a big step along the way to a great WYSIWYG experience.

- You can now use Safari to create and edit Confluence pages.
- The styles dropdown list now illustrates the format of each style, such as 'Heading 1', 'Heading 2' and so on.
- Tables are easier to insert and edit.

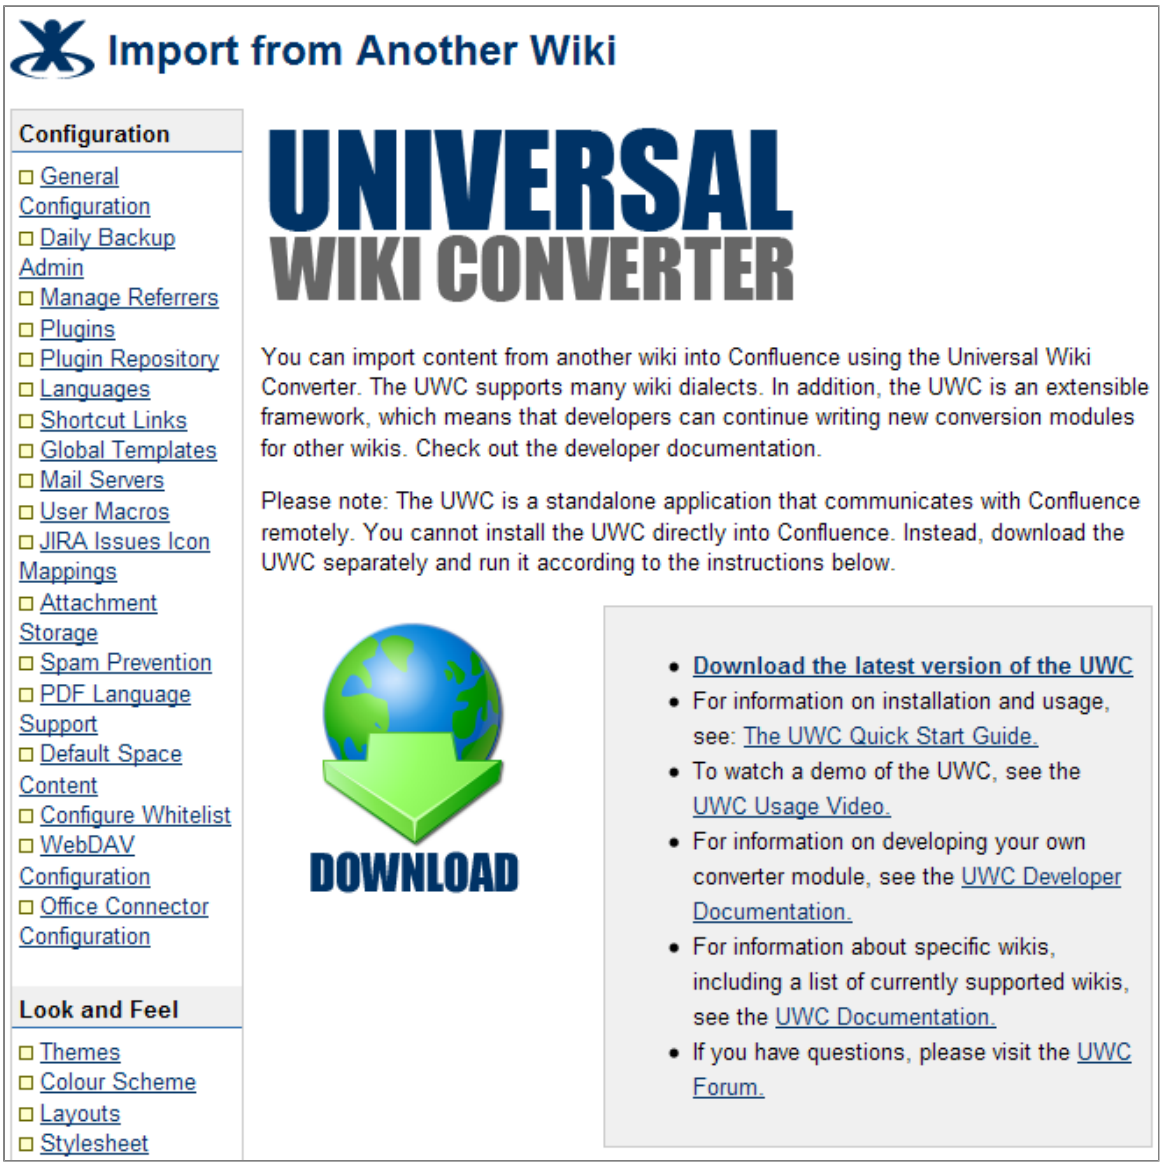




Universal Wiki Converter now with SharePoint Import and More

The Universal Wiki Converter (UWC) allows you to import content from other wikis into Confluence. Coinciding with the release of Confluence 2.10, there are some great improvements to this useful tool.

- You can now import pages from SharePoint wiki libraries, as well as other wikis, into Confluence.
- The UWC's enhanced user interface allows you to drag and drop wiki pages onto the UWC screen for conversion to Confluence.
- A new link on the Confluence Administration Console gives easy access the Universal Wiki Converter documentation and download pages.
- There is a new command-line interface to the UWC.
- The UWC also offers a new converter for Vqwiki, improvements to the Swiki converter, and more.





Improved Activity Macros

The Blog Posts macro, Recently Updated macro and Content by Label macro now support a common set of parameters, making it easier to code the macros and display the content you need.

- Filter content by author, label, space or content type.
- Use a minus sign (-) to exclude specific values. For example, using the 'author' parameter you can specify `author=-hpotter,hgranger,adumbledore`. You will get content which has been created/updated by either 'hgranger' or

'adumbledore' (or both) but 'hpotter' has not touched.
- Sort the resulting list of items by title, date created or date modified, in ascending or descending order.

For the HTML Include macro and the RSS macro, you can now specify a 'whitelist' of allowed URLs. This will improve the security of your Confluence site, because it can be dangerous to include content from untrusted external sites. Our documentation shows you how to specify the whitelist.



Plugin Framework 2

Confluence 2.10 comes with Atlassian's new Plugin Framework 2.1, based on Spring Dynamic Modules using an embedded OSGi container. The new framework lays the groundwork for the following improvements:

- More robust Spring component plugins.
- The ability for plugins to depend on each other.
- Control over plugin load order.
- The ability for plugins to define their own extension points.
- More consistent plugin APIs between products.
- More consistent plugin behaviour across different versions of Confluence.

Take a look at our developer documentation. The new plugin framework is under development. Here are some guidelines on converting your existing plugins to the new framework. We'd be delighted to have your feedback via our JIRA project.



More than 250 Fixes and Improvements

- A new attachments icon  on the first line under the title of a page tells you that there are files attached to the page, as well as how many attachments.
- A new lock icon  marks pages which have view or edit restrictions.
- The Demonstration Space included in the Confluence download now has more sample content. We have adapted pages from our own development, human resources and sales teams, to give some ideas on how your organisation might use Confluence.
- The default home page for a space now includes a list of recently updated content, a search input box and a tree view of the pages in the space. The default home page is created when you add a space. You can edit the home page to include or remove any content as required.
- You can now crop and tailor your profile picture and delete any profile pictures that you no longer want.
- There is no longer any need to re-create the database indexes manually during the upgrade procedure. From Confluence 2.10, the upgrade process will automatically re-create the indexes.
- Trusted authentication and other Seraph-based authentication methods are now available for calls to the Confluence RPC methods. This makes it practical to write front-end AJAX functionality which uses the remote API to retrieve or modify Confluence data.
- Take a look at the complete list of issues resolved in Confluence 2.10.

The Confluence 2.10 Team

Development

Bugfixing and Maintenance
Andrew Lynch
Brian Nguyen
Chris Kiehl
Matthew Jensen

Engine Room
Anatoli Kazatchkov
Charles Miller
Christopher Owen
Matt Ryall

Plugins
Ben Speakmon
Cheryl Jerozal
Jonathan Nolen
Nathan Dwyer

Rich Text Editor and Office Connector
Agnes Ro
David Taylor
Don Willis
Ryan Ackley

Search: Quick Navigation, Did-You-Mean and UI
Chris Broadfoot
David Loeng
Dmitry Baranovskiy
Paul Curren

Team Lead
Per Fragemann

Widget Connector, Custom Stylesheets and User/Group Management
Jens Schumacher

Support

Kuala Lumpur
Arie Murdianto
Azwandi Mohd Aris
David Chui
Fennie Ng
Mei Yan Chan
Ming Giet Chong
Tony Cheah Tong Nyee

San Francisco
Jeremy Largman
Maleko Taylor
Tim Wong
Vincent Chang

Sydney
Gurleen Anand
Ivan Benko
James Fleming
Michael Seager
Renan Battaglin
Roy Hartono

Others

Design
Jason Taylor
Stephen Russell

Performance Engineering
George Barnett

Product Management
Adnan Chowdhury

Product Marketing Management
Bill Arconati

Quality Assurance
Mark Hrynczak
Peter de Zwart

Technical Writing
Edwin Dawson
Sarah Maddox

Universal Wiki Converter
Laura Kolker

Confluence 2.10 Upgrade Notes

Below are some essential notes on upgrading to **Confluence 2.10**. For details of the new features and improvements in this release, please read the Confluence 2.10 Release Notes.

On this page:

- Upgrade Notes
  - Enabling the New Quick Navigation Feature on Customised Confluence Sites
  - Crowd and the User Search
  - RSS and HTML macro whitelists
  - Plugins
  - Java Versions
  - Platforms No Longer Supported
  - JavaScript Libraries
  - Confluence Themes

    - End of Life of SnipSnap Import
    - No Need to Re-Create Database Indexes Manually
- Upgrade Procedure

Upgrade Notes

Enabling the New Quick Navigation Feature on Customised Confluence Sites

If you have customised your Main Layout on either the space or the global level, or if you have a custom theme plugin, the new quick navigation feature will not work immediately for you.

To enable quick navigation, you need to add the following statement to your Main Layout anywhere before the `#standardHeader()` statement:

```
1.#requireResourcesForContext("main")
```

Crowd and the User Search

Confluence 2.10 includes a much-enhanced user search. (See the release notes.) If you are using Atlassian's Crowd for user management, you will need Crowd 1.5.1 or later to use the 'Simple' option in the user search. If your version of Crowd does not support the simple user search, you will see only the 'Advanced' search form. Please consider upgrading your Crowd version, to take advantage of the advanced user search.

RSS and HTML macro whitelists

To improve the security of a default installation of Confluence, macros which display HTML from external sources now use a whitelist of URLs configured by the Confluence administrator. This affects the RSS and HTML include macros.

See Configuring a URL Whitelist for instructions on how to configure a list of allowed URL patterns for the RSS and HTML macros in Confluence 2.10.

Plugins

If you are using any third-party plugins on your Confluence instance, please test them thoroughly before rolling 2.10 into production.

Java Versions

Java 1.4 is not supported in Confluence 2.9 and later. Please refer to the Java 1.4 Support Timeline for more information.

Before upgrading to Confluence 2.10, you will need to ensure your environment is running at least Java 5. Confluence 2.10 supports Java 5 and Java 6. We recommend Java 6 because of its increased performance and easier troubleshooting, due to enhanced memory dump and profiling capabilities.

You can check your current Java version in Confluence:

1. Go to the Confluence '**Administration Console**'. To do this:
    - Open the '**Browse**' menu and select '**Confluence Admin**'. The 'Administration Console' view will open.
2. Select '**System Information**' from the '**Administration**' section in the left-hand panel.
3. Refer to '**Java Version**'.
    - If the version is 1.5 or higher, you do not need to do anything.
    - If the version is 1.4, you need to upgrade your JDK before you can upgrade to Confluence 2.10.

If you are running the Confluence EAR-WAR edition against your own application server, you will need to check with your application server vendor about which JDK versions are supported.

Platforms No Longer Supported

The following platforms were on the 'unsupported but working' list for Confluence 2.9. They are not on that list as of Confluence 2.10. We will still test those platforms infrequently, and Confluence will probably still work fine with them for a while, but they will not be officially supported.

- MySQL 4.1 — please upgrade to MySQL 5
- Tomcat 5.0 — please upgrade to Tomcat 5.5 or 6
- Resin 2 — please upgrade to Resin 3
- JBoss 4.0.x — please upgrade to JBoss 4.2.x

Please refer to the list of supported application servers and databases on our Supported Platforms topic for the updated matrix.

There is a workaround to enable MySQL 4.1.x.

JavaScript Libraries

jQuery is the supported JavaScript library for plugin developers.

Advance notice — deprecated libraries: We have decided to standardise on jQuery as the JavaScript library for Confluence. This library will eventually replace all others. For this reason, use of the following JavaScript libraries in Confluence is **deprecated**:

- Prototype
- Scriptaculous

Because there is a lot of legacy plugin code using Prototype and Scriptaculous, these will continue to be available for this release of Confluence.

The **Yahoo! User Interface Library** (YUI) is no longer supported, following its removal in Confluence 2.9.

Confluence Themes

The Confluence Classic Theme is no longer supported, following its deprecation in Confluence 2.9. This theme uses outdated typography and formatting, which was replaced by the new-look Default Theme in Confluence 2.6. At that time, we introduced the Classic Theme to minimise the impact on customers who relied on the older typography.

When you upgrade to Confluence 2.10, an upgrade task runs to do the following:

- Any space with the Classic Theme configured will be migrated to use the Default Theme.
- If your Global Theme is set to use the Classic Theme, the Global Theme is also migrated to the Default Theme.
- The Classic Theme plugin is uninstalled. (The Classic Theme entry in the database is removed.)

End of Life of SnipSnap Import

Advance notice — EOL SnipSnap import: Confluence 2.10 is the last release which will include the SnipSnap Import option in the Confluence Administration Console. From the next release after Confluence 2.10, the SnipSnap import will no longer be available.

No Need to Re-Create Database Indexes Manually

In previous releases of Confluence, you needed to manually re-create the database indexes during the upgrade procedure. For this purpose, we provided a set of SQL statements as an attachment to the Upgrade Guide. From Confluence 2.10, the upgrade process will automatically re-create the indexes. Please refer to the Upgrade Guide.

Upgrade Procedure



Upgrade a test environment first
As always, please test your upgrades in your test environment before rolling into production.

If you are already running a version of Confluence, please follow these instructions to upgrade to the latest version:

1. Before you upgrade, we strongly recommend that you back up your Confluence Home Directory and database. See the documentation on backing up your Confluence site. If you are using an external database, perform a database backup.
2. If your version of Confluence is earlier than 2.9.x, read the release notes and upgrade guides for all releases between your version and the latest version. In particular:
   - Please read the 2.9 upgrade notes.
   - If you are upgrading from 2.1 or earlier, please read the 2.2 release notes.
   - If you are upgrading from 2.2 or earlier, you will need to upgrade to Confluence 2.5.8 first, confirm the upgrade was successful, then upgrade again from version 2.5.8 to the latest. For more details, please refer to CONF-11767.
3. Download the latest version of Confluence.
4. Follow the instructions in the Upgrade Guide.
5. If you encounter a problem during the upgrade, please create a support ticket and one of our support engineers will assist you through the process.

RELATED TOPICS

Confluence 2.10 Release Notes

Issues resolved in Confluence 2.10

Below are the issues resolved in Confluence 2.10, ordered by number of votes. For the full details of the fixes, improvements and new features, please take a look at our issue tracker. You can also take a look at the Confluence 2.10 Release Notes.

JIRA Issues (15 issues)

Type	Key	Summary	Status	Resolution	Votes
	CONF-2407	Add custom fields to the list of columns to display in jiraissues macro	 Resolved	Fixed	78
	CONF-2191	Group management improvements	 Resolved	Fixed	76
	CONF-3864	Support Safari For Wysiwyg	 Resolved	Fixed	71
	CONF-9472	Can't open hyperlinks from Word when a session is already opened	 Resolved	Fixed	37

	CONF-12742	JIRA Portlet Macro displays portlet twice	Resolved	Fixed	32
	CONF-3135	Support all built-in JIRA fields in the {jiraissues} macro	Resolved	Fixed	31
	CONF-8798	Rich text editor adds new line before and after inline user macros	Resolved	Fixed	29
	CONF-5885	Allow removal of uploaded profile pictures	Resolved	Fixed	23
	CONF-10454	Default Space Content does not support macro	Resolved	Fixed	18
	CONF-10177	Cannot export page as PDF when page contains attachments with UTF8 Characters.	Resolved	Fixed	18
	CONF-7223	Color macro unstable near links in rich text editor	Resolved	Fixed	18
	CONF-10535	Alphabetical page list (under Browse Space) is too slow on large spaces	Resolved	Fixed	17
	CONF-5025	Firefox/IE retrieve cached version of attachments	Resolved	Fixed	17
	CONF-7266	jira issues macro column support for the time tracking fields	Resolved	Fixed	14
	CONF-10088	Confluence libraries incompatible with Websphere 6.0 and Weblogic 9: rss feeds or pdf generation	Resolved	Fixed	12

Above are the issues resolved in Confluence 2.10, ordered by number of votes. For the full list of fixes, improvements and new features, please take a look at our issue tracker.

Workaround For Enabling MySQL 4.1.x with Confluence 2.10

With the release of Confluence 2.10, MySQL 4.1.x is no longer supported and will not work by default. See the Supported Platforms topic for further details.

However, there is a workaround to enable MySQL 4.1.x with Confluence 2.10.

To enable MySQL 4.1.x with Confluence 2.10:

1. When launching Confluence from the command line, add the following parameter:

```
1.-Dmysql4Compatibility=true
```

2. This will enable MySQL 4.1.x to work with Confluence 2.10. Please note however, that use of this database will not be supported by Atlassian.

Confluence 2.9.3 Release Notes

Confluence 2.9.3 is the **standard edition** version of Confluence 2.9.x. This version of Confluence is equivalent to Confluence version 2.9.2, but it does not include Oracle's Coherence technology. Due to licensing changes with Oracle over the Coherence clustering and distributed caching technology, Atlassian is no longer able to distribute this technology to customers without a Confluence clustered license.

Therefore, Atlassian has released 'standard editions' for each previous major version of Confluence back to 2.6, which includes Confluence versions 2.10.4, 2.9.3, 2.8.3, 2.7.4 and 2.6.3. Confluence 2.9.3 will be the only Confluence 2.9.x version available to customers with a non-clustered Confluence license.

For its caching functionality, Confluence 2.9.3 utilises Ehcache. This functionality was provided by the Coherence technology in previous Confluence 2.9.x versions.

For more information about the features, updates and fixes in Confluence 2.9.2, please refer to the Confluence 2.9.2 Release Notes.

Upgrading from a Previous Version of Confluence

Upgrading Confluence should be fairly straightforward. Please read the Confluence 2.9.2 Upgrade Notes. However, if you have customised the cache settings in your installation of Confluence (e.g. for performance reasons), then please read the Confluence 3.0.1 Upgrade Notes for important information about transferring your caching layer customisations from Coherence to Ehcache.

We *strongly recommend* that you back up your `confluence.home` directory and database before upgrading.

Confluence 2.9.2 Release Notes

14 October 2008

Confluence 2.9.2 is a recommended upgrade which fixes some security flaws and other issues.

Please refer to the security advisory for details of the security vulnerabilities, risk assessment and mitigation strategies.

You can now view the Wiki Markup code for previous versions of a page as well as the current version. Open a previous version from the page history, then select 'View Wiki Markup' from the 'Tools' menu. Previously, a bug caused Confluence to show the Wiki Markup only for the current version of the page.

Another bug caused an error to occur when you copied a page and tried to add an attachment before saving the page. This is now fixed, along with a number of other issues.

Don't have Confluence 2.9 yet?

Take a look at the new features and other highlights in the Confluence 2.9 Release Notes.

Download Latest Version

Upgrading from a Previous Version of Confluence

Upgrading Confluence should be fairly straightforward. Please read the Confluence 2.9.2 Upgrade Notes. We *strongly recommend* that you back up your `confluence.home` directory and database before upgrading.

Updates and Fixes in this Release

JIRA Issues (29 issues)

Type	Key	Summary	Priority	Status	Resolution
	CONF-13043	XSS in pagetree plugin		Resolved	Fixed
	CONF-13041	XSS in bookmarks plugin		Resolved	Fixed
	CONF-13092	Provide Patch for XWork ParametersInterceptor attacks		Resolved	Fixed
	CONF-13039	Privilege escalation: User is able to add a page to his watchlist without having the permission		Resolved	Fixed
	CONF-13042	XSS in RSS feed creation		Resolved	Fixed
	CONF-12940	Permission problem in preview blog post		Resolved	Fixed
	CONF-13040	Stored XSS in wiki macro search		Resolved	Fixed
	CONF-12944	XSS in site search action		Resolved	Fixed
	CONF-11640	Errors retrieving mail from Microsoft Exchange: "Unable to load BODYSTRUCTURE"		Resolved	Fixed
	CONF-12920	Paging on Undefined Pages does not work with Auto Html Encoding		Resolved	Fixed
	CONF-12827	Paging does not work on People Directory with Auto Html Encoding		Resolved	Fixed
	CONF-12825	Exclamation marks in comments encoded when collapsed		Resolved	Fixed
	CONF-12921	Paging on Orphaned Pages does not work with Auto Html Encoding		Resolved	Fixed
	CONF-12922	Paging on Restricted Pages does not work with Auto Html Encoding		Resolved	Fixed
	CONF-12291	Avalon Logkit library is missing.		Resolved	Fixed

CONF-11022	Saving a copied page with attachments added to draft fails with NullPointerException	↑	Resolved	Fixed
CONF-13051	View Source only shows the latest version	↑	Resolved	Fixed
CONF-13035	View Page History link broken when viewing the latest version	↑	Resolved	Fixed
CONF-12727	When viewing diffs between pages, there is no longer the option to page between them (link missing)	↑	Resolved	Fixed
CONF-11854	c3p0-0.9.1.2: Built-In JMX-Service incompatible with Websphere 6.1.0.9	↑	Resolved	Fixed
CONF-11324	Session isn't invalidated on logout	↑	Resolved	Fixed
CONF-12799	SearchBean pagination does not check the sanity of the start index	↑	Resolved	Fixed
CONF-12889	Creating a new user macro with the same name name as an existing one silently clobbers the old macro	↑	Resolved	Fixed
CONF-12826	Paging on Mail does not work with auto Html Encoding	↑	Resolved	Fixed
CONF-13003	Deleting a label takes you to the top of a page	↓	Resolved	Fixed
CONF-9563	Classpath Duplication Check fails on Resin 3	↓	Resolved	Fixed
CONF-10580	Confluence uses a version of HSQLDB that sometimes returns results in wrong order	↓	Resolved	Fixed
CONF-10364	Page titles with two hyphens in it "--" cause html "end comment tags" ("-->") between page comments	↓	Resolved	Fixed
CONF-12816	ConvertedPath class uses a VelocityEngine that doesn't know about Atlassians velocity.properties file	↓	Resolved	Fixed

Confluence 2.9.2 Upgrade Notes

Confluence 2.9.2 is a recommended upgrade which fixes some security flaws as well as other bugs. You'll find details of the fixes in the release notes.

Upgrade Procedure

If you are already running a version of Confluence, please follow these instructions to upgrade to the latest version:

1. Before you upgrade, we strongly recommend that you **back up your `confluence.home` directory and database**.
2. If your version of Confluence is earlier than 2.9.0, read the release notes and upgrade guides for **all releases** between your version and the latest version. In particular:
   - Please read the 2.9 upgrade notes and the 2.9.1 upgrade notes.
   - If you are upgrading from 2.1 or earlier, please also read the 2.2 release notes.
3. Download the latest version of Confluence.
4. Follow the instructions in the Upgrade Guide.

Confluence 2.9.1 Release Notes

8 September 2008

Confluence 2.9.1 is a recommended upgrade which fixes some security flaws and other issues.

Please refer to the security advisory for details of the security vulnerabilities, risk assessment and mitigation strategies.

This release also addresses an issue relating to contributor filtering with LDAP. In the previous version, the search function was not able to filter results by authors from an LDAP directory. This fix requires an upgrade action; see the Confluence 2.9.1 Upgrade Notes for more information.

There's a complete list of fixes below. Click a specific issue to see details of the fix.

Don't have Confluence 2.9 yet?

Take a look at the new features and other highlights in the Confluence 2.9 Release Notes.

Download Latest Version

Upgrading from a Previous Version of Confluence

Upgrading Confluence should be fairly straightforward. Please read the Confluence 2.9.1 Upgrade Notes. We *strongly recommend* that you back up your `confluence.home` directory and database before upgrading.

Updates and Fixes in this Release

JIRA Issues (28 issues)					
Type	**Key**	**Summary**	**Priority**	**Status**	**Resolution**
	CONF-12911	Pages that inherit page restrictions are not respecting those restrictions after upgrade to Confluence 2.9		Resolved	Fixed
	CONF-12856	Plugins can lose access to resources in inner jars		Resolved	Fixed
	CONF-7615	XSS bug: usernames not HTML-encoded in all places		Resolved	Fixed
	CONF-12860	Hidden pages' content can be viewed without permission using diffpages.action		Resolved	Fixed
	CONF-12785	Confluence 2.9 Installer does not work when installed as service		Resolved	Fixed
	CONF-12859	Hidden pages' content can be viewed without permission using copypage.action		Resolved	Fixed
	CONF-12714	Corrupt plugin jar can cause entire system to fail to start with java.lang.IllegalStateException: error in opening zip file		Resolved	Fixed
	CONF-12770	Update French and German Translations		Resolved	Fixed
	CONF-10636	Labels Lost during Import		Resolved	Fixed
	CONF-12748	2.8 default theme does not render edit, tools, and Add menus - cut off halfway		Resolved	Fixed
	CONF-12845	View Wiki Markup availabe on Page Restricted pages		Resolved	Fixed
	CONF-11949	Some words are not translated in French		Resolved	Fixed
	CONF-13275	Ognl exception and Homepage set to blank when editing space details		Resolved	Fixed
	CONF-10769	LDAP users are added into People Directory only when they click on their Preferences		Resolved	Fixed
	CONF-12710	'Recent Changes' in page info contains same author multiple times		Resolved	Fixed
	CONF-12717	TOC plugin cannot work with umlaute character		Resolved	Fixed
	CONF-12666	Bookmarks labels can't be added/edited		Resolved	Fixed
	CONF-12766	Links to specific comments on a page do not jump to the comment location after the page is loaded		Resolved	Fixed
	CONF-12711	Searches throw IllegalArgumentException when indexing language is French		Resolved	Fixed
	CONF-11815	Select template link not translated in French language pack.		Resolved	Fixed
	CONF-12811	The TrustedApplicationsFilter doesn't work for /rpc/* URLs			Fixed

				Resolved	
	CONF-12813	Recently-updated macro can hang Dashboard if images are Corrupt		Resolved	Fixed
	CONF-12787	Rich Text full screen editor does not display any content or tools in Confluence 2.9		Resolved	Fixed
	CONF-12708	Review and release contributor filter repair note for Confluence 2.9.1		Resolved	Fixed
	CONF-10415	When viewing older version of a page, the breadcrumbs have 'undefined space' in them		Resolved	Fixed
	CONF-12706	'Edit My Profile' breadcrumb menu produces 'com.atlassian.confluence.extra.userinfo.EditUserInfoAction.action.name' in menu instead of friendly name		Resolved	Fixed
	CONF-12790	The "no results found" message for the Search should include the author filter in the suggestions		Resolved	Fixed
	CONF-10558	IE7 insert and delete row problems in the rich text editor		Resolved	Fixed

Confluence 2.9.1 Upgrade Notes

Confluence 2.9.1 is a recommended upgrade which fixes some security flaws as well as other bugs. You'll find details of the fixes in the release notes.

Upgrade Procedure

If you are already running a version of Confluence, please follow these instructions to upgrade to the latest version:

1. Before you upgrade, we strongly recommend that you **back up your `confluence.home` directory and database**.
2. If your version of Confluence is earlier than 2.8.0, read the release notes and upgrade guides for **all releases** between your version and the latest version. In particular:
   - Please read the 2.9 upgrade notes.
   - If you are upgrading from 2.1 or earlier, please also read the 2.2 release notes.
3. Download the latest version of Confluence.
4. Follow the instructions in the Upgrade Guide.

Fix for Contributor Filtering LDAP Users

An issue was fixed in this release relating to search functionality and filtering by contributor. Authors from LDAP directories are now taken into account, however you will need to carry out the instructions to make sure all users are included in this feature.

Enabling Contributor Filtering for Search

Confluence 2.9 includes an enhancement to the search functionality which allows the filtering of search results by author or contributor. Unfortunately an existing bug in Confluence prevented this functionality from being used on sites that use LDAP for their user management.

Monitoring the Upgrade Task in Confluence 2.9.1

Confluence 2.9.1 fixes this defect and also includes an upgrade task that should ensure any existing content contributors are properly indexed and therefore available in the contributor filter auto-complete box.

Enabling this functionality is simply a case of installing Confluence 2.9.1 and watching your log file during start up for specific messages, such as the following.

```
Beginning PersonalInformation repair.
Found <n> usernames that need to be retrieved.
Finished Personal Information repair.
```

On a large external directory of users, it is possible that this task may take many minutes to run. However, the processing of this task will not delay the startup of Confluence.

Enabling Detailed Progress Reporting during Upgrade

If you wish to see more detailed progress reporting during the upgrade, you will need to carry out the following steps.

To enable detailed progress reporting,

1. Enable '**Debug**' level logging for the logger called `com.atlassian.confluence.upgrade.PersonalInformationRepairTask`.
2. During startup of Confluence, you should also see additional progress reporting messages similar to the following:

```
Created 10 from <n> missing Personal Information objects.
Created 20 from <n> missing Personal Information objects.
Created 30 from <n> missing Personal Information objects.
```

Manually Instigating the Fix

Should you discover any problems during the repair task (as reported in the logs), then a Confluence Administrator can manually rerun the task by visiting the following page.

```
1.<confluence base URL>/admin/createMissingPersonalInfo.jsp
```

That page presents a single button that will re-run the repair task. In addition to the logging noted previously, the same messages will be returned to the administrator's browser.

Actions Taken when Running the Fix

There is no functional cost to running this fix. Should you accidentally run it too frequently, there is no functional implication. It is worth being aware of what is happening in the back-end so you can schedule a relevant time to perform it. While the fix is running, the following occurs:

- Two database queries are run to find the missing contributors, each involving an inner join against the same table.
- Each missing user is requested from the LDAP server by username.
- Each missing user will lead to a row insert in the database (these are batched).

Restoring Inherited Page Permissions After 2.9 Upgrade

Confluence instances that were upgraded to Confluence 2.9 are affected by a vulnerability. Child pages under a page protected by permissions are not protected by inherited permissions, as they should be. Please note that new installs of Confluence 2.9 which were not an upgrade from an old version are **not** affected.

Carry out the steps below to rectify the situation.

To Restore Inherited Page Permissions After Upgrading to Confluence 2.9,

1. Log into Confluence as '**Administrator**'.
2. Access this specific page in Confluence:

   ```
   1.CONFLUENCE_HOME/admin/permissions/pagepermsadmin.action
   ```

   (replace 'CONFLUENCE_HOME' with the domain name of your own Confluence instance).

3. On that page, a single button is visible, entitled '**Rebuild Ancestor Table**'. Click that button. It will report its success.
4. Now, go to the '**Confluence Admin**' page. Here, click '**Cache Statistics**' from the left navigation bar. A long list appears.
5. Find '**Inherited Content Permissions**' in the list. Now, click the '**Flush**' button to the right of '**Inherited Content Permissions**'.
6. Inherited permissions will now be applied.

Read more about this vulnerability in the Security Advisory.

Confluence 2.9 Release Notes

7 August 2008
With great pleasure, Atlassian presents **Confluence 2.9**.

First up is the Search. The new screen design focuses the eye on your search term and results. To help you find information more quickly, Confluence now searches all content types by default and puts the most relevant results at the top of the list. Because a wiki is all about people, Confluence treats personal information as the most relevant. With author filtering, you can now find content written by a specific person.

The macros bundled with Confluence have been treated to a major overhaul. The Chart and Gallery macros offer sophisticated new displays and more interactivity. The revised Pagetree macro is now included in the Confluence download. Try using the Pagetree to add a navigation panel to your pages. We have added some popular new features and fixed a number of much-voted-for bugs in other macros too.

When you are editing a page, Confluence now automatically saves your work and displays the time of the last auto-save. If something goes wrong, you can quickly retrieve your work by selecting 'Drafts' from the user menu.

Other features include resetting page order to alphabetical and some much-needed help in the ongoing battle against spam.

Highlights of this Release:

- Streamlined Search
- Auto Save
- Charts
- Page Tree
- Gallery
- New Tutorial
- More in the Menus
- Alphabetical Page Ordering
- Better Spam Prevention
- Plugin Repository
- Engine Room and Developers' Community
- More than 140 Fixes and Improvements
- Special Thanks

Responding to your Feedback:

210 votes satisfied



- Thank you for all your issues and votes. Keep logging, to help us keep improving!
- Below is a list of the highlights in this release.
- Attached is the full list of issues resolved in this release.

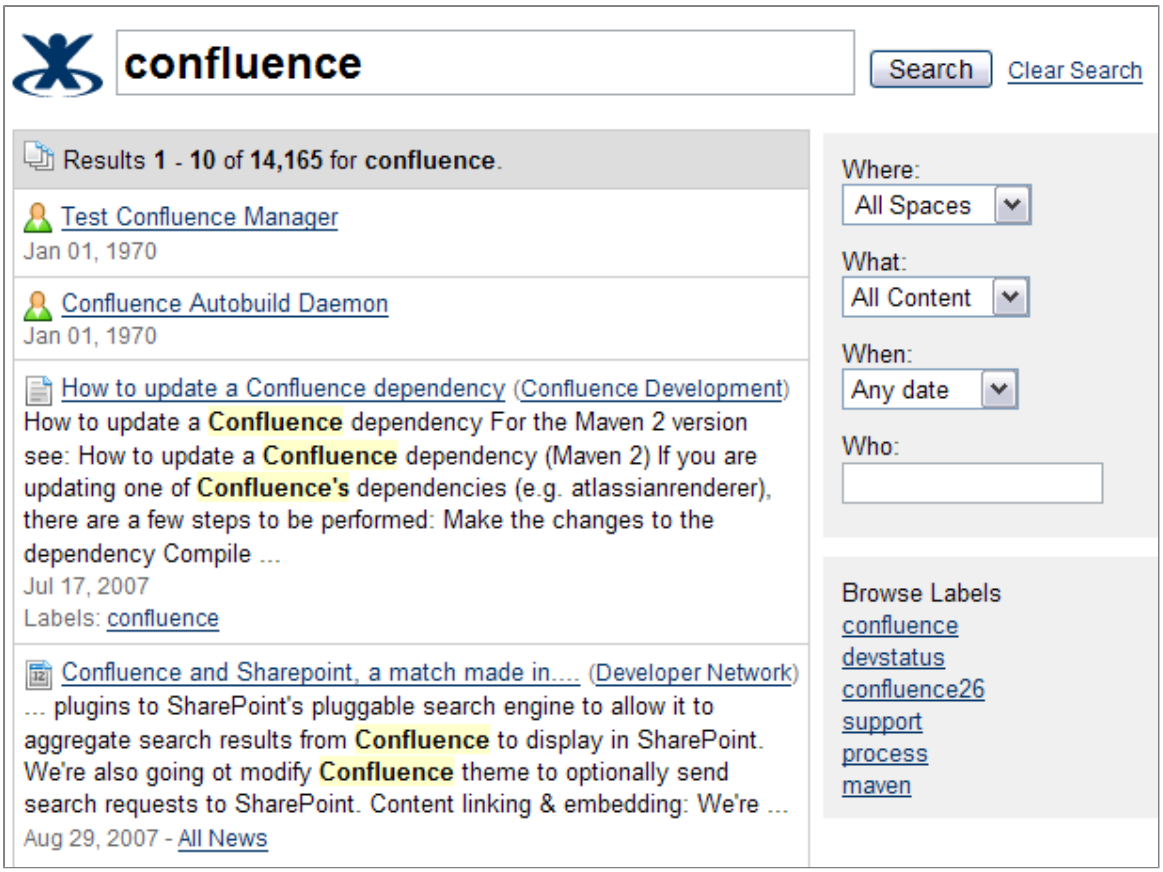
Upgrading from a previous version of Confluence
• **Confluence 2.9 requires Java 5 at a minimum**, and will no longer work with Java 1.4. Before upgrading to Confluence 2.9, please refer to the Confluence 2.9 Upgrade Notes for more details on this change. • Upgrading Confluence should be fairly straightforward. ***We strongly recommend that you back up your Confluence Home directory and your database before upgrading.*** • Please refer to the Confluence 2.9 Upgrade Notes for further essential information about **plugins** and other factors affecting your upgrade.

Highlights of Confluence 2.9



Streamlined Search

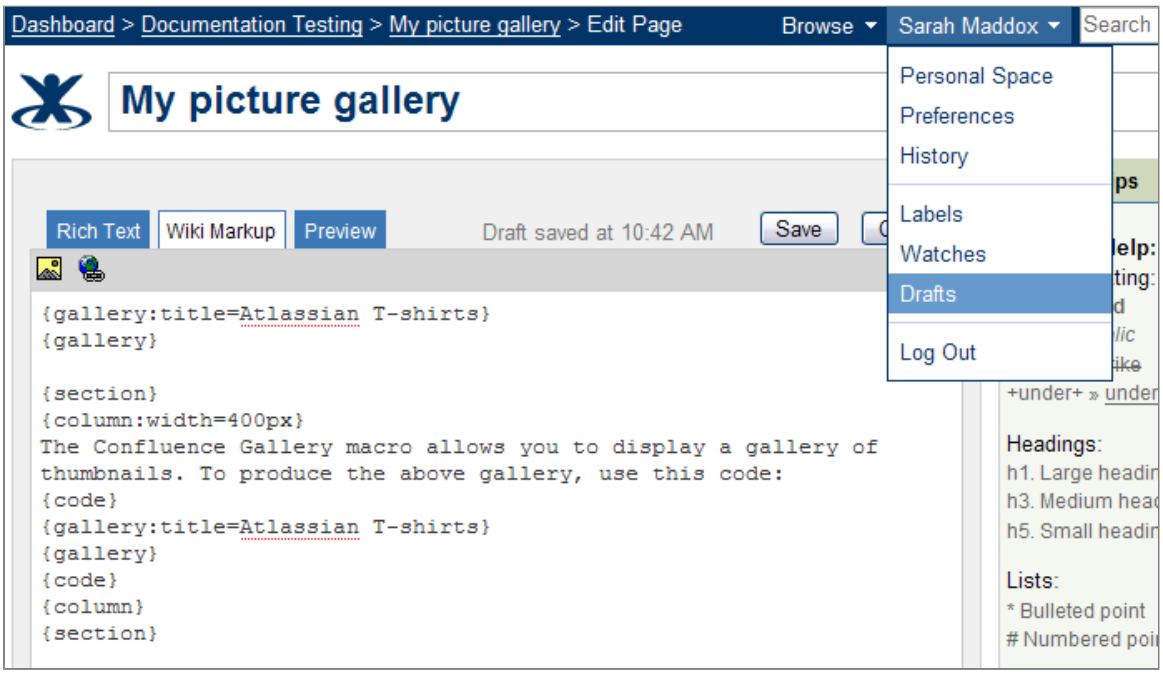
- The redesigned Search screen helps you to focus on the search term and results, by simplifying the filter box and other parts of the screen.
- By default, the search now includes all spaces and all content types: pages, news items, comments and so on.
- You can filter the search results by space, content type and date modified.
- If your Confluence site uses the standard out-of-the-box configuration for user management, you can also filter the search results by contributor.
- The search results are weighted to give personal information the highest relevance, followed by pages, news items and the other content types.
- When using the Search macro or the Search API, you can now also filter by contributor.
- Take a look at our documentation for more details on the search and the ranking of search results.





Auto Save

- Confluence automatically saves your work to a draft while you are adding or editing a page.
- Now with release 2.9, a message near the 'Save' button shows the time of the latest auto-save.
- If something goes wrong, retrieve the saved draft easily from the dropdown menu under your name.
- You can read more about drafts in the documentation.





Charts

- The Chart macro presents a sophisticated new look.
- Below are some example charts displayed on a Confluence page.
- The documentation shows you how to add the macro to your page and choose different formats and types.

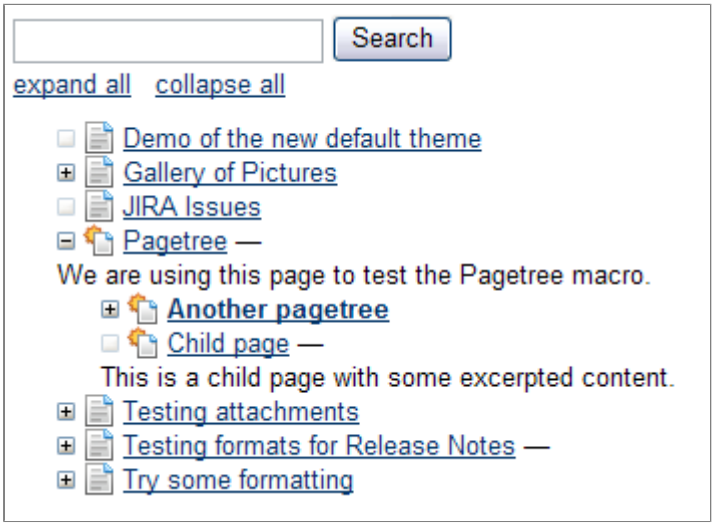




Page Tree

- The Pagetree macro is now bundled as part of your Confluence installation.
- It displays a dynamic, hierarchical list of pages which you can use as a table of contents or a navigation panel.
- You can choose to include a search box above the tree of pages or on its own.
- New links allow viewers to collapse or expand all branches in the tree at once.

- There are a number of options for specifying the root of the page tree. For example, you can show children of the current page, or children of the current page's parent, or all pages in the space.
- Take a look at the documentation to see the complete list of options.





Gallery

- Use the Gallery macro to display a set of pictures on a page.
- Now you can include or exclude individual pictures, or simply display all the images at once.
- Choose pictures from the current page or another Confluence page.
- Take a look at the documentation for help on these and other options.



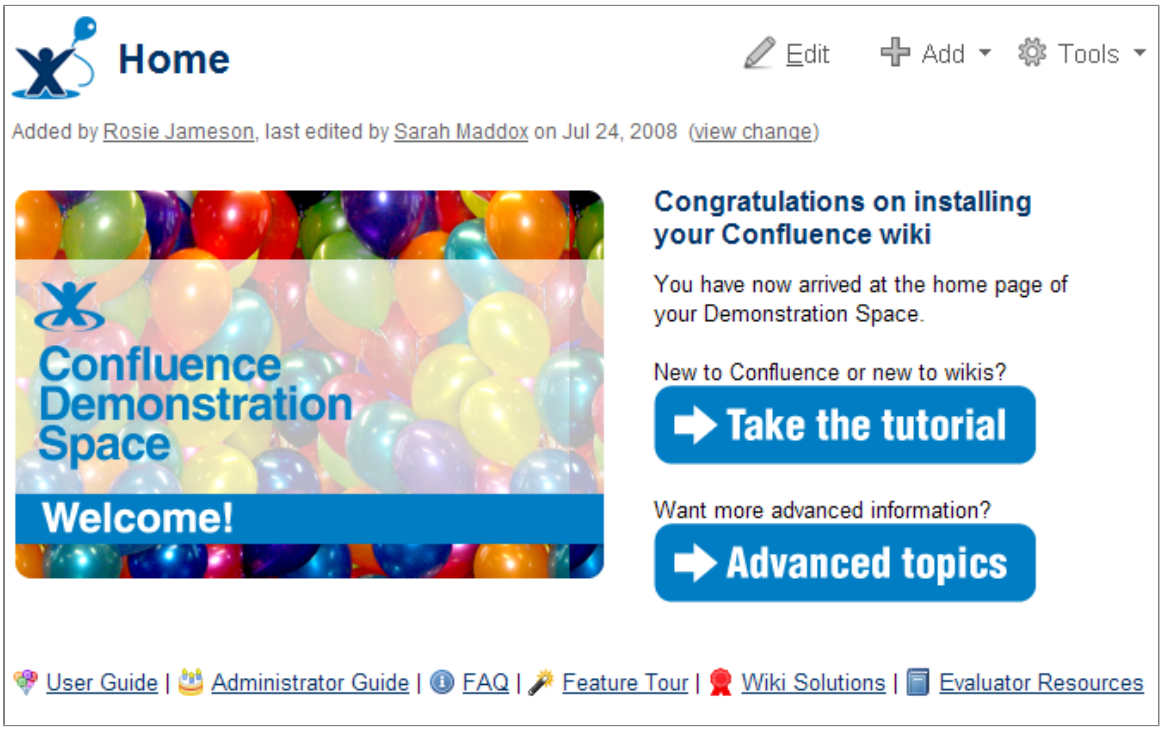
- Viewers can click an image to zoom in and view the gallery as a slide show.



6

New Tutorial

- When you download Confluence and choose to include the sample content, you will receive a new Demonstration Space.
- There's a tutorial for those are new to Confluence or new to wikis.
- Other users will enjoy the advanced topics and links to more information.
- You can use the tutorial as a quick-start guide for new starters in your organisation.



7

More in the Menus

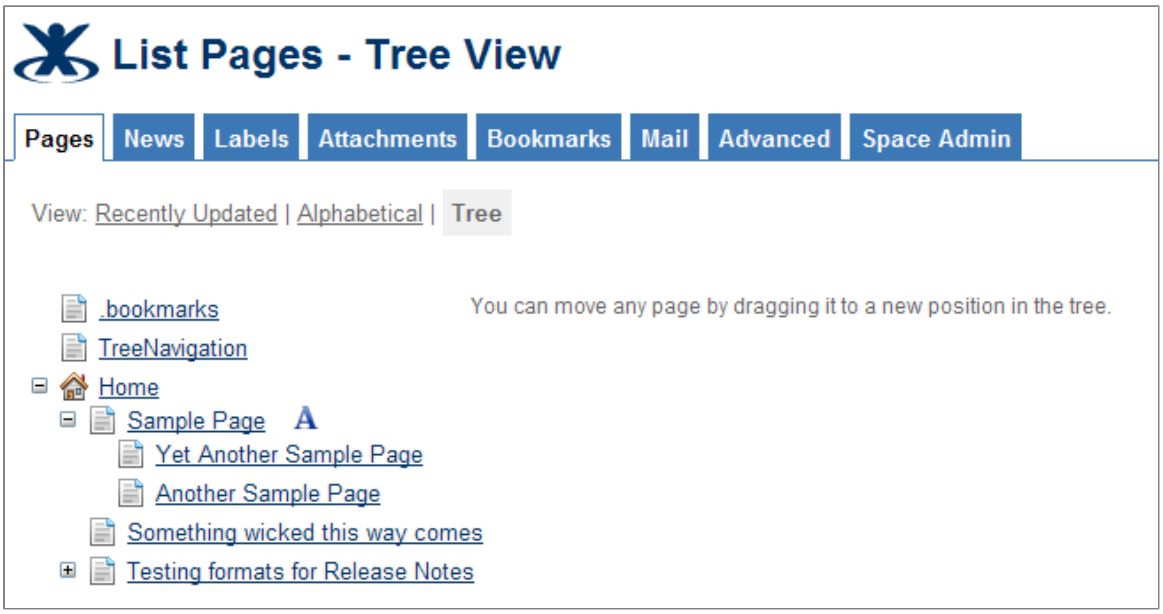
- The Space menu which appeared in Confluence 2.8 has now been renamed to 'Browse'. Faithful Confluence users will recognise and welcome this old friend.
- The Browse menu now includes the People Directory, Space administration and Confluence administration for authorised users.
- You can now reach your personal labels, watches and drafts directly from the dropdown menu under your name on any Confluence page.



8

Alphabetical Page Ordering

- By default Confluence orders your pages alphabetically, but you can drag and drop them into any order you like.
- Now Confluence 2.9 allows you to reset the page order to alphabetical, just by clicking the A icon next to the parent page.
- The documentation tells you more.



9

Better Spam Prevention

- A new link on the user profile screen allows administrators to jump directly to the user management screen — handy for dealing with those pesky spammers of the human variety.
- We have enhanced Confluence's Captcha functionality to strengthen the barrier against non-human spammers such as bots or web spiders.
- Captcha is now active on user profile pages as well as other pages.
- The image which Captcha displays is now even more difficult for non-humans to read. Take a look at the word 'brihter' in the screenshot below.
- You can read the instructions on configuring Captcha behaviour.



10

Plugin Repository

- The Plugin Repository shows more consistent information about the plugins on your Confluence site, including the system and bundled plugins which are shipped with Confluence.

11

Engine Room and Developers' Community

- Confluence's request throughput has been improved, thanks to better class- and resource-loading strategies.
- Startup time has also been reduced. This is particularly good for developers and anyone who needs to restart Confluence often.
- The Search API now allows you to filter search results by contributor.
- As part of the ongoing work to make it easier for internal developers and plugin developers to work on Confluence, this release includes re-factorings of some web action classes.
- To improve Confluence's resistance to cross site scripting security vulnerabilities, we have added an experimental automatic HTML entity encoding feature. This is the first step to providing a more secure product by default.

12

More than 140 Fixes and Improvements

- Take a look at the complete list of issues resolved in Confluence 2.9.

Special Thanks

We'd like to say thank you to Zohar Melamed and Shannon Krebs, who wrote the original Pagetree and Pagetree Search plugins.

The Confluence 2.9 Team

Development

Bugfixing and maintenance
Anatoli Kazatchkov
Brian Nguyen
Chris Broadfoot
Chris Kiehl
Don Willis

Design
Jason Taylor
Stephen Russell

Editor and page tree improvements
Agnes Ro
David Taylor
Dmitry Baranovskiy
Matt Ryall

Engine room
Andrew Lynch
Charles Miller
Christopher Owen
Matthew Jensen

Plugins
Ben Speakmon
Cheryl Jerozal
Jonathan Nolen
Rich Wallace
Ryan Talusan
Tim Moore

David Chui

Product management
Adnan Chowdhury

Search
David Loeng
Paul Curren

Team lead
Per Fragemann

Technical writing
Edwin Dawson
Sarah Maddox

Support

Kuala Lumpur
Arie Murdianto
Azwandi Mohd Aris
David Chui
Fennie Ng
Mei Yan Chan
Ming Giet Chong
Tony Cheah Tong Nyee

San Francisco
Jeremy Largman
Maleko Taylor
Tim Wong
Vincent Chang

Sydney
Ivan Benko
James Fleming
Partha Kamal

Renan Battaglin
Roy Hartono
Ulrich Kuhnhardt

Confluence 2.9 Upgrade Notes

Below are some essential notes on upgrading to **Confluence 2.9**. For details of the new features and improvements in this release, please read the Confluence 2.9 Release Notes.

On this page:

- Upgrade Notes
  - Plugins
  - Java Versions
  - Supported Databases and Application Servers
  - JavaScript Libraries
  - Confluence Themes
- Upgrade Procedure

Upgrade Notes

Plugins

Please check the following if you have added any plugins to Confluence:

- If you are using any third-party plugins, please test them thoroughly before rolling 2.9 into production.
- If you have installed the Gallery plugin onto your Confluence site, please remove it (or do not reinstall it after upgrading Confluence) in order to get the benefit of the new Gallery macro. The Gallery macro has been significantly improved in this release. It now incorporates features which were previously available only in the separate Gallery plugin.
- If you have installed the PageTree plugin or the Pagetree Search plugin onto your Confluence site, please remove them (or do not reinstall them after upgrading Confluence) in order to get the benefit of the new Pagetree macro. The Pagetree macro has been significantly improved in this release. It replaces the previous PageTree plugin, and incorporates the pagetree search option.
- If you are using the Blog Posts macro, be aware that invalid space keys will now be detected and cause the macro to fail. Previously they were ignored and blog posts from all spaces were returned.

Java Versions

Java 1.4 is not supported in Confluence 2.9 and later. Please refer to the Java 1.4 Support Timeline for more information.

Before upgrading to Confluence 2.9, you will need to ensure your environment is running at least Java 5. Confluence 2.9 supports Java 5 and Java 6. We recommend Java 6 because of its increased performance and easier troubleshooting, due to enhanced memory dump and profiling capabilities.

You can check your current Java version in Confluence:

1. Go to the Confluence '**Administration Console**'. To do this:
   - Open the '**Browse**' menu and select '**Confluence Admin**'. The 'Administration Console' view will open.
2. Select '**System Information**' from the '**Administration**' section in the left-hand panel.
3. Refer to '**Java Version**'.
   - If the version is 1.5 or higher, you do not need to do anything.
   - If the version is 1.4, you need to upgrade your JDK before you can upgrade to Confluence 2.9.

If you are running the Confluence EAR-WAR edition against your own application server, you will need to check with your application server vendor about which JDK versions are supported.

Supported Databases and Application Servers Additions to Application Server Support

We have added the following versions to the list of officially supported application servers:

- Resin 3.0 and 3.1
- Tomcat 6 (see known issue with authenticated datasources)

Platforms No Longer Supported

The following platforms were on the 'unsupported but working' list for Confluence 2.8. They are not on that list as of Confluence 2.9. They probably still work, but we have not tested them. Please upgrade soon.

- MySQL 4.0 — please upgrade to MySQL 5
- SQLServer 2000 — please upgrade to SQLServer 2005
- Websphere 6.0 — please upgrade to Websphere 6.1
- Weblogic 8.1 — please upgrade to Weblogic 9.2

Please refer to the list of supported application servers and databases on our Supported Platforms topic for the updated matrix.

Advance Notice — Changes to Supported Platforms in the Next Release

The next major release of Confluence after 2.9 will **not** support the following platforms/versions any more. We will still test those platforms infrequently, and Confluence will probably still work fine with them for a while, but they will be not officially supported.

- MySQL 4.1 — please upgrade to MySQL 5
- Tomcat 5.0 — please upgrade to Tomcat 5.5 or 6
- Resin 2 — please upgrade to Resin 3
- JBoss 4.0.x — please upgrade to JBoss 4.2.x

JavaScript Libraries

jQuery is the supported JavaScript library for plugin developers.

Advance notice — deprecated libraries: We have decided to standardise on jQuery as the JavaScript library for Confluence. This library will eventually replace all others. For this reason, use of the following JavaScript libraries in Confluence is **deprecated**:

- Prototype
- Scriptaculous

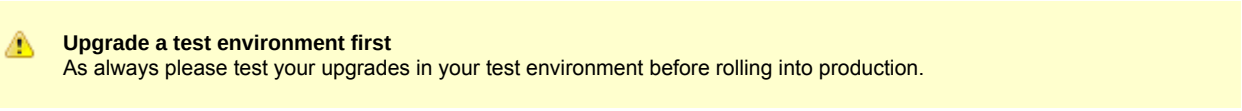
Because there is a lot of legacy plugin code using Prototype and Scriptaculous, these will continue to be available for this release of Confluence.

The **Yahoo! User Interface Library** (YUI) has been removed from this release, following its deprecation in Confluence 2.8.

Confluence Themes

Advance notice — The Confluence Classic Theme will be deprecated in a future release. Confluence 2.9 will be the last version that supports the Confluence Classic Theme. This theme uses outdated typography and formatting, which was replaced by the new-look Default Theme in Confluence 2.6. At that time, we introduced the Classic Theme to minimise the impact on customers who relied on the older typography. After Confluence 2.9, the Classic Theme will no longer be supported.

Upgrade Procedure



If you are already running a version of Confluence, please follow these instructions to upgrade to the latest version:

1. Before you upgrade, we strongly recommend that you back up your Confluence Home Directory and database. See the documentation on backing up your Confluence site. If you are using an external database, perform a database backup.
2. If your version of Confluence is earlier than 2.8.x, read the release notes and upgrade guides for all releases between your version and the latest version. In particular:
   - Please read the 2.8 upgrade notes.
   - If you are upgrading from 2.1 or earlier, please read the 2.2 release notes.
   - If you are upgrading from 2.2 or earlier, you will need to upgrade to Confluence 2.7.0 first, confirm the upgrade was successful, then upgrade again from version 2.7.0 to version 2.8.2. For more details, please refer to CONF-11767.
3. Download the latest version of Confluence.
4. Follow the instructions in the Upgrade Guide.
5. If you encounter a problem during the upgrade, please create a support ticket and one of our support engineers will assist you through the process.

RELATED TOPICS

Confluence 2.9 Release Notes

Issues resolved in Confluence 2.9

Below is the full list of issues resolved by Confluence 2.9. You can read the release notes here.

JIRA Issues (162 issues)				
Type	**Key**	**Summary**	**Status**	**Resolution**
	CONF-1783	Search by user and date/time range	Resolved	Fixed
	CONF-2328	Enable user to search content by author/editor	Resolved	Fixed
	CONF-3640	The mail tab should have a "Search mail" function.	Closed	Fixed
	CONF-5050	IE Crashes when Saving WYSIWYG Content	Resolved	Fixed
	CONF-5061	Space names and keys are not included in search results (not indexed)		Fixed

			Resolved	
	CONF-5383	Jira Portlet Macro Overrides CSS Styles	Resolved	Fixed
	CONF-5444	CSS in Jira Portlet macro overwrites Confluence default styles	Resolved	Fixed
	CONF-6227	Restore this version feature does not restore Page Name changes	Resolved	Fixed
	CONF-6929	Suppress "Unknown Ptg" warnings from re-index of unreadable documents by POI library	Resolved	Fixed
	CONF-7039	Search: Group by Space removes results from favourite spaces	Resolved	Won't Fix
	CONF-7468	/users/viewmydrafts.action is generating broken HTML	Resolved	Fixed
	CONF-7485	Rendering search results very slow if source document is large	Resolved	Fixed
	CONF-7638	Broken plugins disable Confluence on login with MethodInvocationException in StaticPlugin.getResourceDescriptors	Closed	Fixed
	CONF-7648	Disable Compatibility Macros (confluence.macros.compatibility) by Default	Resolved	Fixed
	CONF-8216	Non transaction cache updates should have clearer warning message	Closed	Fixed
	CONF-8217	When moving attachments, the "update links" option updates ALL pages with reference to ANY attachment on the page	Resolved	Fixed
	CONF-8637	Resin 3 responses are being concatenated sometimes when a custom logo is specified	Resolved	Fixed
	CONF-8699	Silent fail of Draft Saving can cause losing typed in information	Resolved	Fixed
	CONF-8756	HTML markup created from {table}{table-row}{table-cell} inserts extra <TD/>'s which interfere with macro-generated HTML contained in tables, unless SnipSnap compatibility macro is disabled.	Resolved	Fixed
	CONF-8783	Comments of attachments are not searchable	Resolved	Fixed
	CONF-8835	Change the sample URL for jiraissues macro usage in the Confluence's Notation Guide	Resolved	Fixed
	CONF-8923	"View changes" link missing from daily emails	Resolved	Fixed
	CONF-8975	Error creating Indexes during setup or site restore: "Column name 'ENTITY_ID' does not exist in target table"	Resolved	Fixed
	CONF-8998	Plugin inner dependency JARs fail to load if JAR does not contain directory entries	Closed	Fixed
	CONF-9050	Remove the space-list from the 404-error-page to reduce load on server	Resolved	Fixed
	CONF-9308	Unwanted Access to File System via Import Pages Functionality	Resolved	Fixed
	CONF-9343	Short wildcard queries fail with org.apache.lucene.search.BooleanQuery$TooManyClauses exception	Resolved	Fixed
	CONF-9390	Adding the same membership twice via UserAccessor throws Hibernate runtime exception	Resolved	Fixed
	CONF-9437	Incorrect Version Comment Behaviour when Reverting a Page to an old Version	Resolved	Fixed
	CONF-9490	Jira portlet (jiraportlet) macro doesn't rewrite relative stylesheet URLs	Resolved	Fixed

	CONF-9547	Viewing page source causes the source to be rendered	Resolved	Fixed
	CONF-9627	Velocity does not automatically escape HTML entities when substituting variables	Resolved	Fixed
	CONF-9677	Notifications of replies to comments contain unsubstituted value for author of original comment	Resolved	Fixed
	CONF-9706	Destination lost without any cookies	Closed	Fixed
	CONF-9720	Safari doesn't show a scroll bar for code panels that stretch wider than the window.	Resolved	Fixed
	CONF-9833	Search not working for Powerpoint or PDF files containing Japanese text	Resolved	Fixed
	CONF-9927	Personal Spaces should be included (or configurable to include) in quick search	Resolved	Fixed
	CONF-10016	Inconsistent search results for numbers within Excel files	Resolved	Fixed
	CONF-10116	Whitespace after name on view attachments page	Resolved	Fixed
	CONF-10135	ProfilingFilter writes to System.out instead of to the proper logfile	Resolved	Fixed
	CONF-10142	Improve performance by specifying a larger size for the Content Permission Set cache	Resolved	Fixed
	CONF-10266	"View changes link" in notification-email has unreplaced variable $originalBlogPost.id and does not work	Resolved	Fixed
	CONF-10365	Make Java 5 the minimum supported Java version	Closed	Fixed
	CONF-10371	Update pom to default to 1.5 compile target	Resolved	Fixed
	CONF-10372	Update system requirement documentation	Closed	Fixed
	CONF-10499	Page Ordering: Reorder Alphabetically	Closed	Fixed
	CONF-10512	Index optimization produce an OutOfMemoryError due to large textual attachments	Resolved	Fixed
	CONF-10539	Improve Confluence start up performance	Resolved	Fixed
	CONF-10710	Improve plugin provided WebWork action resolution speed	Closed	Fixed
	CONF-10775	Link Properties loses reference of the Alias in wiki-markup mode	Resolved	Fixed
	CONF-10838	Group by 'Type' returns Mail results when searching for All Content (No Mail)	Closed	Won't Fix
	CONF-10854	Accepting an Invitation leads to a NonUniqueObject Exception	Resolved	Fixed
	CONF-11023	Increase default database connection pool size from 15 to 30	Closed	Fixed
	CONF-11093	Add Drafts menu item to User menu in header	Closed	Fixed
	CONF-11156	Name of new page should be shown in the tree on edit	Resolved	Fixed
	CONF-11182	Confluence Demonstration Space needs enhancement	Closed	Fixed
	CONF-11231	Remove the YUI JavaScript-Library from Confluence		Fixed

			Resolved	
	CONF-11292	Page Ordering tree: Location label updates wrongly after clicking "done"	Resolved	Fixed
	CONF-11300	If you collapse a loading tree node the node stays collapsed	Resolved	Fixed
	CONF-11303	Improve Confluence startup time by avoiding Spring's createBean() post-processing	Resolved	Fixed
	CONF-11304	Upgrade xerces to latest to resolve resin XML parsing problem	Closed	Fixed
	CONF-11360	Exporting space in PDF - SOAP API doesn't interpret $dateFormatter, while exporting using a browser does	Resolved	Fixed
	CONF-11407	Mark core plugins as system plugins to prevent them from being disabled by administrators	Resolved	Fixed
	CONF-11410	Some users' logins are not remembered using Tomcat	Resolved	Fixed
	CONF-11438	Previewing a page with a tasklist shows error popup	Closed	Fixed
	CONF-11444	As 'Administrator' (not system admin) I cannot change parameter in General Configuration	Resolved	Fixed
	CONF-11462	Run as Service should not be a default - installer	Resolved	Fixed
	CONF-11482	Mail server MBean does not unregister correctly after changing name	Resolved	Fixed
	CONF-11489	Update bundled social bookmarking plugin	Closed	Fixed
	CONF-11494	Space Logos are stored and displayed using badly maintained mechanisms	Resolved	Fixed
	CONF-11503	Confluence Classic Theme - Tabs missing in User Preference mode.	Resolved	Fixed
	CONF-11528	Gallery macro is broken	Closed	Won't Fix
	CONF-11530	Overwriting changes from the editing conflict screen also abandons any additional changes made after the conflict was detected	Resolved	Fixed
	CONF-11539	JIRA issues macro should cache results for anonymous users when using a trusted connection	Resolved	Fixed
	CONF-11543	Improve log4j files	Resolved	Fixed
	CONF-11582	Macro for printable icon should not include tags	Resolved	Fixed
	CONF-11590	Space logo is not restored on import	Closed	Fixed
	CONF-11591	Confluence Fails to deploy into IBM Websphere Application Server with AnnotationException	Resolved	Fixed
	CONF-11602	user-breadcrumbs deleted in Confluence 2.8	Resolved	Fixed
	CONF-11613	Rendering of h4. tags in the PDF export sometimes has no formatting	Closed	Fixed
	CONF-11632	Remove "group by" functionality on search	Resolved	Fixed
	CONF-11737	Spaces tabs on dashboard are not displayed correctly using the Clickr Theme	Resolved	Fixed
	CONF-11743	Link from Page Tree node to Remove Page Action to enable easier removal of pages when		Fixed

		browsing space	Resolved	
	CONF-11765	"unable to find resource" error logged during an export	Resolved	Fixed
	CONF-11776	PDF export does not render panel macro background-color or title alignment	Resolved	Fixed
	CONF-11810	Add links to Order-form and Why-renew-page that notify customers of their expired maintenance	Resolved	Fixed
	CONF-11814	Double escaped ampersands in confluence feeds	Closed	Fixed
	CONF-11816	Improve performance by specifying a larger size for the UI Templates cache	Resolved	Fixed
	CONF-11821	Space decorators should only be loaded from suitable velocity resource loaders	Resolved	Fixed
	CONF-11822	Add captcha to "edit my profile"	Resolved	Fixed
	CONF-11824	Editing a page should not display the pagetree until the user chooses to edit the page's location.	Resolved	Fixed
	CONF-11825	Table layout issue in PDF generation	Closed	Fixed
	CONF-11831	Daily notification email: The space URL isn't interpreting ${space.key}	Resolved	Fixed
	CONF-11833	Global logos don't refresh/no longer display in Confluence 2.8	Resolved	Fixed
	CONF-11835	Users can't reply or edit comments in IE6 with the Default Theme	Resolved	Fixed
	CONF-11852	Config Page in Confluence 2.8.0 still refers to CONF-7620 even though it's fixed in 2.8	Resolved	Fixed
	CONF-11857	Performance degrades after upgrade - Some internal caches are always 1000 items big, reconfiguring them doesn't work	Resolved	Fixed
	CONF-11865	Add to new menu items to "User"-menu: Labels, Watches, Drafts	Resolved	Fixed
	CONF-11866	Improve captcha words and distortion of image	Resolved	Fixed
	CONF-11909	"View change"-link from email links to incorrect set of changes if the page got updated afterwards	Resolved	Fixed
	CONF-11921	Merging two versions with non-conflicting changes loses earlier changes	Resolved	Fixed
	CONF-11932	'Read more' link on blogposts contains unicode garbage	Resolved	Fixed
	CONF-11937	Just show 'Add Labels' when there are no labels	Closed	Fixed
	CONF-11938	While editing labels, Edit link should change to Done link	Closed	Fixed
	CONF-11942	Notation guide sometimes appears for rich text editor	Resolved	Fixed
	CONF-11950	Confluence should ship with transaction isolation level set to 1 for DB2	Resolved	Fixed
	CONF-11964	Add "People directory" menu item to "Browse" menu	Resolved	Fixed
	CONF-11972	Clustered installation stops sending notification emails	Resolved	Fixed
	CONF-11982	Some macros are displayed unstyled and ugly in Preview Mode		Fixed

			Resolved	
	CONF-11994	Race condition in HiLoIdRepairUpgradeTask and ResetHiLoAfterImportListener can lead to primary key violations in upgrade tasks	Closed	Fixed
	CONF-11996	Styles are lost in html, pdf exports and mail notifications	Resolved	Fixed
	CONF-12007	Visual rework of search page	Resolved	Fixed
	CONF-12013	Error page when visiting drafts	Resolved	Fixed
	CONF-12014	Upgrade JavaMail dependency to 1.4.1	Closed	Fixed
	CONF-12020	Export Space from remote API will produce an exception for anonymous users.	Resolved	Fixed
	CONF-12044	Javascript error when editing location for pages that are deep descendants and have been moved	Resolved	Fixed
	CONF-12055	Page Tree loses manual order after operations	Resolved	Fixed
	CONF-12067	HSQL warning should include proper hyperlink to the FAQ	Closed	Fixed
	CONF-12118	Create a Velocity logger to record all invocations on templates retrieved via Actions	Resolved	Fixed
	CONF-12121	Update RPC search documentation to include contributor field	Closed	Fixed
	CONF-12141	Search results with empty content body inherits exerpt from previous search result	Resolved	Fixed
	CONF-12155	TangosolClusteredFifoBuffer does not recover gracefully when items are dropped from underlying store	Resolved	Fixed
	CONF-12171	Production (standalone and war) distributions have old log4j.properties	Resolved	Fixed
	CONF-12172	Comment text is all rendered in italic when using French language pack	Resolved	Fixed
	CONF-12185	Numbered list renders differently in IE6/7 vs Safari 3.1.1 and Firefox	Resolved	Fixed
	CONF-12186	Members of confluence-admins can view restricted pages if they know the page URL	Closed	Fixed
	CONF-12197	Refactor and sanitize the action class hierarchy to make it easier to develop web action for Confluence	Resolved	Fixed
	CONF-12215	Moving a page with more than one level deep of children doesn't update the ancestor table correctly	Resolved	Fixed
	CONF-12228	Security Vulnerability in xwork, need to update to fixed version	Resolved	Fixed
	CONF-12230	Create Warning message in Admin console for users that have Daily Backup enabled.	Resolved	Fixed
	CONF-12247	Menu dropdowns can appear behind content	Resolved	Fixed
	CONF-12250	Tasks on Tasklist Cannot be deleted using Firefox 3.0	Closed	Fixed
	CONF-12252	Wysiwyg editor doesn't load up when gliffy plugin fails to load	Closed	Fixed
	CONF-12255	Include new demo space (DS) in Confluence download	Closed	Fixed
	CONF-12257	Extend breadcrumb generator to allow for custom breadcrumbs		Fixed

			Resolved	
	CONF-12260	Do not display profile userinfo items in recently updated when "Shared Mode" is enabled	Resolved	Fixed
	CONF-12271	Huge whitespace above editor when editing a page anonymously with the wiki markup editor	Resolved	Fixed
	CONF-12281	Incorporate the 2.8 translations into 2.9	Resolved	Fixed
	CONF-12293	Installing Confluence as a service on XP-64 fails due to 32-bit binaries of Tomcat	Closed	Fixed
	CONF-12335	Write developer documentation for HtmlSafe Velocity encoding	Closed	Fixed
	CONF-12339	Fire notification events when watches are added to or removed from pages and/or spaces	Closed	Fixed
	CONF-12341	Sort the list of installed plugins on the '500' error-page	Resolved	Fixed
	CONF-12360	Improve 500 error page formatting and markup	Resolved	Fixed
	CONF-12390	Improve startup time by not reloading complete xwork configuration when plugin state changes	Resolved	Fixed
	CONF-12391	HTML space export loses formatiing with tables	Closed	Fixed
	CONF-12393	Officially support Tomcat 6 by having automated build	Resolved	Fixed
	CONF-12401	Page view breaks if the 'recently-updated' macro is used with a non-existent 'types' parameter set	Resolved	Fixed
	CONF-12406	Labels should not be updated because they are added to a page, because this could potentially lead to a deadlock	Resolved	Fixed
	CONF-12432	Heading ones underlined	Resolved	Fixed
	CONF-12457	Document support for Tomcat 6	Closed	Fixed
	CONF-12461	Changing space layouts sometimes give NullPointerException error page	Resolved	Fixed
	CONF-12464	Incompatible plugins with velocity modules can crash Confluence	Resolved	Fixed
	CONF-12525	Upgrade Confluence to Use Atlassian-Plugins 2.0.0.beta6	Resolved	Fixed
	CONF-12590	Confluence path converter throwing ClassCastException as module uninstalled	Closed	Fixed
	CONF-12615	Blog calendar heading link changes colour on hover	Resolved	Fixed
	CONF-12655	Confluence velocity context used in decorators contains stale application and plugin state	Resolved	Fixed
	CONF-12684	Documentation readme points to almost EOL version of J2SE	Resolved	Fixed
	CONF-12694	Upgrading Confluence - check and update indexes.ddl file	Resolved	Fixed
	CONF-12752	Threads stuck attempting to index Excel attachments due to bug in POI	Resolved	Fixed
	CONF-12854	Lotus Notes html emails not applying styles	Closed	Fixed
	CONF-13782	Rendering search results very slow if source document is large in Confluence 2.8.2		Fixed

			Resolved	
	CONF-16400	Confluence 2.10 documentation for performance tuning has an incorrect path...	Resolved	Fixed

Confluence 2.8.3 Release Notes

Confluence 2.8.3 is the **standard edition** version of Confluence 2.8.x. This version of Confluence is equivalent to Confluence version 2.8.2, but it does not include Oracle's Coherence technology. Due to licensing changes with Oracle over the Coherence clustering and distributed caching technology, Atlassian is no longer able to distribute this technology to customers without a Confluence clustered license.

Therefore, Atlassian has released 'standard editions' for each previous major version of Confluence back to 2.6, which includes Confluence versions 2.10.4, 2.9.3, 2.8.3, 2.7.4 and 2.6.3. Confluence 2.8.3 will be the only Confluence 2.8.x version available to customers with a non-clustered Confluence license.

For its caching functionality, Confluence 2.8.3 utilises Ehcache. This functionality was provided by the Coherence technology in previous Confluence 2.8.x versions.

For more information about the features, updates and fixes in Confluence 2.8.2, please refer to the Confluence 2.8.2 Release Notes.

Upgrading from a Previous Version of Confluence

Upgrading Confluence should be fairly straightforward. Please read the Confluence 2.8.2 Upgrade Notes. However, if you have customised the cache settings in your installation of Confluence (e.g. for performance reasons), then please read the Confluence 3.0.1 Upgrade Notes for important information about transferring your caching layer customisations from Coherence to Ehcache.

We *strongly recommend* that you back up your `confluence.home` directory and database before upgrading.

Confluence 2.8.2 Release Notes

3 July 2008

Confluence 2.8.2 is a recommended upgrade which fixes some security flaws and other bugs. Please refer to the security advisory for details of the security vulnerabilities, risk assessment and mitigation strategies.

This release fixes the 'remember me' problem encountered when using Confluence with Tomcat 5.5.26 or Tomcat 6, where logins are not remembered.

Confluence administrators will see a new link on the user profile screen, allowing them to move directly to the user management screen for that user.

This release also addresses some performance bugs:

- Label links, as well as the label summary pages, now include the 'nofollow' attribute to prevent search engines like Google from indexing them.
- We have increased the size of the UI templates cache, which should reduce the number of times Confluence needs to load resources.
- CSS caching has been improved. (Refer to CONF-11755 if you'd like to know the details.)
- When generating a URL for the multiple label filter, Confluence now sorts the labels in the URL alphabetically, ensuring that there is just one URL for each filter instead of possible multiple URLs. This should reduce the load on the server when search engine crawlers visit the Confluence site, because the crawlers no longer need to index multiple URLs.

The JIRA Issues macro now has improved caching. When trusted communication was first introduced, Confluence did not cache results for anonymous users or results retrieved using trusted communications. With Confluence 2.8.2, caching is implemented for both those cases. This fix should improve the performance of your JIRA site, because Confluence no longer needs to access the JIRA site as often.

There's a complete list of fixes below. Click a specific issue to see details of the fix.

Don't have Confluence 2.8 yet?

Take a look at the new features and other highlights in the Confluence 2.8 Release Notes.

Download Latest Version

Upgrading from a Previous Version of Confluence

Upgrading Confluence should be fairly straightforward. Please read the Confluence 2.8.2 Upgrade Notes. We *strongly recommend* that you back up your `confluence.home` directory and database before upgrading.

Updates and Fixes in this Release

JIRA Issues (25 issues)					
Type	**Key**	**Summary**	**Priority**	**Status**	**Resolution**

	CONF-12228	Security Vulnerability in xwork, need to update to fixed version		Resolved	Fixed
	CONF-12056	Hot Referrers section in page info has broken links		Resolved	Fixed
	CONF-12011	Multiple-label filter generates redundant URLs		Closed	Fixed
	CONF-11985	XSS vulnerability in create/edit/copy page and blogpost actions		Resolved	Fixed
	CONF-11977	Customised Left Navigation Theme displays "test" next to user icon.		Resolved	Fixed
	CONF-11940	Add nofollow to label links so search engines like google don't index them		Resolved	Fixed
	CONF-11864	Sort order in People Directory makes no sense		Resolved	Fixed
	CONF-11833	Global logos don't refresh/no longer display in Confluence 2.8		Resolved	Fixed
	CONF-11831	Daily notification email: The space URL isn't interpreting ${space.key}		Resolved	Fixed
	CONF-11816	Improve performance by specifying a larger size for the UI Templates cache		Resolved	Fixed
	CONF-11776	PDF export does not render panel macro background-color or title alignment		Resolved	Fixed
	CONF-11765	"unable to find resource" error logged during an export		Resolved	Fixed
	CONF-11755	Caching of some CSS imported by combined.css isn't ideal, uses space cache counter instead of global		Closed	Fixed
	CONF-11587	Add link from User homepage to user admin page if current user is admin		Resolved	Fixed
	CONF-11579	Space Admin cause RuntimeException: Unable to find resource '/spaces/space-admin-breadcrumbs.vm'		Resolved	Fixed
	CONF-11539	JIRA issues macro should cache results for anonymous users when using a trusted connection		Resolved	Fixed
	CONF-11410	Some users' logins are not remembered using Tomcat		Resolved	Fixed
	CONF-11360	Exporting space in PDF - SOAP API doesn't interpret $dateFormatter, while exporting using a browser does		Resolved	Fixed
	CONF-11338	Custom global logo image height is not respected		Closed	Fixed
	CONF-11137	XSS vulnerability in pagepicker.action and spacepagepicker.action		Resolved	Fixed
	CONF-9273	headlines (h1, h2,..) included inside code macro are rendered from toc and toc-zone macro		Resolved	Fixed
	CONF-8975	Error creating Indexes during setup or site restore: "Column name 'ENTITY_ID' does not exist in target table"		Resolved	Fixed
	CONF-8835	Change the sample URL for jiraissues macro usage in the Confluence's Notation Guide		Resolved	Fixed
	CONF-8220	adding option of title to noformat macro breaks the toc macro and interferes with toc-zone macro		Resolved	Fixed
	CONF-4157	CamelCase links are not detected by Incoming Links		Resolved	Fixed

Confluence 2.8.2 Upgrade Notes

Confluence 2.8.2 is a recommended upgrade which fixes some security flaws as well as other bugs. You'll find details of the fixes in the release notes.

Upgrade Procedure

If you are already running a version of Confluence, please follow these instructions to upgrade to the latest version:

1. Before you upgrade, we strongly recommend that you **back up your `confluence.home` directory and database**. See the documentation on backing up your Confluence site.
2. If your version of Confluence is earlier than 2.8.0, read the release notes and upgrade guides for **all releases** between your version and the latest version. In particular:
   - Please read the 2.8 upgrade notes.
   - If you are upgrading from 2.1 or earlier, please read the 2.2 release notes.
   - If you are upgrading from 2.2 or earlier, you will need to upgrade to Confluence 2.7.0 first, confirm the upgrade was successful, then upgrade again from version 2.7.0 to version 2.8.2. For more details, please refer to CONF-11767.
3. Download the latest version of Confluence.
4. Follow the instructions in the Upgrade Guide.

Confluence 2.8.1 Release Notes

21 May 2008
Confluence 2.8.1 is a recommended upgrade which fixes some security flaws as well as other bugs in Confluence 2.8. Please refer to the security advisory for details of the security vulnerabilities, risk assessment and mitigation strategies.

Two of the bug fixes resolve problems with rendering a wiki page in Internet Explorer 6. Additionally, these fixes will significantly improve performance in some configurations of Confluence 2.8.

We have created some performance testing scripts, which will be particularly useful for large or mission-critical Confluence installations.

Using a custom space logo caused some problems, particularly with Resin application server, where the page would sometimes hang and then display incorrectly. This is now fixed.

There's a complete list of fixes below. Click a specific issue to see details of the fix.

Don't have Confluence 2.8 yet?
Take a look at the new features and other highlights in the Confluence 2.8 Release Notes.

Download Latest Version

Upgrading from a Previous Version of Confluence

Upgrading Confluence should be fairly straightforward. Please read the Confluence 2.8.1 Upgrade Notes. We *strongly recommend* that you back up your `confluence.home` directory and database before upgrading.

Updates and Fixes in this Release

JIRA Issues (47 issues)

Type	Key	Summary	Priority	Status	Resolution
	CONF-12153	Rollover help text for comment action links is misleading		Closed	Fixed
	CONF-11779	Collapse all comments link collapses add, reply and edit comment forms		Resolved	Fixed
	CONF-11774	IE6 clients download spacer.gif from wrong URL many times		Resolved	Fixed
	CONF-11625	Attachments view can't collapse		Resolved	Fixed
	CONF-11621	Numbered Lists getting cut off in tables when using IE 6 and IE 7		Resolved	Fixed
	CONF-11612	Add note to top of DOC and archived CONFnn spaces about documentation versioning		Closed	Fixed
	CONF-11605	Bad link in description for Plugin Repository plugin		Resolved	Fixed
	CONF-11585	Mailbox Import errors not escaped		Resolved	Fixed

	CONF-11582	Macro for printable icon should not include tags		Resolved	Fixed
	CONF-11532	Mailbox Import errors appear twice		Resolved	Fixed
	CONF-11524	XSS vulnerability in viewinfo.action		Resolved	Fixed
	CONF-11512	requests to fake.gif (defined in menu.css) causes SocketException in server when requested by IE6		Resolved	Fixed
	CONF-11491	Invalid tooltip on Edit | Remove | Reply in comments section - always says "Permanent link to this comment"		Closed	Fixed
	CONF-11485	backslash missing in ConfluenceActionSupport.properties		Resolved	Fixed
	CONF-11482	Mail server MBean does not unregister correctly after changing name		Resolved	Fixed
	CONF-11479	Add left-hand pagetree navigation to Confluence documentation DOC space		Closed	Fixed
	CONF-11468	Drop services and JavaWrapper from installer, or fix documentation		Resolved	Fixed
	CONF-11463	Error message styles lost in setup		Resolved	Fixed
	CONF-11462	Run as Service should not be a default - installer		Resolved	Fixed
	CONF-11456	Installer deletes old Confluence installation without warning		Resolved	Fixed
	CONF-11452	Users can move attachments to a space they have no permission for		Resolved	Fixed
	CONF-11444	As 'Administrator' (not system admin) I cannot change parameter in General Configuration		Resolved	Fixed
	CONF-11435	Restore and deprecate general-breadcrumbs.vm		Closed	Fixed
	CONF-11414	The first comment action at the bottom left should not have a border		Resolved	Fixed
	CONF-11412	Site Import leaves key generation strategy in an inconsitent state when running in a cluster		Resolved	Fixed
	CONF-11399	blog-posts macro did not render "Read More..." link after excerpt		Resolved	Fixed
	CONF-11385	Cancelling raise support request screen takes you to a blank page		Resolved	Fixed
	CONF-11378	Draft spacekey index name is too long for DB2		Resolved	Fixed
	CONF-11368	Provide load-testing scripts for Large or Mission-Critical Confluence Installations		Resolved	Fixed
	CONF-11358	Tools link appears in wrong place on first page load		Resolved	Fixed
	CONF-11331	OSUser entity migration fails due to case differences in database columns		Resolved	Fixed
	CONF-11325	Refreshing a page hangs and corrupts page results when a custom logo is used on Confluence running behind mod_jk or mod_proxy_ajp		Resolved	Fixed
	CONF-11300	If you collapse a loading tree node the node stays collapsed		Resolved	Fixed
	CONF-11297	Custom end of body HTML appears twice		Resolved	Fixed

	CONF-11292	Page Ordering tree: Location label updates wrongly after clicking "done"		Resolved	Fixed
	CONF-11183	Numbered list entries are too close to left-hand margin		Resolved	Fixed
	CONF-11156	Name of new page should be shown in the tree on edit		Resolved	Fixed
	CONF-11133	Changing passwords in Confluence does not update the user cache		Resolved	Fixed
	CONF-11073	Menus wrap when font size is increased		Resolved	Fixed
	CONF-10931	New add menu still appears when user has no permissions to add		Resolved	Fixed
	CONF-10854	Accepting an Invitation leads to a NonUniqueObject Exception		Resolved	Fixed
	CONF-10335	Mail connections do not specify a timeout, leading to monopolization of resources		Resolved	Fixed
	CONF-10142	Improve performance by specifying a larger size for the Content Permission Set cache		Resolved	Fixed
	CONF-9390	Adding the same membership twice via UserAccessor throws Hibernate runtime exception		Resolved	Fixed
	CONF-8637	Resin 3 responses are being concatenated sometimes when a custom logo is specified		Resolved	Fixed
	CONF-7760	Make the plus-sign a legal character for attachments		Resolved	Fixed
	CONF-2487	Attachments containing spaces download with + in their filenames.		Resolved	Fixed

Confluence 2.8.1 Upgrade Notes

Confluence 2.8.1 is a recommended upgrade which fixes some security flaws as well as other bugs. You'll find details of the fixes in the release notes.

Upgrade Procedure

If you are already running a version of Confluence, please follow these instructions to upgrade to the latest version:

1. Before you upgrade, we strongly recommend that you **back up your `confluence.home` directory and database**.
2. If your version of Confluence is earlier than 2.8.0, read the release notes and upgrade guides for **all releases** between your version and the latest version. In particular:
   - Please read the 2.8 upgrade notes.
   - If you are upgrading from 2.1 or earlier, please read the 2.2 release notes.
3. Download the latest version of Confluence.
4. Follow the instructions in the Upgrade Guide.

Confluence 2.8 Release Notes

10 April 2008

The Atlassian Confluence team is delighted to present **Confluence 2.8**.

We have simplified the screen design to focus on content. There's a totally new menu structure for editing or adding content and many other actions. We've also grouped the menu functions so that they are easier to use and understand. This will help new users start using Confluence quickly. For our faithful customers, you may find existing features you didn't know Confluence had!

The much voted for page-ordering feature allows you to define the order of your wiki pages yourself — just drag and drop your pages into the right position.

We've beautified the display of comments on pages and news items. With the multiple-label filter, you can combine more than one label in your label searches and surf labels directly via a sensible URL. Getting Confluence up and running is easier than ever with our new installer. An enhanced task list is bundled with Confluence, bringing faster response times and a simple yet powerful user interface.

There are some great performance enhancements and a lot for administrators and developers too.

- Thank you for all your issues and votes. Keep on logging, to help us keep improving!
- Below is a list of the highlights in this release.
- Attached is the full list of issues resolved in this release.

Upgrading to Confluence 2.8?
• Upgrading Confluence should be fairly straightforward. ***We strongly recommend that you back up your Confluence Home directory and your database before upgrading.*** • If you are using any **third-party plugins**, please test them thoroughly before rolling 2.8 into production. • Please refer to the Confluence 2.8 Upgrade Notes for further essential information about upgrading.

Responding to your feedback:

14 new feature requests implemented

300+ votes satisfied



Highlights of this release:

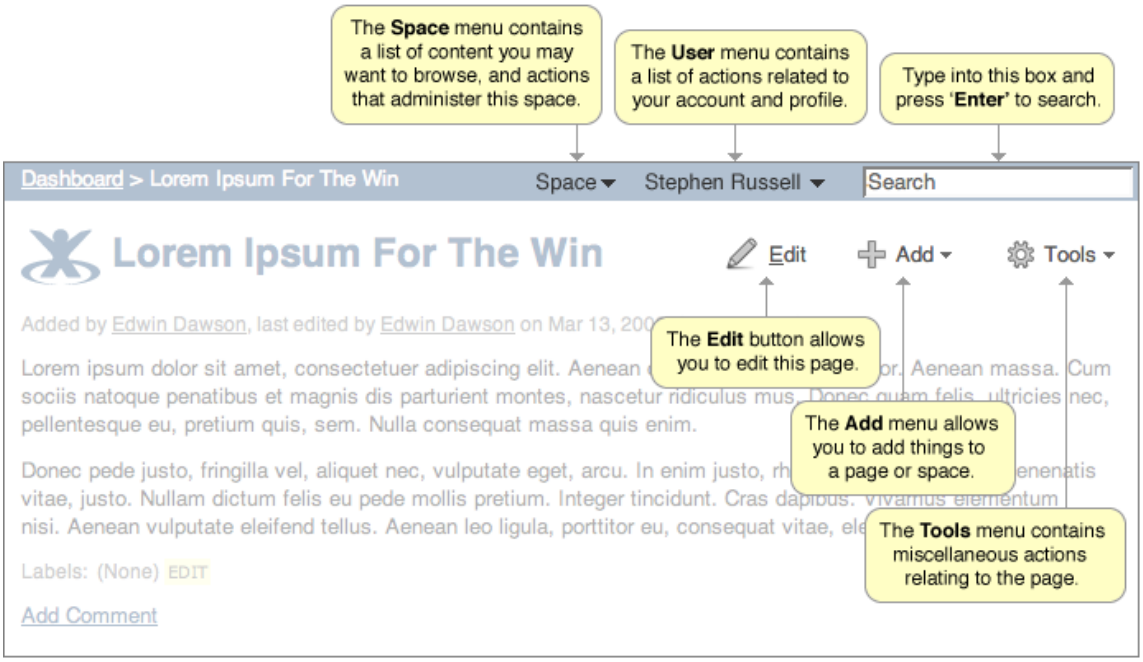
- Dynamic menus and simplified screen design
- Page ordering
- Collapsible comments
- Multiple-label filter
- Confluence installer
- Task list
- Performance enhancements
- Administration, management and monitoring
- Highlights for developers
- Over 100 fixes and improvements
- Special thanks

Highlights of Confluence 2.8



Dynamic menus and simplified screen design

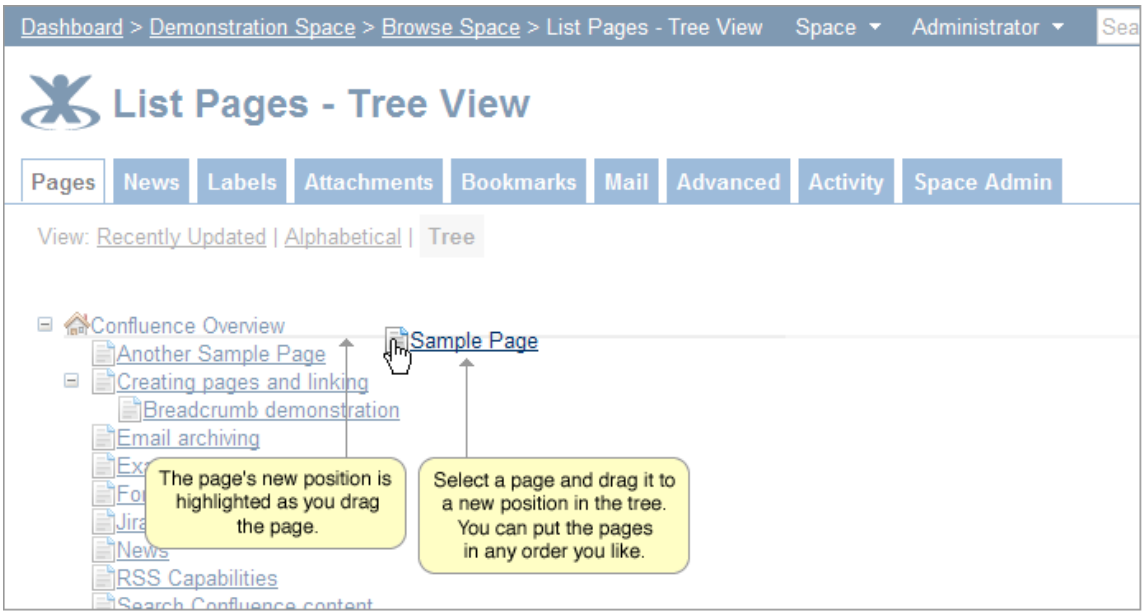
- Drop-down menus replace tabs and links.
- Grouping of functions is more intuitive.
- Features are more visible than before — you may even come across things you didn't know Confluence had!
- A cleaner screen design lets you focus on the page's content.
- It's simpler than ever to edit a page.
- There is more space for entering text.
- See an overview of the new menu structure and what's changed.





Page ordering

- One of the most popular feature requests is now a reality — Confluence allows manual page ordering.
- You can choose the order in which wiki pages are displayed.
- Use a dynamic tree view to drag and drop your pages into the right position.
- Page order is reflected in all tree views for a space or a page family, including the PageTree macro and exports to PDF, HTML and XML.
- Take a look at the documentation to learn more.

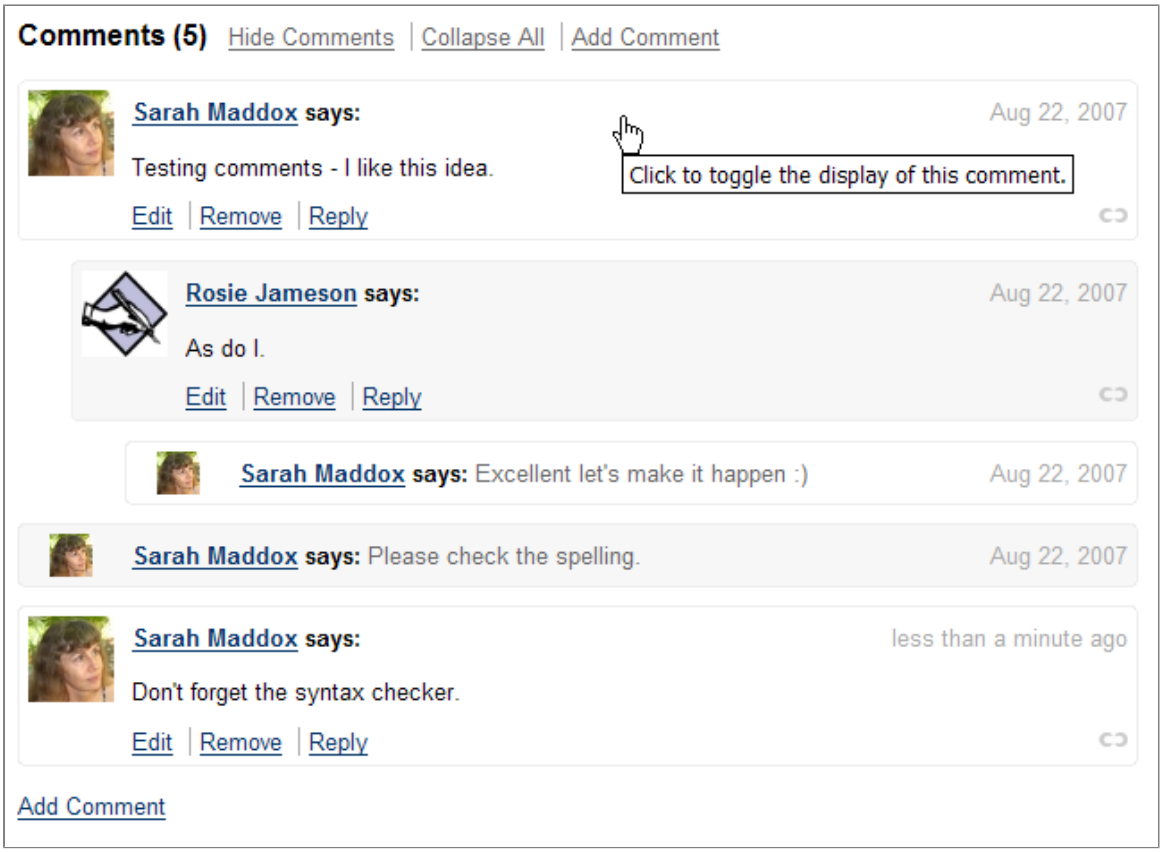




Collapsible comments

- We've beautified the display of comments on pages and news items.
- You can collapse comments to a single line.
- It's easier to keep track of multi-level comment threads.

- When collapsed, a comment shows a single-line excerpt.
- The time stamp is relative for recent comments — for example, it might say 'less than a minute ago'.
- The 'permanent link' icon is back, after a temporary absence in Confluence 2.6. Instead of lurking behind the date, it's now a link icon at the bottom right of each comment.



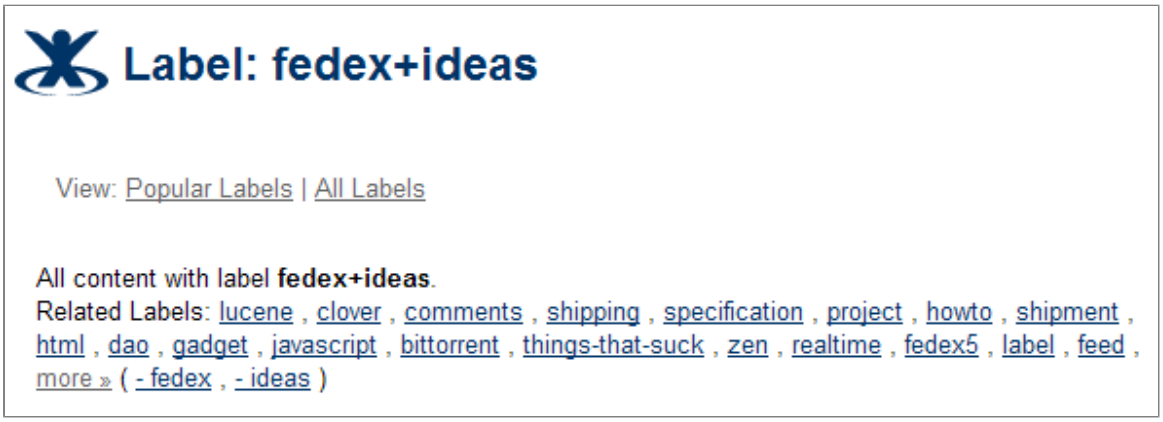


Multiple-label filter

- Combine more than one label in your label searches.
- Use the '-' link next to the label name to subtract a label from your multiple-label search.
- Browse labels simply by typing in a URL, such as:

```
1.http://CONFLUENCE_HOSTNAME/label/shpit+ideas
```

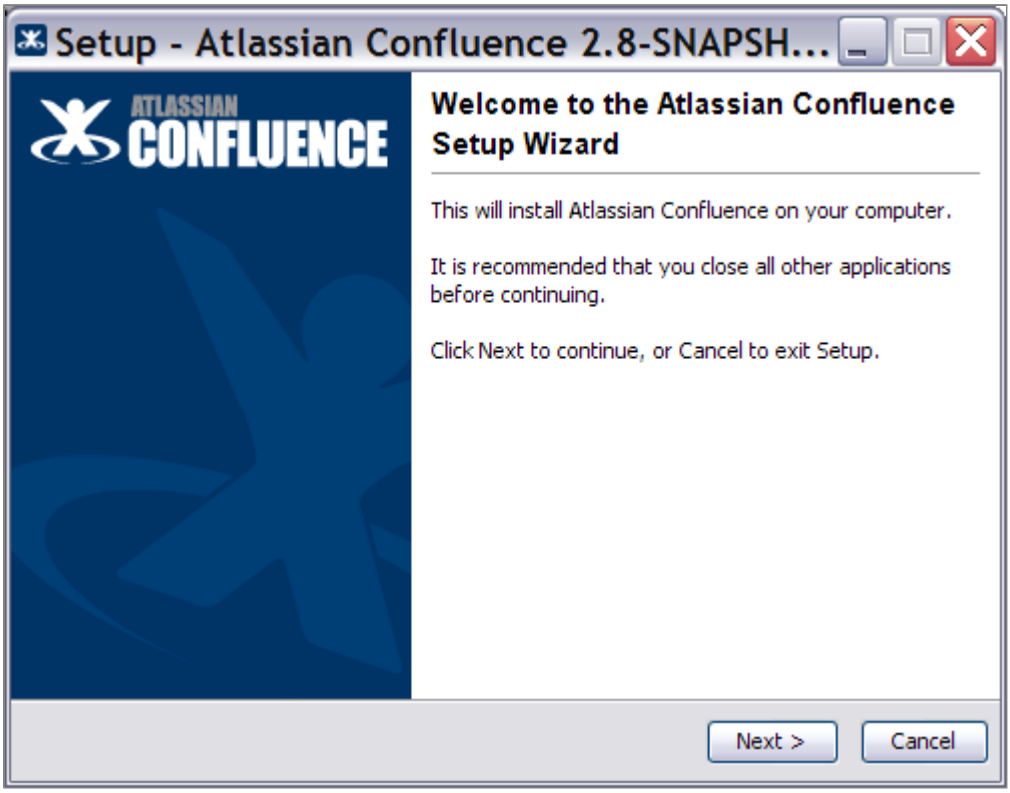
- Take a look at the documentation to learn more.





Confluence installer

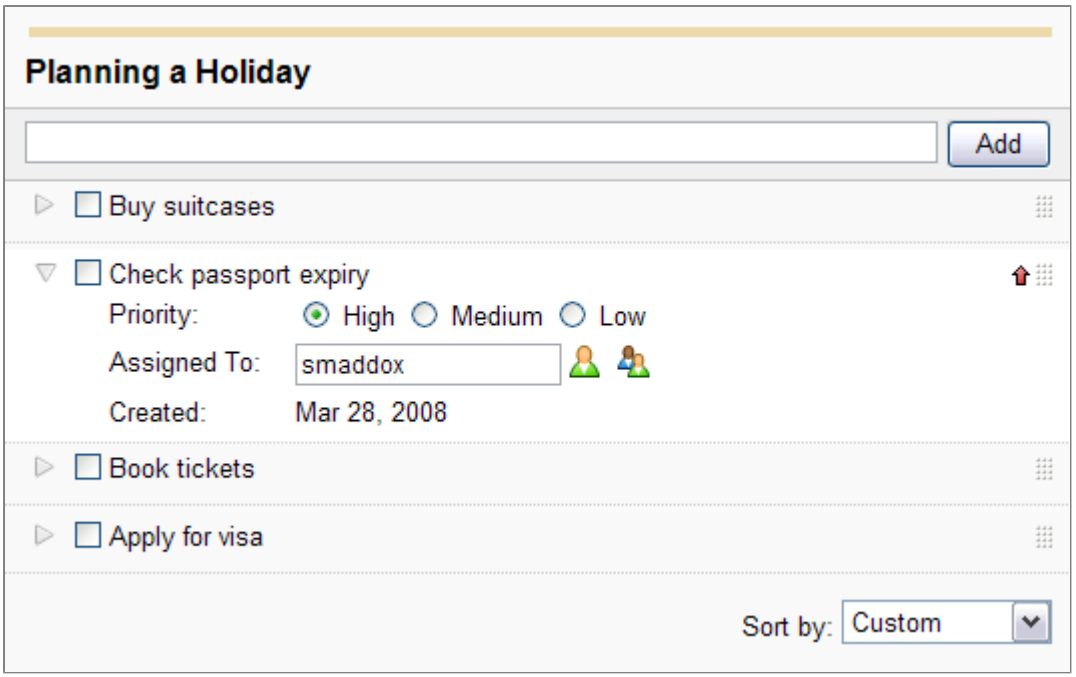
- The installation wizard lets you install Confluence Standalone edition without fuss or bother.
- Your Java environment is automatically configured.
- Confluence starts up in your browser after installation and leads you straight into the Setup Wizard.
- You can choose to install Confluence as a Windows service.
- Confluence appears in your Windows Start Menu.
- Read the documentation.



6

Task list

- The enhanced task list plugin is now bundled with Confluence.
- Permissions in the task list match the permissions of the page containing it.
- New user and group pickers help you to choose the right person or group.
- The sophisticated visual design suits the Confluence look and feel — take a look at the borders, icons, radio buttons and the handy visible cue on drag-and-drop.
- The data is stored in the page, and you can edit it there too.
- Performance on list operations is vastly improved.
- Existing task lists are upgraded on first edit. Read more in the upgrade notes.
- Here's the guide to using the macro.



7

Performance enhancements

- You will notice significant performance improvements in this release.
- A new gzip compression filter speeds up the transfer of data from Confluence and uses far less memory than the earlier implementation.
- The PDF space export uses less server memory.
- And more for the technically-minded:
  - The commonly-accessed resources use permanent client-side caching.
  - We have optimised some operations that were performed on every request, decreasing the average time taken to load a page.
  - We have optimised the database access for labels and attachments.
  - Where users belong to many groups, the retrieval of groups from LDAP is faster.
  - Access to Confluence's data storage mechanism (Bandana) has been made more granular, so that we don't load the entire context when retrieving individual keys from large contexts.

8

Administration, management and monitoring

- Confluence now supplies statistics and other information via a Java Management Extensions (JMX) interface. A number of third-party consoles will allow you to view the information and monitor your Confluence instance. There's more information in the documentation.
- The Plugin Repository client shows you which plugins are officially supported.
- You can raise a support request via the Administration Console.
- Troubleshoot your LDAP user management directly from the Administration Console.
- A new field on the System Information page displays the current access latency to the Confluence database — useful for diagnosing database network problems.

9

Highlights for developers

- Install custom path-mapping as part of your plugin, using the new Path Converter plugin modules — prettify your URLs.
- Include your JavaScript and CSS resources neatly, using the Web Resource plugin type.
- The Velocity template engine has been upgraded to version 1.5.
- We have begun a cleanup of the HTML and CSS for viewing and editing a page, moving towards semantic markup and web standards.
- We are moving towards standardisation on a single JavaScript library, based on jQuery. This is the supported JavaScript library for

plugin developers. The benefits? No more conflicts between libraries. jQuery is fast. And there's less to learn!



Over 100 fixes and improvements

- The People Directory uses the hCard microformat for simple integration with a variety of microformat-enabled tools.
- And more.

Special thanks

We'd like to thank Shannon Krebs, David Peterson and David Chui in particular for their continuing contributions to the open-source Confluence plugin library.

The Confluence 2.8 team

Development

Bugfixing and maintenance
Chris Broadfoot
David Loeng
Paul Curren

Engine room
Andrew Lynch
Charles Miller
Christopher Owen
Don Willis

Page ordering
Anatoli Kazatchkov
Matthew Jensen

Team lead
Per Fragemann

UI overhaul
Agnes Ro
David Taylor
Dmitry Baranovskiy
Matt Ryall

Design

Jason Taylor
Stephen Russell

Oversight and ~~mis~~management

Brett Jackson
Mike Cannon-Brookes
Scott Farquhar
Soren Harner
Tony Dagger

Product management

Adnan Chowdhury

Technical writing

Edwin Dawson
Sarah Maddox

Support

Kuala Lumpur
Mei Yan Chan
Ming Giet Chong
Tony Cheah Tong Nyee

San Francisco
Jeremy Largman

Tim Wong
Vincent Chang

Sydney
Ivan Benko
James Fleming
Kevin Tran
Partha Kamal
Roy Hartono
Ulrich Kuhnhardt

Confluence 2.8 Upgrade Notes

Below are some essential notes on upgrading to **Confluence 2.8**. For details of the new features and improvements in this release, please read the Confluence 2.8 Release Notes.

On this page:

- Upgrade Notes
  - Crowd Integration
  - Default Order of Pages Changed from Alphabetical to Natural
  - Plugins
  - Velocity Template Engine
  - Dynamic Tasklist 2
  - Customised Page Layouts
  - JavaScript Libraries
  - Java Versions
- Upgrade Procedure

Upgrade Notes

Crowd Integration

If you are using Atlassian Crowd as your Confluence user management and single sign-on solution, please upgrade to Crowd 1.3.2 or later. With Confluence 2.8 the atlassian-user interface has changed, and Crowd 1.3.2 provides the required update to Crowd's atlassian-user integration module.

Default Order of Pages Changed from Alphabetical to Natural

Confluence 2.8 introduces the ability to move pages into any order you choose.

As part of the above feature, we have changed the default page order in Confluence, from simple alphabetical ordering to a 'natural' ordering. The natural ordering handles numeric values correctly when doing string comparisons.

Impact:

- The new natural ordering is the same as the ordering already used by the PageTree plugin, which some customers use to create a left-hand navigation panel.
- The change to natural ordering should have little effect on most users because, under most situations, natural ordering and alphabetical ordering will produce the same results.
- For customers who have inserted chapter or page numbers to force the correct order, the new natural ordering will show the existing pages in the correct order.

 If you do find that the order of your pages is adversely affected, you can use the new page-ordering function to move the pages.

Plugins

Please check the following if you have added any plugins to Confluence:

- If you are using any third-party plugins, please test them thoroughly before rolling 2.8 into production.
- The PageTree plugin has been updated. If you are using this plugin, please download the latest version to ensure compatibility with Confluence 2.8.

Velocity Template Engine

Confluence's Velocity template engine has been upgraded from 1.3 to 1.5. Please test carefully for compatibility problems with existing third-party themes and plugins. For developers, there's more information about Migrating to Velocity 1.5.

Dynamic Tasklist 2

The Dynamic Tasklist 2 plugin is now bundled with Confluence. The new `tasklist` macro replaces the older `tasklist` and `dynamictasklist` macros. What happens to existing tasklists?

- By default, the new macro will be enabled and the older macros disabled in your Plugin Repository.
- When someone views a page containing an older version of the task list, the display will show the new format but the page will not be updated.
- When someone first adds a task or changes anything in the task list, the data will be converted to the new format.

Customised Page Layouts

The Confluence Upgrade Guide includes instructions on re-applying your customisations after the upgrade. We're repeating some of that information here, because it's particularly important due to the UI changes in this release.

If a space uses a customised page layout, the new Confluence 2.8 layout will not be applied. This means that you will not see the new menu structure within that space. For example, this will happen if you are using pagetree navigation to form a table of contents on the left.
Fix: Apply the Default Page Layout, then re-insert your custom code.

Steps in detail:

1. Go to the Space Admin screen and click 'Edit' to view your customised Page Layout.
2. Copy the customised code.
3. Cancel the edit.
4. Click 'Reset Default' to apply the new Confluence 2.8 default page layout.
5. On the Space Admin screen, click 'Create Custom' to create a custom page layout.
6. Reinsert your customised code and click 'Save'.

JavaScript Libraries

jQuery is the supported JavaScript library for plugin developers.

Advance notice — deprecated libraries: We have decided to standardise on jQuery as the JavaScript library for Confluence. This library will eventually replace all others. For this reason, use of the following JavaScript libraries in Confluence is **deprecated**:

- The Yahoo! User Interface Library (YUI)
- Prototype
- Scriptaculous

Because there is a lot of legacy plugin code using Prototype and Scriptaculous, these will continue to be available for at least one more major release of Confluence.

Java Versions

Confluence 2.8 supports Java 1.4, Java 5 and Java 6. We recommend Java 6 because of its increased performance and easier troubleshooting, due to enhanced memory dump and profiling capabilities.

Advance notice — Java 1.4 will be deprecated in a future release. Confluence 2.8 will be the last version that supports Java 1.4. Please refer to the Java 1.4 Support Timeline for more information.

Upgrade Procedure



As always please test your upgrades in your **TEST environment** before rolling into **PRODUCTION**.

If you are already running a version of Confluence, please follow these instructions to upgrade to the latest version:

1. Before you upgrade, we strongly recommend that you **back up** your Confluence Installation Directory and your Confluence Home Directory, as directed in the Upgrade Guide. If you are using an external database, perform a database backup.
2. If your version of Confluence is earlier than 2.7.x, read the release notes and upgrade guides for all releases between your version and the latest version.
    - If you are upgrading from 2.1 or earlier, please read the 2.2 release notes.
    - If you are upgrading from 2.2 or earlier, you will need to upgrade to Confluence 2.7.0 first, confirm the upgrade was successful, then upgrade again from version 2.7.0 to version 2.8. For more details, please refer to CONF-11767.
3. Download the latest version of Confluence.
4. Follow the instructions in the Upgrade Guide.
5. If you encounter a problem during the upgrade, please create a support ticket and one of our support engineers will assist you through the process.

RELATED TOPICS

Confluence 2.8 Release Notes

Issues resolved in Confluence 2.8

Below is the full list of issues resolved by Confluence 2.8. You can read the release notes here.

JIRA Issues (124 issues)				
Type	**Key**	**Summary**	**Status**	**Resolution**
	CONF-1031	Page ordering	Resolved	Fixed
	CONF-1075	User print-only stylesheet to make printing easier, more beautiful and more configurable	Resolved	Fixed

	CONF-3525	View News needs to be more prominent	Resolved	Fixed
	CONF-3719	Drag and Drop functionality in the Tree Viewer for moving pages	Resolved	Fixed
	CONF-4002	Reduce memory footprint of file-upload by writing uploaded file directly to disk	Resolved	Fixed
	CONF-4577	Filter by multiple labels	Resolved	Fixed
	CONF-4939	Send notifications when deleting content (e.g. when moving it to the trash) instead of sending them when purging the trash	Resolved	Fixed
	CONF-4961	Blog-posts macro limits number of results before applying label filter rather than after	Resolved	Fixed
	CONF-5019	Catch back button (of just leaving the page) when editing to ensure auto-save captures latest edits	Closed	Fixed
	CONF-5320	Javascript error (autosave) when saving page	Resolved	Fixed
	CONF-5663	Access keys for "search" and "save" conflict (both use "s"), this leads to inconsistent behavious across browsers	Resolved	Fixed
	CONF-5717	Confusing error message when viewing mail archive with no space key	Closed	Fixed
	CONF-5735	Improve handling of downtime of an LDAP repository	Closed	Fixed
	CONF-5795	Breadcrumbs shouldn't hyperlink the current page	Closed	Fixed
	CONF-5799	Hide trash when user can't remove page	Resolved	Fixed
	CONF-6532	Stack trace when LDAP user tries to reset password in "Forgot my password" action	Resolved	Fixed
	CONF-6702	Clean Confluence main stylesheet for readability	Resolved	Fixed
	CONF-6713	Clicking 'Edit' while editing loses recent changes - currently selected tab should not be clickable	Resolved	Fixed
	CONF-7409	Allow plugins to include Javascript and CSS resources in the HTML header	Resolved	Fixed
	CONF-7436	CSS selectors in themes can clash with defined Velocity macros	Resolved	Fixed
	CONF-7451	Many LDAP groups make 'Edit Groups' very slow	Resolved	Fixed
	CONF-7465	Confluence responds extremely slowly if LDAP connection times out	Closed	Fixed
	CONF-7620	Downloaded zip files with IE7 is corrupted (not able to open the archive) from confluence.atlassian.com	Resolved	Fixed
	CONF-8138	Notifications for file attachment removal states the file attachers name, not the file removers name	Resolved	Fixed
	CONF-8293	Removed content notifications should include "removed by" user	Resolved	Fixed
	CONF-8483	under 'Space > List Attachments', make 'Filter By File Extension' non-case-sensitive	Closed	Fixed
	CONF-8805	when you add a new image using wiki markup, then switch to the Rich Text Editor, 'width=32,height=32' gets appended to the new image	Closed	Fixed
	CONF-8921	Show All Comments Doesn't	Resolved	Fixed

	CONF-8971	Inefficient LDAP group lookup	Resolved	Fixed
	CONF-8972	NullPointerException when accessing a Confluence instance via tinyURL link wihtout identifier	Resolved	Fixed
	CONF-9046	Functional tests dump a load of Javascript errors to the logs	Closed	Fixed
	CONF-9093	Max label limit can be passed by adding labels via ajax	Resolved	Fixed
	CONF-9114	support for password protected proxies	Resolved	Fixed
	CONF-9206	Patch for resource handling to set proper expiry dates	Resolved	Fixed
	CONF-9220	Create a support request page similar to JIRA	Resolved	Fixed
	CONF-9241	Comment collapsing or hiding hides comment just added	Closed	Fixed
	CONF-9243	New spaces in the dashboard are listed out of order	Resolved	Fixed
	CONF-9291	Pluggable wiki renderers	Resolved	Fixed
	CONF-9302	Blog-Post macro parameter 'match-labels=all' incorrect if a label does not exist	Resolved	Fixed
	CONF-9304	File-uploading improvement: Remove dependency on pell-multipart library	Resolved	Fixed
	CONF-9419	Confluence 2.6-dr1 main-action.css tweaks	Resolved	Fixed
	CONF-9420	Combine duplicated ordered list CSS entries	Resolved	Fixed
	CONF-9422	Confusing heading CSS	Resolved	Fixed
	CONF-9428	Remove greylinks duplictation	Resolved	Fixed
	CONF-9429	More style duplication gardening...	Resolved	Fixed
	CONF-9430	Separate out left nav theme styles in to the left nav theme heh	Resolved	Fixed
	CONF-9433	Upgrade Atlassian Common modules for Confluence 2.7	Resolved	Fixed
	CONF-9462	main-action.css in conf 2.6-dr1 sets body to text-align:center!!!! RARGH!!! NASTY!!!	Closed	Fixed
	CONF-9554	Capture OutOfBounds when accessing users outside of the total list	Resolved	Fixed
	CONF-9571	Upgrade tangosol to a new version	Resolved	Fixed
	CONF-9609	Space key with "-" (hyphen) thows exception	Resolved	Fixed
	CONF-9640	Upgrade Confluence's Velocity dependency to 1.5	Resolved	Fixed
	CONF-9666	SQLException when trying to purge external referrers	Resolved	Fixed
	CONF-9695	Table styles in preview are different to normal page view	Closed	Fixed

	CONF-9723	remove comment notification wrongly claims that the comment author (or perhaps page author?) is the comment remover	Resolved	Fixed
	CONF-9750	IllegalStateException from JDK hook on shutdown	Closed	Fixed
	CONF-9790	Check for database latency and display it in system information page	Resolved	Fixed
	CONF-9806	Pagination on Attachments tab does not reflect filtered results	Closed	Fixed
	CONF-9832	Ability to 'Show Comment' even when No comment set	Resolved	Fixed
	CONF-9861	Provide JMX interface for Confluence core statistics	Resolved	Fixed
	CONF-9885	Admin group membership URL fails for groups with parentheses in group id	Resolved	Fixed
	CONF-9892	Notification mail for page-removal displays wrong user	Resolved	Fixed
	CONF-9893	Pagination in the User Picker(Search) is broken when using the group filter	Resolved	Fixed
	CONF-9895	Using the ampersand ('&') in user group name causes problems	Resolved	Fixed
	CONF-9930	Gzip filter (used for http-compression between client and server) creates very large temporary objects in memory	Resolved	Fixed
	CONF-9983	Editing News items (blog posts) allows you to save them without a title	Resolved	Fixed
	CONF-10024	Clean up breadcrumbs templates to use the new BreadcrumbManager	Resolved	Fixed
	CONF-10029	Include JDBC database info in automatic support request email	Resolved	Fixed
	CONF-10037	Alphabetical View of a space only displays pages that start with the letter 'a'	Resolved	Fixed
	CONF-10213	XML RPC server uses platform default character encoding to decode requests	Resolved	Fixed
	CONF-10222	"Diff" option in the RSS Feed Builder does not take effect	Resolved	Fixed
	CONF-10250	Confusion when using Edit's "Shortcut Key" in Confluence and end of line key for a MacOS machine	Resolved	Fixed
	CONF-10260	Blank daily notification emails when the wiki content has not changed at all, e.g. on weekends	Resolved	Fixed
	CONF-10275	Sort order of People Directory is not alphabetical for case sensitive entries	Resolved	Fixed
	CONF-10305	Multimedia *.swf files are not able to render in Internet Explorer browser	Resolved	Fixed
	CONF-10392	Confluence javac should not fork	Closed	Fixed
	CONF-10394	Bloated logfiles due to warning message from [apache.commons.httpclient.HttpMethodBase] getResponseBody()	Closed	Fixed
	CONF-10402	View Problems with IE while editing with RTE	Resolved	Fixed
	CONF-10449	No whitespace before the word 'updated' on comments	Resolved	Fixed
	CONF-10463	RPC method 'convertToPersonalSpace' ignores parameter 'newName' for space's new name	Resolved	Fixed

	CONF-10467	Multimedia WMV files embedded into a Confluence page fail to load	Closed	Fixed
	CONF-10484	Update http compression documentation	Resolved	Fixed
	CONF-10524	Unable to edit an attachment comment to blank (effectively removing it)	Resolved	Fixed
	CONF-10552	Improve configuration loading performance in ConfluenceCachingBandanaPersister	Resolved	Fixed
	CONF-10554	Improve instructions in README text and HTML files in Confluence Standalone zip	Resolved	Fixed
	CONF-10574	Typo in attachment logging statement - "Cound" instead of "Could"	Closed	Fixed
	CONF-10582	"Printable" version of a page cuts off text when printed.	Resolved	Fixed
	CONF-10584	Integrate LDAP troubleshooting tool ('Paddle') into Confluence	Resolved	Fixed
	CONF-10587	Include Path Conversion Plugin Points	Resolved	Fixed
	CONF-10605	Groups listed in user profile should link to group search page	Resolved	Fixed
	CONF-10611	Improve user listing navigation and make number of users displayed configurable	Resolved	Fixed
	CONF-10645	class cast exception when viewing users in groups and Userlister macro throws ClassCastException when using osuser for user management.	Resolved	Fixed
	CONF-10650	No groups are displayed under the "Manage Groups" page if Confluence fails to query an LDAP repository	Resolved	Fixed
	CONF-10659	CLONE -pagination problem: 5 pages found but only first 4 are linked linked	Resolved	Fixed
	CONF-10660	Renaming a blogpost to use an illegal character causes exception instead of a simple error/try-again page	Resolved	Fixed
	CONF-10686	Some space permissions lost after upgrade to 2.7.1	Resolved	Fixed
	CONF-10736	Anonymous users can't remove news / blogs even though there is an explicit permissions allowing them to do so	Resolved	Fixed
	CONF-10739	Classic Theme and Attachment Tab	Closed	Fixed
	CONF-10777	Stop processing views of historical content in the usage tracking plugin	Resolved	Fixed
	CONF-10851	Renderer component plugin documentation	Closed	Fixed
	CONF-10905	Can edit a news post to have the same title as another post in the same day	Resolved	Fixed
	CONF-10910	Confluence Setup Wizard points to incorrect documentation on CAC	Resolved	Fixed
	CONF-10912	Improve comment threading	Resolved	Fixed
	CONF-10913	Remove GlobalHelper dependency from Condition implementations	Resolved	Fixed
	CONF-10936	Incorporate improved German and French translations	Resolved	Fixed
	CONF-11016	Toggle area for collapsing comment should have a tooltip	Closed	Fixed

	CONF-11040	Grouppicker and Userpicker display unescaped user-entered content	Resolved	Fixed
	CONF-11058	Latest Confluence API javadoc not available (404 Not Found)	Resolved	Fixed
	CONF-11081	URL not encoded for group browser	Resolved	Fixed
	CONF-11139	Improve custom profile pic detection logic in templates	Closed	Fixed
	CONF-11144	Attachments deleted due to google crawled and no 'nofollow' link	Resolved	Fixed
	CONF-11147	Results of query for latest versions of attachments should be cached	Closed	Fixed
	CONF-11152	User with null email address breaks daily report job	Resolved	Fixed
	CONF-11215	Use hCard microformat for people directory	Closed	Fixed
	CONF-11239	Anonymous rpc access causes large amount of WARN messages in logs	Resolved	Fixed
	CONF-11259	Cannot switch language in the setup wizard prior to the database setup step	Resolved	Fixed
	CONF-11274	Confluence should display server ID on "outdated license" page	Resolved	Fixed
	CONF-11279	Adding a comment on the clickr theme produces a javascript alert of a SQL exception	Resolved	Fixed
	CONF-11310	Please change destination email address of "Support Request" page	Resolved	Fixed
	CONF-11328	Improve label for license generation in Setup Wizard	Resolved	Fixed
	CONF-11483	new version 2.8 paddle jar throws ClassNotFound exception	Closed	Fixed
	CONF-12362	Upgrade c3p0 to 0.9.1 or higher to avoid Acquisition Attempt Failed error	Resolved	Fixed
	CONF-12375	Add support for specifying css media for plugins	Resolved	Fixed
	CONF-14992	When hitting th ecancel button after starting to add a comment, the page no longer shows 0 comments	Resolved	Fixed

Confluence 2.8 Screen and Menu Changes

Can't find a menu item in Confluence 2.8? That's understandable, because the user interface (layout of the screens and menus) in Confluence 2.8 has changed dramatically. This page contains a quick summary of the changes. Please read on for a quick and easy introduction.

Refer to Using the Confluence Screens for an overview of the new Confluence menu structure.

On this page:

- What's on this page?
- Confluence 2.8 User Interface in Detail
- New Menu Structure in Confluence 2.8
- Changes to Menu Items in Confluence 2.8
  - RELATED TOPICS

Some of the Confluence documentation will not immediately reflect the new user interface in Confluence 2.8. We are working to rectify the situation as quickly as possible. In the interim, please use this page as a guide.

What's on this page?

The image below shows a Confluence 2.8 screen, with numbers added to the image. The numbers relate to a table of menu commands below the image, showing the new menu structure. Finally, there is a comprehensive table of menu commands, their new locations and how to get there.

Confluence 2.8 User Interface in Detail

The following screenshot and numbered list show the new menu layout in Confluence 2.8.

Screenshot: Confluence 2.8 Annotated User Interface Changes



New Menu Structure in Confluence 2.8

The numbers in this list relate to the numbers added onto the image above.

1	'**Space Menu**' containing Pages, News, Labels, Attachments, Bookmarks, Mail, Advanced and Space Admin.
2	'**User Menu**' containing Personal Space, Preferences, History, Administration (Confluence Administrators only) and Log Out.
3	'**Edit Button**' which opens the current page for editing.
4	'**Add Menu**' containing Page, News, Comment, Attachment, Add Bookmark. There may also be other items in this menu, depending upon the plugins you have installed on your Confluence site. For example, Add Diagram (Gliffy plugin) and Add Spreadsheet (EditGrid plugin).
5	'**Tools Menu**' containing Attachments, History, E-mail, Favourite, Watch, Info, View Wiki Markup, Export to PDF, Export to Word, Copy, Move and Remove.
6	'**Labels**'. This line shows tags attached to the current page.

Changes to Menu Items in Confluence 2.8

This table lists every change to the menu structure that has occurred in Confluence 2.8. Names for some commands have changed, so the old menu option names are listed, next to the new names, with instructions for finding them. (You can search the text on this page for the name of a menu item you used to use in an earlier version of Confluence – you'll find instructions next to it.)

Old items, now under Space Menu	New instructions
Browse Space — Pages	Click '**Space**', '**Pages**'

Browse Space — News	Click '**Space**', '**News**'
Browse Space — Labels	Click '**Space**', '**Labels**'
Browse Space — Attachments	Click '**Space**', '**Attachments**'
Browse Space — Bookmarks	Click '**Space**', '**Bookmarks**'
Browse Space — Mail	Click '**Space**', '**Mail**'
Browse Space — Advanced options	Click '**Space**', '**Advanced**'
Space Admin	Click '**Space**', '**Space Admin**'
Old items, now under Add Menu	**New instructions**
Add Page	Click '**Add**', '**Page**'
Add News	Click '**Add**', '**News**'
Add Comment	Click '**Add**', '**Comment**'
Attachments tab	Click '**Add**', '**Attachment**'
Bookmarks	Click '**Add**', '**Add Bookmark**'
Add other items to the page, as provided by plugins such as Gliffy diagrams and spreadsheets	Click '**Add**' and then select the relevant option (available only if you have the relevant plugin installed)
Add Spreadsheet	Click '**Add**', '**Add Spreadsheet**'
Old items, now under Tools Menu	**New instructions**
Attachments tab	Click '**Tools**', '**Attachments**'
View Page History	Click '**Tools**', '**History**'
Email the page to someone	Click '**Tools**', '**E-mail**'
Favourite button	Click '**Tools**', '**Favourite**'
Watch button	Click '**Tools**', '**Watch**'
Info tab	Click '**Tools**', '**Info**'
View Wiki Markup	Click '**Tools**', '**View Wiki Markup**'
View a printable version of the current page	Click '**File**', '**Print Preview**' (in your browser menu)
Export to PDF	Click '**Tools**', '**Export to PDF**'
Export to Word	Click '**Tools**', '**Export to Word**'
Copy (page)	Click '**Tools**', '**Copy**'
Move (page)	Click '**Tools**', '**Move**'
Remove (page)	Click '**Tools**', '**Remove**'
Miscellaneous page element(s)	**New instructions**
Labels	(moved to the bottom of page content)

RELATED TOPICS

Using the Confluence Screens

Take me back to Confluence User's Guide

Confluence 2.8 Beta Release Notes

28 March 2008

Confluence 2.8 will be launched in a few weeks' time. These release notes apply to **Confluence 2.8 Beta**, which is currently undergoing internal testing. These release notes show the highlights of the upcoming release, although we have not included all the enhancements and bug fixes. We'll publish the final and complete release notes with the release of Confluence 2.8.0.

If you are interested in trying out a developer's release, please take a look at the information and warnings in Development Releases.

What's Coming in Confluence 2.8

1 Page Ordering

- One of the most popular feature requests is now reality — Confluence allows manual page ordering.
- You can choose the order in which wiki pages are displayed.
- Use a dynamic tree view to drag and drop your pages into the right position.
- Page order is reflected in all tree views for a space or a page family, including the PageTree macro and exports to PDF, HTML and XML.
- Read the documentation.

2 Dynamic Menus and Simplified Editing

- Drop-down menus replace tabs and links.
- Grouping of functions is more intuitive.
- Features are more visible than before — you may even come across things you didn't know Confluence had!
- Simplified screen design allows you to focus on editing the page content.
- There is more space for entering text.
- See an overview of the new menu structure and what's changed.

3 Collapsible Comments

- We've beautified the display of comments on pages and news items.
- You can collapse comments to a single line — just click the subject line of a single comment, or 'Collapse All'.
- It's easier to keep track of multi-level comment threads.
- When collapsed, a comment shows a single-line excerpt.
- The time stamp is relative for recent comments — for example, it might say 'less than a minute ago'.
- The new 'permanent link' icon makes it easier to link directly to a comment from another page.

4 Multiple Label Filter

- Combine more than one label in your label searches.
- Use '**+**' and '**-**' links to add or subtract labels from your search.
- Browse labels simply by typing in a URL, such as:

```
1.http://CONFLUENCE_HOSTNAME/label/shpit+ideas
```

- Read the documentation.

5 And Lots More

- Enhanced Dynamic Tasklist.
- An installation wizard for Confluence Standalone edition on Windows and Mac.
- Significant performance improvements.
- More administration, management and monitoring tools.
- Plenty of bug fixes to keep everyone happy.

Confluence 2.7.4 Release Notes

Confluence 2.7.4 is the **standard edition** version of Confluence 2.7.x. This version of Confluence is equivalent to Confluence version 2.7.3, but it does not include Oracle's Coherence technology. Due to licensing changes with Oracle over the Coherence clustering and distributed caching technology, Atlassian is no longer able to distribute this technology to customers without a Confluence clustered license.

Therefore, Atlassian has released 'standard editions' for each previous major version of Confluence back to 2.6, which includes Confluence versions 2.10.4, 2.9.3, 2.8.3, 2.7.4 and 2.6.3. Confluence 2.7.4 will be the only Confluence 2.7.x version available to customers with a non-clustered Confluence license.

For its caching functionality, Confluence 2.7.4 utilises Ehcache. This functionality was provided by the Coherence technology in previous Confluence 2.7.x versions.

For more information about the features, updates and fixes in Confluence 2.7.3, please refer to the Confluence 2.7.3 Release Notes.

Upgrading from a Previous Version of Confluence

Upgrading Confluence should be fairly straightforward. Please read the Confluence 2.7.3 Upgrade Notes. However, if you have customised the cache settings in your installation of Confluence (e.g. for performance reasons), then please read the Confluence 3.0.1 Upgrade Notes for important information about transferring your caching layer customisations from Coherence to Ehcache.

We *strongly recommend* that you back up your `confluence.home` directory and database before upgrading.

Confluence 2.7.3 Release Notes

19 March 2008
Confluence 2.7.3 is a recommended upgrade which focuses on fixing a number of security flaws. Please refer to the security advisory for details of the vulnerabilities, risk assessment and mitigation strategies.

There's a complete list of fixes below. Click a specific issue to see details of the fix, and to download patches where relevant.

Don't have Confluence 2.7 yet?
Take a look at the new features and other highlights in the Confluence 2.7 Release Notes.

Download Latest Version

Upgrading from a Previous Version of Confluence

Upgrading Confluence should be fairly straightforward. Please read the Confluence 2.7.3 Upgrade Guide. We *strongly recommend* that you back up your `confluence.home` directory and database before upgrading.

Updates and Fixes in this Release

JIRA Issues (13 issues)

Type	Key	Summary	Priority	Status	Resolution
	CONF-11316	Removing user throws NullPointerException		Closed	Fixed
	CONF-11153	XSS vulnerability in social bookmarking plugin bundled in Confluence		Resolved	Fixed
	CONF-11149	XSS vulnerability in browseusers.vm		Resolved	Fixed
	CONF-11141	XSS vulnerabilities in insert image and link actions		Resolved	Fixed
	CONF-11081	URL not encoded for group browser		Resolved	Fixed
	CONF-11042	XSS vulnerabilities in create space action		Resolved	Fixed
	CONF-11040	Grouppicker and Userpicker display unescaped user-entered content		Resolved	Fixed
	CONF-11027	XSS vulnerabilities in create/edit/copy page and blogpost actions		Resolved	Fixed
	CONF-11026	username not validated in add user to favourites action		Resolved	Fixed
	CONF-11019	Fix XSS vulnerabilities in the stacktraces and cause by's displayed on the 500 error page		Resolved	Fixed
	CONF-11005	XSS vulnerability in signup actions		Resolved	Fixed
	CONF-11002	viewuser.action has an XSS problem around username		Resolved	Fixed
	CONF-9559	Cross-site scripting vulnerability in /dashboard.action		Resolved	Fixed

Confluence 2.7.3 Upgrade Guide

Confluence 2.7.3 is a recommended upgrade which focuses on fixing a number of security flaws. You'll find details of the fixes in the release notes.

Upgrade Procedure

If you are already running a version of Confluence, please follow these instructions to upgrade to the latest version:

1. Before you upgrade, we strongly recommend that you **back up your `confluence.home` directory and database**.
2. If your version of Confluence is earlier than 2.7.2, read the release notes and upgrade guides for **all releases** between your version and the latest version. In particular:
   - Please read the 2.7 upgrade notes.
   - If you are upgrading from 2.1 or earlier, please read the 2.2 release notes.
3. Download the latest version of Confluence.
4. Follow the instructions in the Upgrade Guide.

Confluence 2.7.2 Release Notes

6 March 2008
Confluence 2.7.2 is a recommended upgrade which fixes a security flaw and other bugs.

We have fixed a security flaw which allowed users who have 'View' permission (or higher) on a space to purge any page in that space. For details, please refer to the security advisory and related JIRA issue.

Other good fixes in this point release:

- The Rich Text editor no longer breaks links to pages in other spaces.
- Some customers have reported problems with permissions after upgrading to Confluence **2.7.1**, where some space permissions or global permissions were lost if using a case-sensitive database. From Confluence **2.7.2**, the space permissions and global permissions screens will display a message highlighting any case-sensitivity problems. We have also provided a routine to fix existing permissions affected by this issue — read the detailed instructions on running the routine.
- The JIRA Portlet macro now displays correctly when using trusted communication between JIRA and Confluence.
- Confluence 2.7.2 also clarifies the procedures around renewing your license before upgrading.

There's a complete list of fixes below. Click a specific issue to see details of the fix, and to download patches where relevant.

Don't have Confluence 2.7 yet?
Take a look at the new features and other highlights in the Confluence 2.7 Release Notes.

Download Latest Version

Upgrading from a Previous Version of Confluence

Upgrading Confluence should be fairly straightforward. Please read the Confluence 2.7.2 Upgrade Guide. We *strongly recommend* that you back up your `confluence.home` directory and database before upgrading.

Updates and Fixes in this Release

JIRA Issues (23 issues)

Type	Key	Summary	Priority	Status	Resolution
	CONF-11638	Formatting of Recently-updated macro broken at the space level if not using Default Confluence Theme		Closed	Fixed
	CONF-10807	Users with view permissions on a space are able to delete (purge) pages they don't have permission to edit/access		Resolved	Fixed
	CONF-10777	Stop processing views of historical content in the usage tracking plugin		Resolved	Fixed
	CONF-10686	Some space permissions lost after upgrade to 2.7.1		Resolved	Fixed
	CONF-10621	Improve logging in DefaultSpacePermissionManager		Resolved	Fixed
	CONF-10534	Fix documentation links in EAR/WAR README file		Resolved	Fixed
	CONF-10467	Multimedia WMV files embedded into a Confluence page fail to load		Closed	Fixed
	CONF-10465	JIRA Portlet Macro not displaying when authenticating using the trusted application between JIRA and Confluence		Resolved	Fixed
	CONF-10458	WYSIWYG rich text editor breaks links to pages		Resolved	Fixed
	CONF-10431	Userlister macro no longer seems to work if you don't supply a "group" parameter		Resolved	Fixed
	CONF-10410	Distortion of dashboard while using Left Navigation Theme		Resolved	Fixed
	CONF-10399	formatting problems of recentlyupdated.action with leftnavigation theme		Resolved	Fixed
	CONF-10305	Multimedia *.swf files are not able to render in Internet Explorer browser		Resolved	Fixed
	CONF-10205	Page permissions info on page info screen doesn't display permission creator			Duplicate

		correctly		Closed	
	CONF-10189	Reference to plugin administration should not be viewable by confluence admin, only by sysadmins		Resolved	Fixed
	CONF-10188	'Attachment storage page' link on EDIT general configuration page is broken		Resolved	Fixed
	CONF-10181	Global permissions 'Help Tips' disappear.		Resolved	Fixed
	CONF-10145	PDF generation does not preserve whitespaces defined in code and noformat macros		Resolved	Fixed
	CONF-10125	Link in Help Tips on General Configuration page goes to wrong doc space		Resolved	Fixed
	CONF-9959	Cannot upgrade from 2.2.10 and earlier to 2.6 without first having to upgrade to 2.5		Resolved	Fixed
	CONF-9723	remove comment notification wrongly claims that the comment author (or perhaps page author?) is the comment remover		Resolved	Fixed
	CONF-7919	Using the JIRAPORTLET macro does not render any HTML at all		Resolved	Fixed
	CONF-7021	Disabling trackback configuration option does not disable display of incoming pings		Resolved	Fixed

Confluence 2.7.2 Upgrade Guide

Confluence 2.7.2 is a recommended upgrade which fixes a security flaw and other bugs. You'll find details of the fixes in the release notes.

Upgrade Procedure

If you are already running a version of Confluence, please follow these instructions to upgrade to the latest version:

1. Before you upgrade, we strongly recommend that you **back up your `confluence.home` directory and database**.
2. If your version of Confluence is earlier than 2.7.0, read the release notes and upgrade guides for **all releases** between your version and the latest version. In particular:
   - Please read the 2.7 upgrade notes.
   - If you are upgrading from 2.1 or earlier, please read the 2.2 release notes.
3. Download the latest version of Confluence.
4. Follow the instructions in the Upgrade Guide.

Confluence 2.7.1 Release Notes

24 January 2008

Presented with pleasure by the Atlassian Confluence team: **Confluence 2.7.1** is a recommended upgrade which fixes a security flaw and other bugs, and brings a couple of improvements.

We have identified and fixed an XSS (cross-site scripting) security flaw which may affect Confluence instances in a public environment. For details, please refer to the security advisory.

Recording of authorship and history for page attachments is improved, so that attachment history is now retained after operations such as editing the attachment or moving it to a new page.

A new option on the Export Space screen allows administrators to export all pages to XML even when page-restrictions deny the administrator access to some of the pages.

This release also fixes problems in the following areas:

- The SOAP API, which was broken in Confluence 2.7.0.
- Logging.
- Internationalisation (support for different languages).
- Case-sensitivity for usernames and group names when using LDAP integration, and problems with upper-case letters in usernames when applying space permissions.

There's a complete list of fixes below. Click a specific issue to see details of the fix, and to download any patches where relevant.

Confluence 2.7.1 is available from the download centre.

Upgrading from a previous version of Confluence

Upgrading Confluence should be fairly straightforward. Please read the upgrade instructions. *We strongly recommend that you back up your `confluence.home` directory and database before upgrading!*

Updates and fixes in this release

JIRA Issues (27 issues)					
Type	**Key**	**Summary**	**Priority**	**Status**	**Resolution**
	CONF-13848	CLONE -Space export does not export restricted pages.		Resolved	Incomplete
	CONF-10436	Document update to attachment details		Closed	Fixed
	CONF-10434	Document update to export space screen		Closed	Fixed
	CONF-10392	Confluence javac should not fork		Closed	Fixed
	CONF-10359	http.proxyPort System property is ignored		Resolved	Fixed
	CONF-10311	Improve WebLinks in Confluence to allow for customizing the baseURL		Resolved	Fixed
	CONF-10302	Poor handling of null request in DashboardMacroSupport.getRequestParameter() during exports (such as HTML and PDF)		Resolved	Fixed
	CONF-10289	Security vulnerability with Dashboard spacesSelectedTab		Resolved	Fixed
	CONF-10269	ContentEntityObject.getAttachmentNamed() doesn't return the latest version of attachment		Resolved	Fixed
	CONF-10268	Hibernate SQL logging does not show up on confluence 2.6.2 and 2.6.1		Closed	Fixed
	CONF-10245	{note} {warning} {info} {tip} macros do not display a standard width when using Classic Theme		Closed	Fixed
	CONF-10235	Confluence 2.7 SOAP API is broken		Closed	Fixed
	CONF-10213	XML RPC server uses platform default character encoding to decode requests		Resolved	Fixed
	CONF-10201	"org.apache.commons.logging.impl.Jdk14Logger does not implement Log" on Websphere due to commons-logging-1.0.jar		Resolved	Fixed
	CONF-10140	XML-RPC does not handle Japanese characters in page title and content		Resolved	Fixed
	CONF-9923	Cannot switch to english language after upgrading to 2.6.0 from 2.5.4 with german language pack		Resolved	Fixed
	CONF-9867	Entering invalid space key in URL allows setting of invalid global permissions		Resolved	Fixed
	CONF-9609	Space key with "-" (hyphen) thows exception		Resolved	Fixed
	CONF-9488	The Recently Updated Content Macro fails to export to HTML		Closed	Fixed
	CONF-9469	Upper case letters in user names don't work with space permissions		Resolved	Fixed
	CONF-9265	Attachment history gets lost		Resolved	Fixed
	CONF-8863	Retain the Original Create Date on Attachments despite Edits to FileName, Comment Type, Content Type, or Page location		Resolved	Fixed
	CONF-8848	Editing an attachment to removes author and date information from previous versions		Resolved	Fixed

	CONF-8475	Missing localisation in attachments macro		Resolved	Fixed
	CONF-7342	Confluence uses user.language and user.country to choose default language		Resolved	Fixed
	CONF-7283	Space export does not export restricted pages.		Resolved	Fixed
	CONF-3504	A Space belonging to an old session is sometimes being cached.		Closed	Cannot Reproduce

Confluence 2.7.1 Upgrade Guide

Confluence 2.7.1 is a recommended upgrade which fixes a security flaw and other bugs, and brings a couple of improvements. You'll find details of the fixes in the release notes.

Upgrade Notes

As part of the fix for case-sensitivity in usernames and group names when using LDAP integration (CONF-9469), an upgrade task will consolidate permissions for the same user or groups where they differ only by case.

Depending on the size of your Confluence installation (number of spaces and the permissions applied to them) the upgrade task could result in a longer-than-usual delay when starting Confluence for the first time after the upgrade. Please be patient while this happens — it could take a few minutes.

During this process there will be regular progress reports in the log, something like this:

```
[atlassian.confluence.upgrade.ConsolidatePermissionsUpgradeTask] doUpgrade Consolidating
SpacePermissions for Space Monkeys in a Barrel (key=MONBAR)
```

Backup essential, because permissions will be modified
The upgrade task will modify permissions in your Confluence database. So that our usual 'strong recommendation' to back up the database becomes even stronger.

Upgrade Procedure

If you are already running a version of Confluence, please follow these instructions to upgrade to the latest version:

1. Before you upgrade, we strongly recommend that you **back up your `confluence.home` directory and database**.
2. If your version of Confluence is earlier than 2.7.0, read the release notes and upgrade guides for **all releases** between your version and the latest version. In particular:
   - Please read the 2.7 upgrade notes.
   - If you are upgrading from 2.1 or earlier, please read the 2.2 release notes.
3. Download the latest version of Confluence.
4. Follow the instructions in the Upgrade Guide.

Confluence 2.7 Release Notes

12 December 2007: With great pleasure, the Atlassian Confluence team presents Confluence 2.7.

Confluence 2.7 has improvements for administrators and end-users alike. Administrators can set up trusted communication between Confluence and JIRA. The result? The JIRA Issues and JIRA Portlet macros will show only the issues which the Confluence user is authorised to see. You no longer need to include a username and password in the markup code.

The two-tier administrator permissions allow system administrators to delegate some functions to team leaders or managers. Logging is simplified, and administrators can change logging levels at runtime. The improved user management framework speeds up your user searches.

Sorting of images is now possible in the Gallery macro. And when creating a page or news item, you can attach images or files immediately, without waiting until you have saved the page.

- Upgrading to Confluence 2.7 is free for all customers with active Confluence software maintenance as at 12 December 2007.
- Thank you for all your issues and votes. Keep on logging!
- We've highlighted our favourite bits of this new release below.
- And there's more.

Upgrading to Confluence 2.7

Upgrading Confluence should be fairly straightforward. ***We strongly recommend that you back up your `confluence.home` directory and database before upgrading!***
Please refer to the upgrade instructions. There you will find instructions on upgrading, and details of the following:

- If you are using any third-party plugins, please test them thoroughly before rolling 2.7 into production.
- If you are using the Resin application server, you will need to configure Resin to use an XML parser that is XSD-aware.
- When upgrading, you should consider turning off the new JIRA/Confluence trusted communication feature and/or warnings.
- Your users will be automatically migrated to AtlassianUser during the Confluence upgrade process.
- All users and groups with the old 'Administer Confluence' permission will be converted to the new 'System Administrator' permission.
- By default, all installations of Confluence will now write log messages to the Confluence home directory instead of the application server's log file.
- Java 1.4 will be deprecated after Confluence 2.8.
- Read more information.

Responding to your feedback:

4 new feature requests implemented

380 votes satisfied

DOWNLOAD latest version

Highlights of this release:

- JIRA Issues and Portlet macros use new trusted authentication
- Two-tier administrator permissions
- Inserting images and attaching files during page creation
- Sorting of images in Gallery macro
- Simplified and improved logging
- Performance, maintainability and administration
- Plus over 90 fixes and improvements

Highlights of Confluence 2.7

1

JIRA Issues and Portlet macros use new trusted authentication

- The JIRA Issues macro and the JIRA Portlet macro allow you to display a list of JIRA issues on your Confluence page.
- Prior to Confluence 2.7, you had to include a username and password in the markup code if you needed to display issues with restricted viewing. This release and JIRA 3.12 solve the problem.
- Read more about trusted communication.



Example markup – User will see the issues they are authorised to see (provided trusted communication is enabled):

```
1.{jiraissues:url=CONTENT|columns=type;key;summary}
```

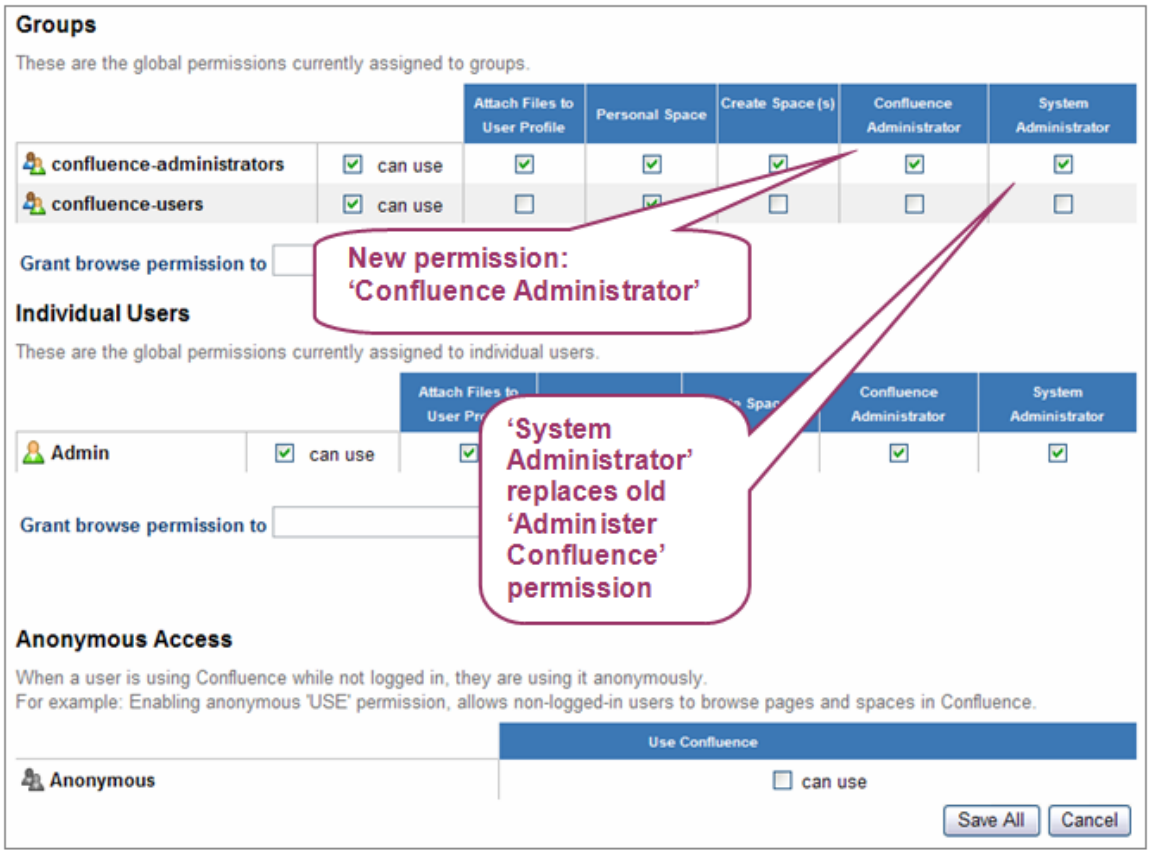
Example markup – User will see only the issues authorised for anonymous viewing:

```
1.{jiraissues:url=CONTENT|columns=type;key;summary|anonymous=true}
```



Two-tier administrator permissions

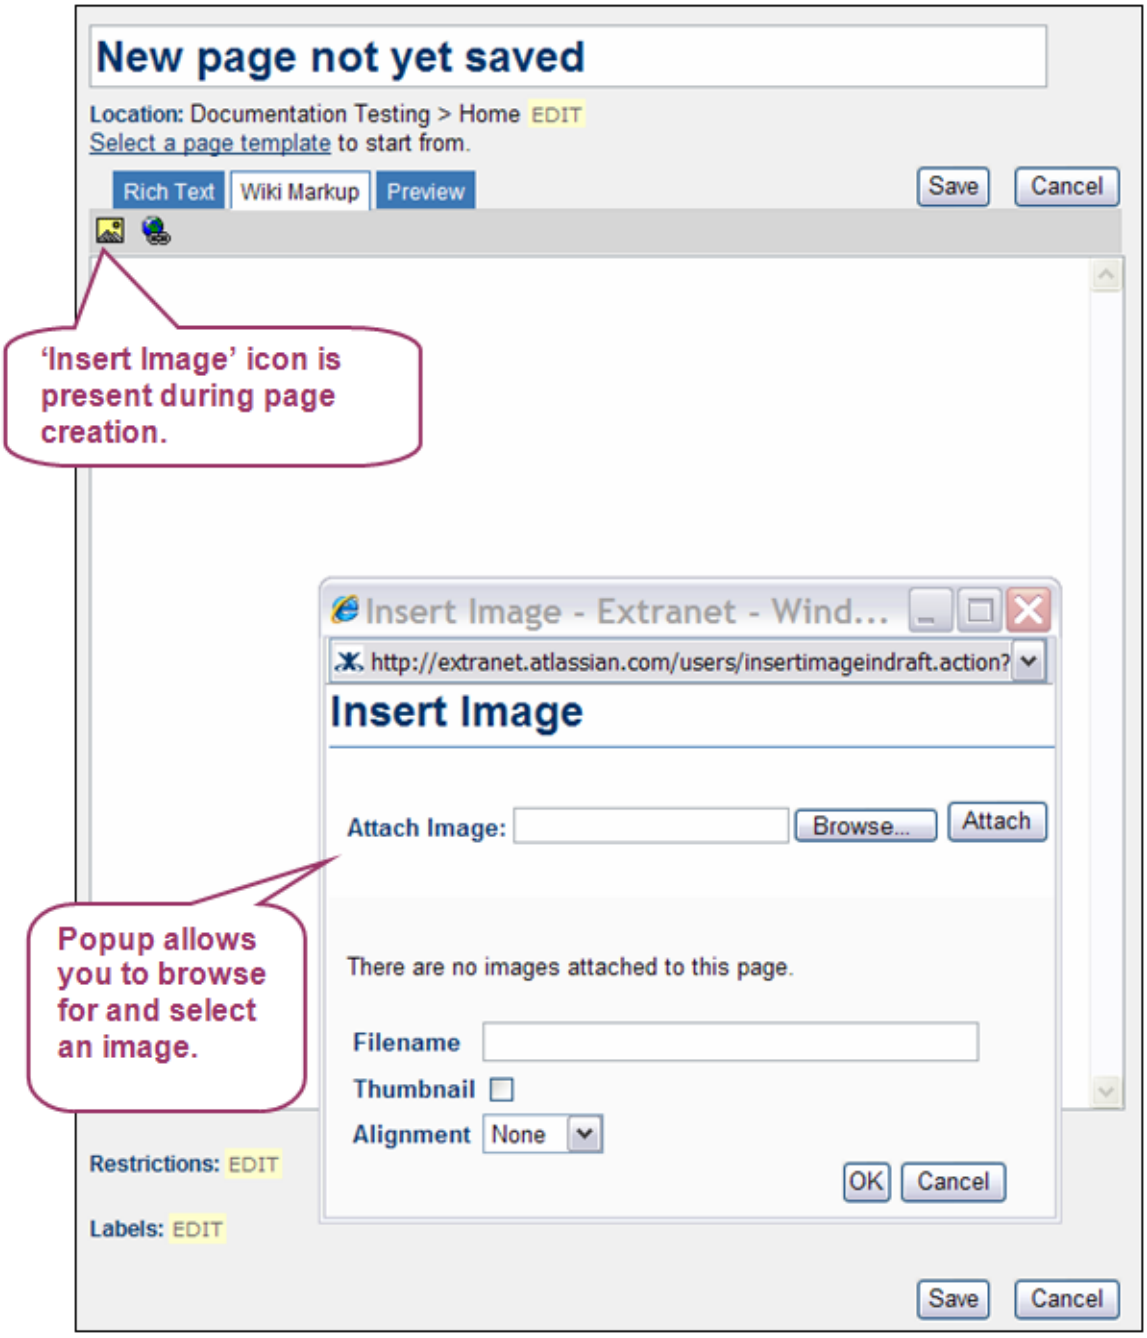
- The original 'Administer Confluence' permission is now called 'System Administrator'.
- A new permission level, called 'Confluence Administrator', is similar to 'System Administrator' but excludes the functions which may compromise the security of the Confluence system.
- You can delegate administrator privileges to project managers or team leaders while preserving the security of your Confluence site, by granting the managers the new 'Confluence Administrator' permission.
- Read the documentation.





Inserting images and attaching files during page creation

- You can now attach an image or other file during creation of new page – before you have saved the page.
- This applies to pages and news items.
- We have standardised the options for Wiki Markup mode and the Rich Text editor:
  - The 'Insert Image' popup allows you to select thumbnails and alignment.
  - The 'Insert Link' popup allows aliases and tooltips.
- Read more about inserting an image and attaching a file to a page.



4

Sorting of images in Gallery macro

- The new 'sort' argument allows you to order the images by file name, comment or date last modified.
- Read the documentation.

Example: Sorting the images by file name

```
1.{gallery:title=Some office photos, and a waterfall|sort=name}
```

Example: Sorting the images by date and showing the most-recently-modified first

```
1.{gallery:title=Some office photos, and a waterfall|sort=date|reverseSort}
```

5

Simplified and improved logging

- Confluence now writes its logs to the Confluence home directory. Both the Standalone and the EAR/WAR editions behave in the same way. For more information, see the logging documentation.
- We have rationalised the reporting to the different levels (ERROR, INFO, WARN, etc) and removed many unnecessary exceptions and stacktraces from the logs.
- You can change the logging levels while Confluence is running. Read more information.



Performance, maintainability and administration

- If you are currently using the standard configuration for user management, your users will be automatically migrated to the AtlassianUser framework on upgrade to Confluence 2.7. See the upgrade notes. This will result in a dramatic increase in the speed of user searches.
- Confluence now supports Java 6, allowing you to take advantage of its performance improvements.
- A further set of improvements to the stability of your Confluence system.



Plus over 90 fixes and improvements

- You'll no longer see those annoying browser messages when you click the browser's 'Back' button, for example after viewing search results.
- We've fixed some problems with the image and file attachment popups, and made them behave in the same way for both the Rich Text Editor and Wiki Markup.
- See the list of features, improvements and bug fixes.

The Confluence 2.7 team

Development
Agnes Ro
Anatoli Kazatchkov
Samuel Le Berrigaud
Andrew Lynch
Charles Miller
Christopher Owen
Dave Loeng
David Taylor
Dmitry Baranovskiy
Don Willis
Matt Ryall
Matthew Jensen
Paul Curren
Sam Le Berrigaud

Oversight & ~~Mis~~management
Adnan Chowdhury
Mike Cannon-Brookes
Per Fragemann
Scott Farquhar
Soren Harner
Tony Dagger

UI
Jason Taylor
Stephen Russell

Technical Writing
Sarah Maddox

Support
Choy Li Tham
Emily Stumpf
Gurleen Anand
Ivan Benko
James Fleming
Jeremy Largman
Kevin Tran
Mei Yan Chan
Ming Giet Chong

Partha Kamal
Roy Hartono
Tony Cheah Tong Nyee
Vincent Chang
Wayne Tombo

Confluence 2.7 Upgrade Guide

Confluence 2.7 has improvements for administrators and end-users alike. Administrators can set up trusted communication between Confluence and JIRA. The result? The JIRA Issues and JIRA Portlet macros will show only the issues which the Confluence user is authorised to see. You no longer need to include a username and password in the markup code.

The two-tier administrator permissions allow system administrators to delegate some functions to team leaders or managers. Logging is simplified, and administrators can change logging levels at runtime. The improved user management framework speeds up your user searches.

Sorting of images is now possible in the Gallery macro. And when creating a page or news item, you can attach images or files immediately, without waiting until you have saved the page.

Refer to the release notes for details of the features, improvements and bug fixes in this release.

On this page:

- Upgrade notes
  - Plugins
  - Configuring Resin
  - Trusted communication between JIRA and Confluence
  - WebDAV attachment manager deprecated
  - User migration
  - System Administrator and Confluence Administrator permissions
  - Location of Confluence logs
  - Java versions
- Upgrade procedure

Upgrade notes

Plugins

If you are using any third-party plugins, please test them thoroughly before rolling 2.7 into production.

Configuring Resin

This note applies if you are using the Resin application server. Due to an upgrade to one of the core components of Confluence (namely Spring), it will no longer work against an out-of-box configuration of Resin. To resolve this, you will need to configure Resin to use an XML parser that is XSD-aware. More information:

- Troubleshooting a Resin configuration
- Spring Framework documentation

Trusted communication between JIRA and Confluence

Confluence 2.7 provides a new trusted communication protocol, to allow secure authentication for the JIRA Issues and JIRA Portlet macros.

When upgrading, consider whether to **disable warning messages or turn off trusted communication altogether**. You may like to do this if you have a number of existing JIRA Issues or JIRA Portlet macros in your Confluence pages, and you do not intend to set up trusted communications soon. Read more about configuring trusted communications.

By default, trusted communication will be enabled when you upgrade to Confluence 2.7. This will affect your existing macros as follows:

- If the macro markup contains a username and password, the functionality is unchanged.
- If you have already set up trusted communication with a JIRA instance, the macro output will be as described in the guide on trusted communication.
- If the macro markup does not contain a username and password and you have not set up trusted communication with a JIRA instance, your Confluence pages will show a warning message above the macro output. See more information on troubleshooting.

WebDAV attachment manager deprecated

The option to store Confluence **attachments** on a WebDAV server has never worked in a useful fashion, and has not been maintained for many versions.

The WebDAV attachment manager is **deprecated** from Confluence 2.7, and will be removed from a future version of Confluence. If you store attachments on external WebDAV servers, we recommend that you migrate to file-system or database-backed attachment storage immediately. Refer to CONF-9313 and CONF-2887.

 This DOES NOT affect the operation of the WebDAV plugin.

User migration

Confluence 2.7 replaces OSUser with AtlassianUser as the underlying user management framework, greatly improving performance. Read more information about user management in Confluence.

When you upgrade from an earlier version of Confluence to release 2.7.0 or later, your users will be automatically migrated to AtlassianUser (but see the exceptions in the next paragraph). You may notice that your upgrade takes longer than usual, due to this migration process.

Automatic migration will **not** occur if any of the following is true:

- You have delegated user management to either JIRA or LDAP, or you have changed your user management from the standard configuration. If this is the case, you should upgrade as usual, ensuring that you retain your existing `atlassian-user.xml` or `osuser.xml` file.
- You have already migrated to AtlassianUser.

The progress of the migration will be shown in your log files. For example, a successful migration will show:

```
01.2007-10-08 21:33:07,979 INFO [main]
[atlassian.confluence.upgrade.OSUserToAtlassianUserMigrationUpgradeTask] userMigrationStarted
Starting user migration. 12288 users to migrate.
02.2007-10-08 21:33:09,784 INFO [main]
[atlassian.confluence.upgrade.OSUserToAtlassianUserMigrationUpgradeTask] userMigrated 100 users
migrated out of 12288.
03....
04.2007-10-08 21:36:18,304 INFO [main]
[atlassian.confluence.upgrade.OSUserToAtlassianUserMigrationUpgradeTask] userMigrated 12100 users
migrated out of 12288.
05.2007-10-08 21:36:20,112 INFO [main]
[atlassian.confluence.upgrade.OSUserToAtlassianUserMigrationUpgradeTask] userMigrated 12200 users
migrated out of 12288.
06.2007-10-08 21:36:21,562 INFO [main]
[atlassian.confluence.upgrade.OSUserToAtlassianUserMigrationUpgradeTask] userMigrationComplete
User migration complete.
07.2007-10-08 21:36:35,957 INFO [main]
[atlassian.confluence.upgrade.OSUserToAtlassianUserMigrationUpgradeTask] groupMigrationStarted
Group migration started. 44 groups to migrate.
08.2007-10-08 21:36:35,960 INFO [main]
[atlassian.confluence.upgrade.OSUserToAtlassianUserMigrationUpgradeTask] groupMigrated 10 groups
migrated out of 44.
09...
10.2007-10-08 21:36:38,055 INFO [main]
[atlassian.confluence.upgrade.OSUserToAtlassianUserMigrationUpgradeTask] groupMigrated 40 groups
migrated out of 44.
11.2007-10-08 21:36:38,208 INFO [main]
[atlassian.confluence.upgrade.OSUserToAtlassianUserMigrationUpgradeTask] groupMigrationComplete
Group migration complete.
```

System Administrator and Confluence Administrator permissions

Confluence 2.7 brings two administrator-level permissions in place of one. The new permissions are '**System Administrator**' and '**Confluence Administrator**'. Refer to the documentation for details.

When you upgrade to version 2.7, all users and groups with the old 'Administer Confluence' permission will be converted to the new 'System Administrator' permission. The powers of the 'confluence-administrators' group remain unchanged.

Location of Confluence logs

Confluence's default logging behaviour has changed with Confluence 2.7. Both Standalone and EAR/WAR distributions follow the same default behaviour:

- When you start Confluence, log entries will be sent to the application server logs until Confluence has completed its initial bootstrap. Any log entries will be repeated into the `<confluence-home>` log described below.
- Once the initial startup sequence is complete, all logging will be to `<confluence-home>/logs/atlassian-confluence.log`. For example: `c:/confluence/data/logs/atlassian-confluence.log`.

  Note that the default location is now the Confluence **home directory** instead of the **application server's log file**.

Java versions

Confluence 2.7 supports Java 1.4, Java 5 and Java 6. We recommend Java 6 because of its increased performance and easier troubleshooting, due to enhanced memory dump and profiling capabilities.

Advance notice: Java 1.4 will be deprecated in a future release. Confluence 2.8 will be the last version that supports Java 1.4.

Upgrade procedure

If you are already running a version of Confluence, please follow these instructions to upgrade to the latest version:

1. Before you upgrade, we strongly recommend that you **back up your `confluence.home` directory and database**.
2. If your version of Confluence is earlier than 2.6.x, read the release notes and upgrade guides for all releases between your version and the latest version. In particular:
   - If you are upgrading from 2.1 or earlier, please read the 2.2 release notes.
   - If you are upgrading from a version earlier than 2.6.0, please read the 2.6 upgrade notes.
3. Download the latest version of Confluence.
4. Follow the instructions in the Upgrade Guide.

Issues Resolved in Confluence 2.7

Below is the full list of issues resolved by Confluence 2.7. You can read the release notes here.

Errors were reported by the JIRA trusted connection.

- APP_UNKNOWN; Unknown Application: {0}; ["confluence:4557196"]

JIRA Issues (118 issues)				
Type	**Key**	**Summary**	**Status**	**Resolution**
	CONF-1339	WebDAV support pass-thru permissions	Resolved	Won't Fix
	CONF-1545	Add New User form too small (too narrow)	Resolved	Fixed
	CONF-2127	A page containing the rss-macro is not displayed if the requested rss-feed is "down"	Resolved	Fixed
	CONF-2219	Enable adding attachments while creating new page or news	Resolved	Fixed
	CONF-2285	AbstractHttpRetrievalMacro is too noisy in logs	Resolved	Fixed
	CONF-2711	Better authentication method for jiraissues and jiraportlet macros	Resolved	Fixed
	CONF-2866	Enable administrator to change loglevels at runtime	Resolved	Fixed
	CONF-3232	Insert-image-popup makes cursor position jump to top of text area	Resolved	Fixed
	CONF-4100	Return to search for Confluence	Closed	Fixed
	CONF-4738	Make search redirect to GET, so that back button works	Closed	Fixed
	CONF-4778	Allow {gallery} macro to sort by filename or by date	Resolved	Fixed
	CONF-5218	Permissions adminstration permits adding a user as a group	Closed	Fixed
	CONF-5405	Don't allow creation of a user with the same name as an existing group	Closed	Fixed
	CONF-5457	Indexing errors and imprecise logging ("Looks like an invalid size") with certain Powerpoint (ppt) attachments	Resolved	Fixed
	CONF-5749	Migration tool used to move from OSUSer to AtlassianUser should distinguish external users	Resolved	Fixed
	CONF-6582	Write logs to $confluence_home/logs	Resolved	Fixed
	CONF-6597	Image-attachment popup should not ask user for filename	Resolved	Fixed
	CONF-6766	The Recently-used-labels-Macro displays duplicate labels	Resolved	Fixed

	CONF-7048	Profile user searching and address search bottlenecks	Closed	Fixed
	CONF-7303	Searching for users by name should not retrieve every user	Resolved	Fixed
	CONF-7457	Remove meta cache control elements from Confluence HTML	Resolved	Fixed
	CONF-7481	Officially support Java 1.6	Resolved	Fixed
	CONF-7587	Lists don't have same amount of whitespace beneath headings (compared to normal text beneath headings)	Resolved	Fixed
	CONF-7599	Automatically migrate Confluence installations to using Hibernate storage for atlassian-user users	Resolved	Fixed
	CONF-7600	Servlet plugin does not unload classes cleanly	Closed	Fixed
	CONF-7975	Migrate c.a.c to hibernate user repository	Closed	Fixed
	CONF-8020	Missing context on indexing exceptions / errors	Resolved	Fixed
	CONF-8280	Export code fails when executed in non-servlet context	Resolved	Fixed
	CONF-8311	User picker: Javascript error occurs for users whose names contain double-quotes	Resolved	Fixed
	CONF-8745	Javascript (and logfiles) error in Rich-Text-Editor for anonymous users: "couldn't find resource '/plugins/tinymce/css/table.css"	Resolved	Fixed
	CONF-8879	User permissions are cached beyond user removal, so a deleted-and-recreated user may temporarily have wrong permissions	Resolved	Fixed
	CONF-8887	Incorporate warning message into footer when confluence installation is using HSQL	Resolved	Fixed
	CONF-8902	Attachment-operation separator ("pipe symbol") is displayed even when some operations are not visible	Resolved	Fixed
	CONF-8910	ArrayIndexOutOfBoundsException in MsPowerpointContentExtractor	Closed	Fixed
	CONF-8922	WYSIWYG editor: language values for german are missing	Closed	Fixed
	CONF-9074	Mail links broken in PDF export	Resolved	Fixed
	CONF-9079	webdav as a attachment store does not work	Resolved	Fixed
	CONF-9117	Confluence API supports adding user with null password, but users will null passwords produce NullPointerException when using the osuser to atlassian-user migration utility jsp	Resolved	Fixed
	CONF-9129	Introduce 'Confluence Administrator' permission (less powerful than 'System Administrator' permission)	Resolved	Fixed
	CONF-9178	LDAP groups sometimes not shown for a user when Confluence is clustered	Resolved	Fixed
	CONF-9195	Confluence 2.5.6 ldap configuration failing on osuser2atluser.jsp migration	Resolved	Fixed
	CONF-9289	Resources served from /display/* are not sent with correct cache headers	Resolved	Fixed
	CONF-9290	Improve browser-caching and back-navigation by removing the "no-store" cache control headers	Resolved	Fixed
	CONF-9337	Include tImestamp on the "Recently Updated" section on Dashboard	Resolved	Fixed

	CONF-9367	Log all available data displayed on the 500 page to the logfile as well, and provide a unique identifier on page for easy bug retrieval from logs	Resolved	Fixed
	CONF-9373	XWork plugin load failure can cause subsequent XWork plugins to be unloadable	Resolved	Fixed
	CONF-9379	Organise spring config files more sanely	Resolved	Fixed
	CONF-9394	Option to disable "secure" cookie when using HTTPS just for login page	Resolved	Fixed
	CONF-9409	Image-attachment popup design causes image attachments to fail	Resolved	Fixed
	CONF-9432	Document (better) where the log files can be located	Closed	Fixed
	CONF-9434	Enabling Caching for Hibernate Repository causes net.sf.hibernate.LazyInitializationException error	Resolved	Fixed
	CONF-9442	Make "insert-image-popup" and "link-popups" consistent across rich text and wiki editors	Resolved	Fixed
	CONF-9444	Clean up default link bodies for "raw" links.	Resolved	Fixed
	CONF-9445	Logfiles: Incomplete and missing logging of exceptions	Resolved	Fixed
	CONF-9447	Log important information to INFO (instead of to DEBUG, or not at all)	Resolved	Fixed
	CONF-9466	Hibernate exception when removing space, due to cascade relationship	Closed	Fixed
	CONF-9468	Deprecate com.atlassian.confluence.renderer.WikiStyleRenderer and its default Confluence implementation	Closed	Fixed
	CONF-9473	Character-Encoding ISO-8859-1 does not work in jiraissues macro	Resolved	Fixed
	CONF-9520	Update documentation for logging to home directory feature	Closed	Fixed
	CONF-9530	Merge changes from tiny-mce stable branch back to main trunk before releasing trunk changes into confluence	Closed	Fixed
	CONF-9531	Document new parameters on gallery macro	Closed	Fixed
	CONF-9566	Breadcrumbs for attachment page displays error	Resolved	Fixed
	CONF-9583	After deleting a user (and using OSUser repository), some of the user details still remain in the database	Resolved	Fixed
	CONF-9591	Confluence doesn't start when only os user repository is configured in atlassian-user.xml	Closed	Fixed
	CONF-9610	Remove empty *Action.properties files	Resolved	Fixed
	CONF-9623	userAccessor.removeGroup(...) fails in Confluence 2.6.0	Resolved	Fixed
	CONF-9626	Logging of job-progress should be done with an aspect or wrapper	Resolved	Fixed
	CONF-9658	"Link Properties-Popup" looks broken in classic-based themes (left nav, clickr, classic)	Resolved	Fixed
	CONF-9700	Plugin repository not available on websphere 5.1.10	Resolved	Fixed
	CONF-9717	Exceptions thrown while retrieving internationalised text halt rendering of entire page	Resolved	Fixed

	CONF-9734	Thumbnails are not regenerated if a new attachment of the same name is uploaded	Resolved	Fixed
	CONF-9738	Move to new pom architecture	Resolved	Fixed
	CONF-9743	RPC plugins cannot be loaded dynamically	Closed	Fixed
	CONF-9778	Changes to user-preferences sometimes don't propagate across a cluster to the other nodes	Resolved	Fixed
	CONF-9794	Images attached to pages (with very large page ids) don't render in preview mode	Resolved	Fixed
	CONF-9811	User migration from an OSUser to a Crowd repository fails.	Resolved	Fixed
	CONF-9836	Remove WebDAV attachment manager documentation	Resolved	Fixed
	CONF-9843	The Warning-, Info-, Tip- and Note-Macros do not expand to full column/screen width	Resolved	Fixed
	CONF-9852	Upgrade to Atlassian-Extras 1.10	Resolved	Fixed
	CONF-9876	The URLs in the footer are rendered with a missing /	Resolved	Fixed
	CONF-9894	Group picker: Javascript error occurs for groups whose names contain double-quotes	Resolved	Fixed
	CONF-9896	Rework the Space Export functionality - broken by 2.6 changes.	Resolved	Fixed
	CONF-9899	Confluence throws ClassCastException when you are trying to edit a blogpost	Resolved	Fixed
	CONF-9908	Border underneath the main page tabs is not rendered on IE	Resolved	Fixed
	CONF-9914	Breadcrumb on "Running Tasks" doesn't list the Dashboard	Resolved	Fixed
	CONF-9938	Prevent calls to external websites from blocking Confluence	Resolved	Fixed
	CONF-9943	Filter and thread-local for "request timestamp"	Resolved	Fixed
	CONF-9944	Functional test RPC plugin should provide getText(key) method	Resolved	Fixed
	CONF-9947	TestBean prevents execution of units tests for confluence module when run individually	Resolved	Fixed
	CONF-9948	Incorrect comments link on blog post (in clickr theme)	Resolved	Fixed
	CONF-9950	Enable standards-compliant HTML-rendering by changing DOCTYPE from HTML 4.0 to HTML 4.01	Resolved	Fixed
	CONF-9951	Integrate Seraph's trusted applications infrastructure	Resolved	Fixed
	CONF-9957	Document the new API to export resource content for plugin developers	Resolved	Fixed
	CONF-9964	Textual changes on Attachment Storage admin page to deprecate WebDAV	Resolved	Fixed
	CONF-9968	Spelling mistake in line 79 of Page-hierarchy.pdfexport.vm file	Resolved	Fixed
	CONF-9984	Superfluous exception logging from createrssfeed.action (java.io.IOException: WEB8001: Write failed)	Resolved	Fixed

	CONF-9989	Using the attachments macro in a comment is broken and disables access to the whole page	Resolved	Fixed
	CONF-9994	Built-in profile picture breaks on upgrade due to change in file format from .png to .gif	Resolved	Fixed
	CONF-9995	Shared Mode setting gets overwritten by general configuration	Resolved	Fixed
	CONF-10003	Unit tests create a SpringContext for every test case	Resolved	Fixed
	CONF-10028	RSVP macro problem when page title contains unsafe url characters	Resolved	Fixed
	CONF-10032	Database driver version not displayed when using datasource	Resolved	Fixed
	CONF-10041	Add/watch icon layout broken in Safari 3.0	Resolved	Fixed
	CONF-10056	i18n error in email subject	Resolved	Fixed
	CONF-10094	Password reset e-mail does not give new password; instead gives $action.newPassword	Resolved	Fixed
	CONF-10111	Missing i18n strings	Resolved	Fixed
	CONF-10120	Users listed in a group (in admin console) are not sorted alphabetically with Hibernate-based user management	Resolved	Fixed
	CONF-10122	Insert link popup does not have its alias field populated with the text the user has highlighted	Resolved	Fixed
	CONF-10143	ClassCastException creating RSS feed because a Mail instance being cast to an AbstractPage	Resolved	Fixed
	CONF-10164	XSS vulnerability in recently updated and configure RSS feed actions	Resolved	Fixed
	CONF-10173	Supply new indexes.ddl file for Upgrade Guide	Resolved	Fixed
	CONF-10175	Information breadcrumb on certain page types is a broken link	Resolved	Fixed
	CONF-10185	General configuration displays $timeouts.heading instead of Connection Timeouts	Resolved	Fixed
	CONF-10211	Missing internationalization on global layouts configuration page	Resolved	Fixed
	CONF-11210	Improve speed of "Atlassian Plugin Repository"	Resolved	Fixed
	CONF-11832	Http Timeout Documentation is Out Of Date	Closed	Fixed
	CONF-12177	Duplicate user accounts allowed in database because no unique constraint configured on USER table	Closed	Fixed
	CONF-12509	Check array size exceeding VM limit during indexing	Closed	Fixed

Confluence 2.6.3 Release Notes

Confluence 2.6.3 is the **standard edition** version of Confluence 2.6.x. This version of Confluence is equivalent to Confluence version 2.6.2, but it does not include Oracle's Coherence technology. Due to licensing changes with Oracle over the Coherence clustering and distributed caching technology, Atlassian is no longer able to distribute this technology to customers without a Confluence clustered license.

Therefore, Atlassian has released 'standard editions' for each previous major version of Confluence back to 2.6, which includes Confluence versions 2.10.4, 2.9.3, 2.8.3, 2.7.4 and 2.6.3. Confluence 2.6.3 will be the only Confluence 2.6.x version available to customers with a non-clustered Confluence license.

For its caching functionality, Confluence 2.6.3 utilises Ehcache. This functionality was provided by the Coherence technology in previous Confluence 2.6.x versions.

For more information about the features, updates and fixes in Confluence 2.6.2, please refer to the Confluence 2.6.2 Release Notes.

Upgrading from a Previous Version of Confluence

Upgrading Confluence should be fairly straightforward. Please read the Confluence 2.6.2 Upgrade Notes. However, if you have customised the cache settings in your installation of Confluence (e.g. for performance reasons), then please read the Confluence 3.0.1 Upgrade Notes for important information about transferring your caching layer customisations from Coherence to Ehcache.

We *strongly recommend* that you back up your `confluence.home` directory and database before upgrading.

Confluence 2.6.2 Release Notes

27 November 2007

Atlassian is proud to announce the release of Confluence 2.6.2. This is a **highly recommended** upgrade, because it fixes some security flaws which may affect Confluence instances in a public environment. These flaws are XSS (cross-site scripting) vulnerabilities in some of Confluence's macros and Wiki Markup, which potentially allowed a user to insert malicious HTML tags or script into a Confluence page. Please refer to the Security Advisory for details.

This point release also includes more than 20 other fixes and improvements.

As part of our drive to tighten up the security in Confluence, we have removed support for the 'style' attribute in the Wiki Markup for images. This was an undocumented feature, which is now no longer available. To help those who may have used the 'style' tag to add coloured borders, we have added a new 'bordercolor' attribute to the image markup.

The PDF and HTML space exports are now more reliable than in Confluence 2.6.0 and 2.6.1. We've fixed the failure to send daily digest email notifications. (This problem occurred when the Confluence instance contained draft pages.) This release also contains some improvements in the wiki's support of internationalisation. And you'll be delighted to see that the plus and minus buttons are back, next to the 'Recently Updated' section of the Dashboard – so you can now increase or decrease the number of items you see in that section.

There's a complete list of fixes below. You can download Confluence 2.6.2 from the download centre.

Upgrading from a previous version of Confluence

Upgrading Confluence should be fairly straightforward. Please read the upgrade instructions. *We strongly recommend that you back up your `confluence.home` directory and database before upgrading!*

Updates and fixes in this release

JIRA Issues (26 issues)

Type	Key	Summary	Priority	Status	Resolution
	CONF-10042	log4j references ConfluenceHomeLogAppender class that does not exist in 2.6.1		Resolved	Fixed
	CONF-10008	Document bordercolor attribute on wiki markup for images		Resolved	Fixed
	CONF-9975	error saving new or existing documents when using the Rich Text Editor		Resolved	Fixed
	CONF-9890	Typo in the Wiki Notation guide - {{monspaced}}		Resolved	Fixed
	CONF-9877	CustomCJKAnalyzer does not work with JDK 1.6		Resolved	Fixed
	CONF-9873	Concurrent Modification Exception caused when accessing a space in clustered version of Confluence		Resolved	Fixed
	CONF-9866	Replace System.out, System.err and printStackTrace references with logging		Resolved	Fixed
	CONF-9807	On the Administrators page e-mail addresses are shown "mailto:user@domain.com" instead of "user@domain.com"		Resolved	Fixed
	CONF-9804	Move the XFire dependency for Crowd from 1.2.1 to 1.2.6		Resolved	Fixed
	CONF-9803	Open source Confluence system and bundled plugins		Closed	Fixed
	CONF-9781	Duplicate Webwork JAR in Confluence 2.6-stable		Resolved	Fixed

	CONF-9773	Image thumbnail links on dashboard recently updated don't use context path in URL		Resolved	Fixed
	CONF-9771	NullPointerException when exporting space on Websphere and JBoss		Resolved	Fixed
	CONF-9770	Recent Changes on dashboard doesn't display plus/minus buttons to resize list		Resolved	Fixed
	CONF-9701	Daily report job not being generated when using drafts		Resolved	Fixed
	CONF-9659	PDF space export failing in Resin 3.x due to incorrect handling of HttpServletRequest		Resolved	Fixed
	CONF-9350	Numerous XSS Type 2 vulnerabilities in macros bundled with Confluence		Resolved	Fixed
	CONF-9299	Tiny Link can generate URLs ending with punctuations (which outlook doesn't like)		Resolved	Fixed
	CONF-9258	Incorrect search results for single and double-byte Japanese strings		Resolved	Fixed
	CONF-9238	Annonymous group loses view permission		Resolved	Fixed
	CONF-8495	"Find Groups" Popup not working in IE		Resolved	Fixed
	CONF-7750	extractBundledPlugins Found atlassian-bundled-plugins.zip, but failed to read file		Resolved	Fixed
	CONF-7680	Non-internationalised UI text in comments and space rss links		Resolved	Fixed
	CONF-7601	Images produced by macro plugins like Gliffy are not generated into .doc or .html exports		Closed	Fixed
	CONF-7362	captcha.response.empty key given when user cancels		Resolved	Fixed
	CONF-3427	PDF export doesn't handle {color} tag properly		Resolved	Fixed

Confluence 2.6.2 Upgrade Guide

27 November 2007

Atlassian is proud to announce the release of Confluence 2.6.2. This is a **highly recommended** upgrade, because it fixes some security flaws which may affect Confluence instances in a public environment. These flaws are XSS (cross-site scripting) vulnerabilities in some of Confluence's macros and Wiki Markup, which potentially allowed a user to insert malicious HTML tags or script into a Confluence page. Please refer to the Security Advisory for details.

This point release also includes more than 20 other fixes and improvements.

As part of our drive to tighten up the security in Confluence, we have removed support for the 'style' attribute in the Wiki Markup for images. This was an undocumented feature, which is now no longer available. To help those who may have used the 'style' tag to add coloured borders, we have added a new 'bordercolor' attribute to the image markup.

The PDF and HTML space exports are now more reliable than in Confluence 2.6.0 and 2.6.1. We've fixed the failure to send daily digest email notifications. (This problem occurred when the Confluence instance contained draft pages.) This release also contains some improvements in the wiki's support of internationalisation. And you'll be delighted to see that the plus and minus buttons are back, next to the 'Recently Updated' section of the Dashboard – so you can now increase or decrease the number of items you see in that section.

Upgrade Procedure

If you are already running a version of Confluence, please follow these instructions to upgrade to the latest version:

1. Before you upgrade, we strongly recommend that you **back up your `confluence.home` directory and database**.
2. If your version of Confluence is earlier than 2.6.1, read the release notes and upgrade guides for all releases between your version and the latest version. In particular:
   - If you are upgrading from 2.1 or earlier, please read the 2.2 release notes.
   - If you are upgrading from a version earlier than 2.6.0, please read the 2.6 upgrade notes.
3. Download the latest version of Confluence.

4. Follow the instructions in the Upgrade Guide.

Confluence 2.6.1 Release Notes

9 November 2007

Atlassian is proud to announce the release of Confluence 2.6.1. This is a **highly recommended** upgrade from 2.6, because it fixes some security flaws which may affect Confluence instances in a public environment. We will release more details of the security fixes within a few days, giving our customers a chance to upgrade before the details are made public. This point release also includes more than 40 other fixes and improvements.

We're especially glad that anonymous comments will no longer show the profile picture of the previous commenter, and newly-created spaces no longer seem to have been created by an anonymous user. The {excerpt} and {excerpt-include} macros now behave better. And we've fixed some issues with internationalisation.

There's a complete list of fixes below.

You can download Confluence 2.6.1 from the download centre.

Security Advisory 19 November 2007
We have now published the details of the security vulnerabilities, which are fixed in Confluence 2.6.1. Please refer to the security advisory for more information.

Upgrading from a previous version of Confluence

Upgrading Confluence should be fairly straightforward. Please read the upgrade instructions. *We strongly recommend that you back up your `confluence.home` directory and database before upgrading!*

Updates and fixes in this release

Errors were reported by the JIRA trusted connection.

- APP_UNKNOWN; Unknown Application: {0}; ["confluence:4557196"]

JIRA Issues (50 issues)

Type	Key	Summary	Priority	Status	Resolution
	CONF-9960	Missing files from the Confluence 2.6.1 standalone zip file		Resolved	Fixed
	CONF-9949	Confluence 2.6.1 bundled Tomcat missing the java 1.4 compatibility jar		Resolved	Fixed
	CONF-9824	Macros returning null from getBodyRenderMode() cause page edits to fail		Closed	Fixed
	CONF-9799	atlassian-confluence.log not generated in Confluence 2.6 standalone		Resolved	Fixed
	CONF-9760	Confluence 2.6.0 needs dependency cleanup		Closed	Fixed
	CONF-9732	User migration JSP script for LDAP intergation is not compatible with JDK 1.4		Resolved	Fixed
	CONF-9718	DWR debug mode is enabled		Resolved	Fixed
	CONF-9714	!! wiki notation doesn't seem to work for attached bmps in 2.6		Resolved	Fixed
	CONF-9704	Security Issue: XSS in wiki exception error page		Resolved	Fixed
	CONF-9700	Plugin repository not available on websphere 5.1.10		Resolved	Fixed
	CONF-9686	Mail is not send by confluence for page updates		Closed	Fixed
	CONF-9671	Duplicate javamail / mail JARs in different version		Resolved	Fixed
	CONF-9667	Anonymous comments display user image of previous non-anonymous commenter		Resolved	Fixed

	CONF-9662	Migrating attachments from file system to database presents a javascript error: response has no properties		Closed	Fixed
	CONF-9622	Newly created spaces appear to have been created by "Anonymous"		Resolved	Fixed
	CONF-9590	ConfluenceEntityObject.log is not static?		Resolved	Fixed
	CONF-9568	NullPointerException occuring in DefaultReferralManager		Resolved	Fixed
	CONF-9565	NullPointerException occuring in StylesheetAction		Resolved	Fixed
	CONF-9560	Cross-site scripting vulnerability in 500page.jsp		Resolved	Fixed
	CONF-9536	Selecting "site default" as preferred language results in alternating behavior		Resolved	Fixed
	CONF-9502	Explain inclusion/exclusion of wiki markup (formatting, images, links) in excerpt and excerpt-include macros		Closed	Fixed
	CONF-9497	Notifications for pages whose titles contain special characters have broken "Add Comment" links.		Resolved	Fixed
	CONF-9489	The Excerpt Macro no longer works in comments		Resolved	Fixed
	CONF-9477	Userlister macro ":groups=" option should not display all users when incorrectly specified or no group exists, and should accept ":group=" parameter		Resolved	Fixed
	CONF-9456	XSS Bug in printable link display		Resolved	Fixed
	CONF-9454	Pages created with restrictions show as created by Anonymous		Resolved	Fixed
	CONF-9436	NPE adding page with metadata macro on CAC		Resolved	Fixed
	CONF-9431	Bandana cache not flushed when site backup restored		Resolved	Fixed
	CONF-9392	StackOverflowError in ConfluenceLinkResolver.extractLinkTextList		Resolved	Fixed
	CONF-9330	Display the Server ID string in the Setup Wizard		Resolved	Fixed
	CONF-9209	XSS in image wiki markup		Resolved	Fixed
	CONF-9164	Clean up poorly formatted (bad) properties in ConfluenceActionSupport.properties		Closed	Fixed
	CONF-9151	Excerpt macro now rendered with HTML paragraph tag		Resolved	Fixed
	CONF-9099	the 'Insert Table' popup is full-screen wide		Resolved	Fixed
	CONF-9076	Underlining mnemonic missing from "Edit" and "View" tabs on pages using default theme		Resolved	Fixed
	CONF-8983	Excerpt-include macro does not display images.		Resolved	Fixed
	CONF-8866	Shortcut link values are double URL-escaped		Resolved	Fixed
	CONF-8826	Recently updated doesn't handle anonymous changes correctly		Resolved	Fixed
	CONF-8701	if an {excerpt} includes links, the {excerpt-include} macro does not render them		Resolved	Fixed

	CONF-8673	Missing i18n keys for dashboard's recently updated message		Resolved	Fixed
	CONF-8672	Missing i18 keys for "Choose Owner"		Resolved	Fixed
	CONF-8671	Missing i18 keys for Add label		Resolved	Fixed
	CONF-8669	i18n keys for Confluence Daily Report HTML mail missing		Resolved	Fixed
	CONF-8534	Using excerpt-include doesn't render the body in the {excerpt} macro when it contains {note} or {code} or {noformat}		Resolved	Fixed
	CONF-8004	Missing I18N-key for Index Update Task Items		Resolved	Fixed
	CONF-7378	"Remember me" does not work with sun application servers		Resolved	Fixed
	CONF-7127	Can't edit page when watchers include an LDAP user without an email address		Resolved	Fixed
	CONF-6586	Attachment links broken in text from include-excerpt macro		Resolved	Fixed
	CONF-5088	Functional tests should test typical Atlassian User configuration		Closed	Fixed
	CONF-4753	Excerpt macro does not work when placed in a heading.		Resolved	Fixed

Confluence 2.6.1 Upgrade Guide

9 November 2007

Atlassian is proud to announce the release of Confluence 2.6.1. This is a **highly recommended** upgrade from 2.6, because it fixes some security flaws which may affect Confluence instances in a public environment. We will release more details of the security fixes within a few days, giving our customers a chance to upgrade before the details are made public. This point release also includes more than 40 other fixes and improvements.

We're especially glad that anonymous comments will no longer show the profile picture of the previous commenter, and newly-created spaces no longer seem to have been created by an anonymous user. The {excerpt} and {excerpt-include} macros now behave better. And we've fixed some issues with internationalisation.

Upgrade Procedure

If you are already running a version of Confluence, please follow these instructions to upgrade to the latest version:

1. Before you upgrade, we strongly recommend that you **back up your `confluence.home` directory and database**.
2. If your version of Confluence is earlier than 2.6.0, read the release notes and upgrade guides for all releases between your version and the latest version. In particular:
   - If you are upgrading from 2.1 or earlier, please read the 2.2 release notes.
   - If you are upgrading from a version earlier than 2.6.0, please read the 2.6 upgrade notes.
3. Download the latest version of Confluence.
4. Follow the instructions in the Upgrade Guide.

Confluence 2.6 Release Notes

The Atlassian Confluence team is delighted to present Confluence 2.6.

Confluence 2.6 brings many popular features which save you time and improve the your wiki's usability. Upgrading to Confluence 2.6 is free for all customers with active Confluence software maintenance as at 27th September 2007.

A new theme brings a fresh, clean look and feel. This is for our customers who have asked for a friendlier interface and improved readability. We've included more social networking features, to enhance the sense of community in your wiki. For example, author photos are now shown in comments and in the 'Recent Updates' on the Dashboard. And the Social Bookmarking plugin is now shipped with Confluence, allowing you to share bookmarks with your team.

Other popular new features include default content for spaces, labels on templates, the ability to backdate or rename news items (blog posts), official MySQL 5.0 support and PDF export of images.

- Many thanks for your issues and votes. They help us keep improving our products.
- We've highlighted the main features of this release below.
- Attached is a full list of issues resolved in 2.6.

Upgrading to Confluence 2.6
• Upgrading Confluence should be fairly straightforward. Please refer to the upgrade instructions and notes . ***We strongly recommend that you back up your `confluence.home` directory and database before upgrading!*** • All draft pages will be destroyed during the upgrade process. Confluence administrators should warn users of the Confluence site that drafts will not survive the upgrade. • If you are using any **third-party plugins**, please test them thoroughly before rolling 2.6 into production.

Responding to your feedback:

5 new feature requests implemented

275 votes satisfied

Highlights of this release:

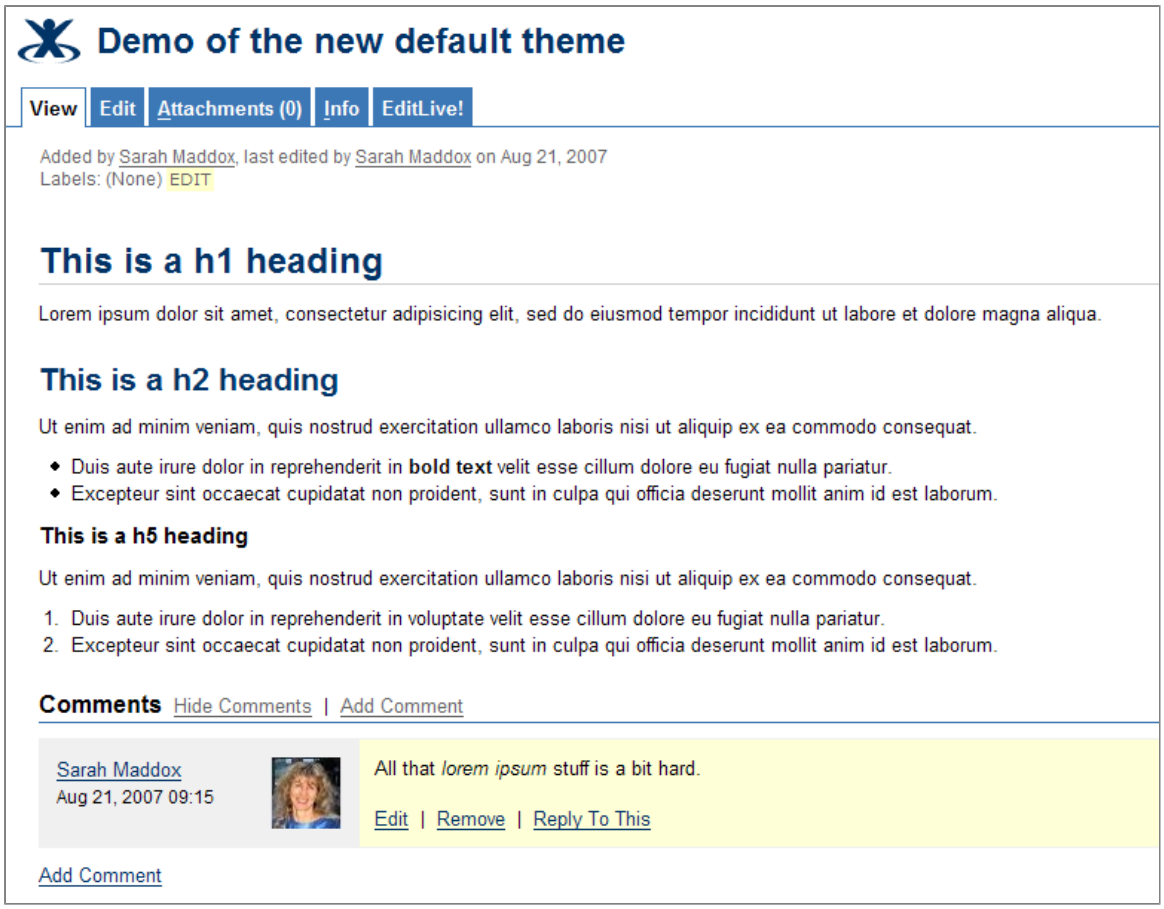
- Fresh look for the Default theme
- Personalised comments and Dashboard
- Space description on Dashboard
- Labels on templates
- Default content for space home pages
- Social Bookmarking plugin now bundled with Confluence
- Back-dating and renaming news items
- Plus over 90 other fixes and improvements
- Special thanks

Highlights of Confluence 2.6



Fresh look for the Default theme

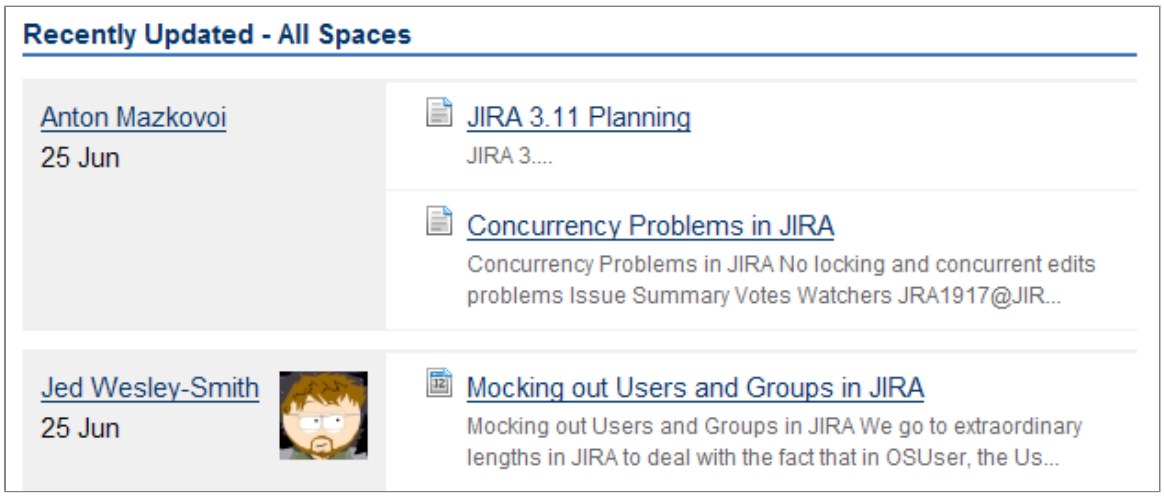
- Default font now Arial instead of Verdana.
- Fresh, clean look and feel - these release notes are an example of the new style.
- If you prefer the original Confluence look and feel, select the **Confluence Classic Theme** when creating a space.
- Improved layouts for email and RSS feeds, helping you to skim-read and classify information quickly.





Personalised comments and Dashboard

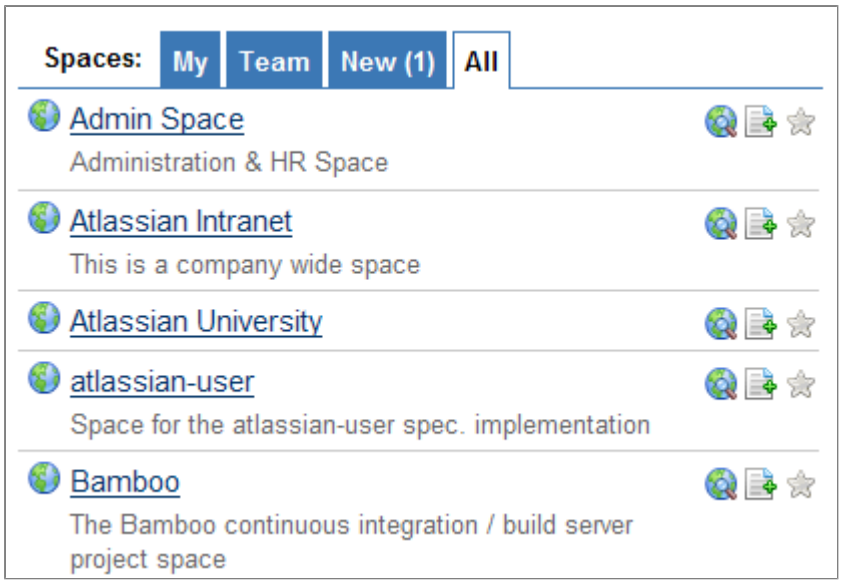
- In the new themes (Confluence Default and Clickr), comments now include the user's profile picture. (See screenshot above.)
- 'Recently Updated' section on Dashboard includes profile picture and a summary of the change.





Space description on Dashboard

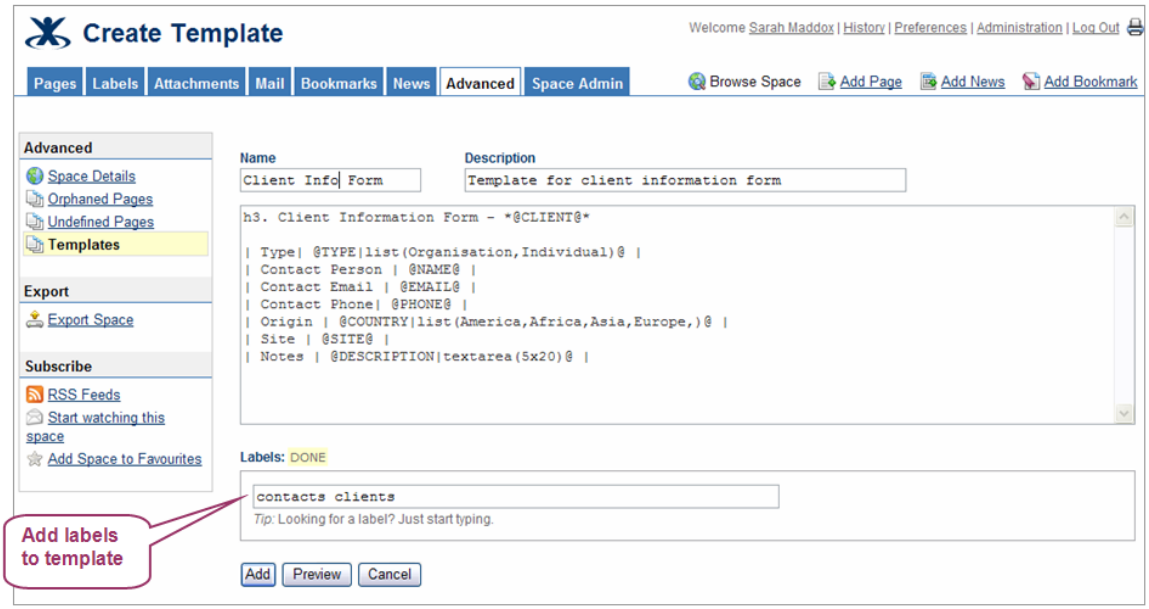
- List of spaces on the Dashboard now includes a short description of the space.



4

Labels on templates

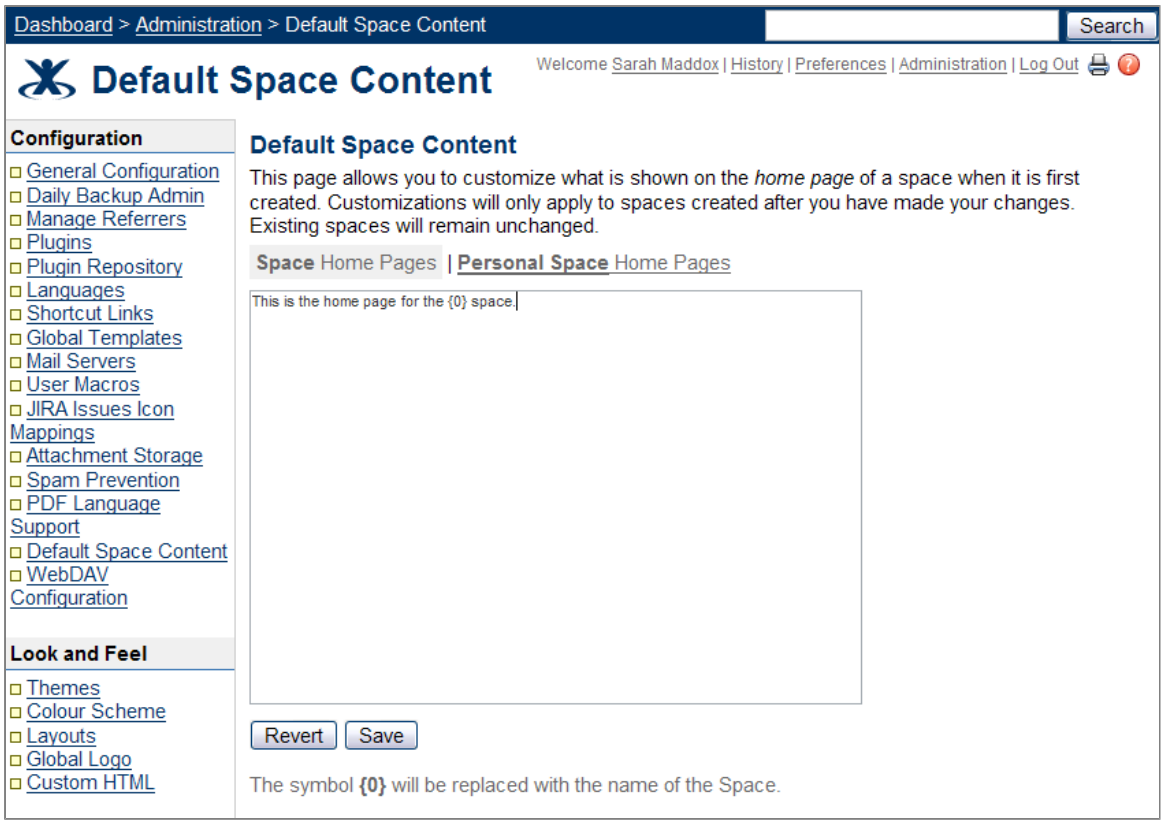
- You can include labels when adding a page template.
- New pages based on the template will automatically include the labels.
- Read the documentation.



5

Default content for space home pages

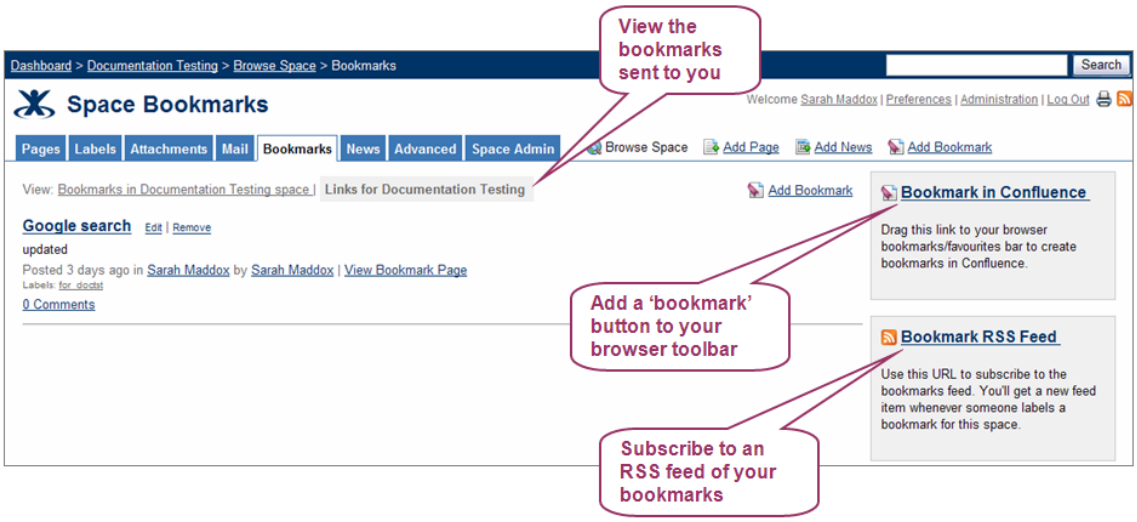
- Confluence site administrators can define default content for a space.
- Home page for newly-added spaces will include the default content.
- Read the documentation.



6

Social Bookmarking plugin now bundled with Confluence

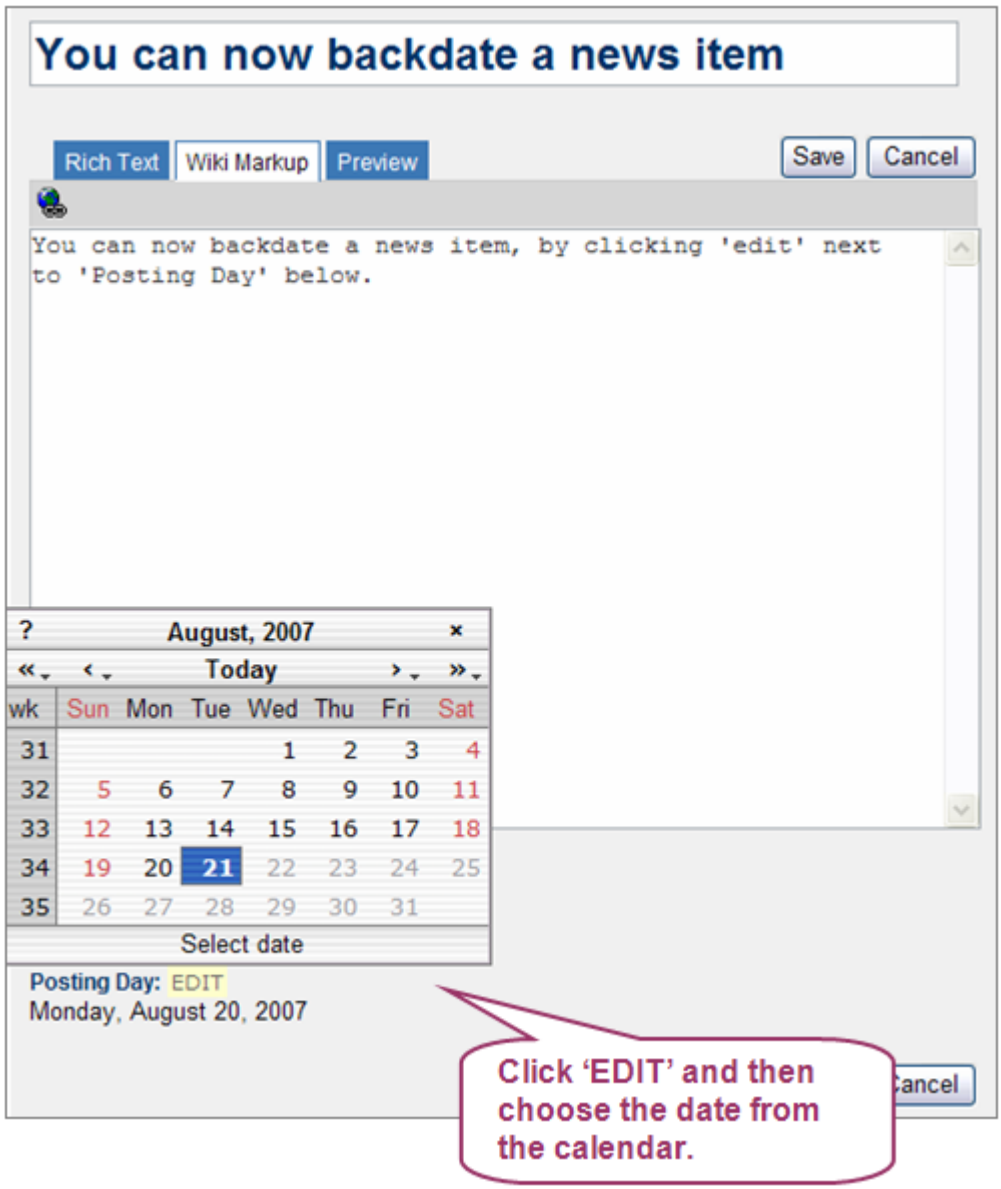
- Use Confluence to share bookmarks with your team.
- Plugin will be **enabled** by default.
- You can create a bookmark for any space in which you have 'create page' permission.
- To view your bookmarks, go to the '**Bookmarks**' tab under '**Browse Space**'.
- Drag the bookmarklet from the 'Bookmarks' tab onto your browser toolbar to create bookmarks any time.
- Use the bookmarks macro to display a list of bookmarks anywhere in Confluence.
- Subscribe to an RSS feed for your bookmarks.
- Read the documentation.



7

Back-dating and renaming news items

- Rename a news item at any time.
- When creating a news item, you can now set the posting date to earlier than today.
- Backdating is also supported by the RPC interface - useful for migrating blog posts from other systems.



8

Plus over 90 other fixes and improvements

- **MySQL 5.0** is officially supported, when used with Confluence 2.5 and above.
- Images generated by macro plugins will now **export to PDF**, .doc and .html formats.
- For plugin developers, the Joda-time library has been upgraded from 0.98 to 1.4 in Confluence 2.6. Plugins that use the date formatting or parsing functionality of Joda-time will need to be recompiled to work with Confluence 2.6.
- Administrators can configure a non-standard port for the Confluence outgoing mail server. The host address can now be specified as **hostname:port**.
- Improved **user migration** when integrating with LDAP: If you have existing Confluence users with the same usernames as LDAP users, you can now avoid duplicate users by configuring the LDAP repository before running the migration. The migration will then ignore users who have the same username as an LDAP user. Read the documentation.
- And more.

Special thanks

We'd like to thank some of our valued community members whose contributions to the open source plugin library have made this version of Confluence even stronger:

- Shannon Krebs, for the BloggingRPC plugin, the Contributors plugin and the Social Bookmarking plugin.
- David Peterson and Bob Swift for the Chart Plugin.
- David Peterson for the IM Presence Plugin, and the Table of Contents plugin.
- Dan Hardiker & Adaptavist for the Plugin Repository Client.

The Confluence 2.6 team

Development
Paul Curren
Tom Davies
Matthew Jensen
Anatoli Kazatchkov
Samuel Le Berrigaud
David Loeng
Charles Miller
Christopher Owen
Agnes Ro
Matt Ryall
Don Willis

UI
Jason Taylor
Stephen Russell

Technical Writing
Rosie Jameson
Sarah Maddox

Oversight & ~~Mis~~management
Mike Cannon-Brookes
Scott Farquhar
Soren Harner
Per Fragemann

Confluence 2.6 Upgrade Guide

Confluence 2.6 brings many popular features which save you time and improve the your wiki's usability. Upgrading to Confluence 2.6 is free for all customers with active Confluence software maintenance as at 27th September 2007.

A new theme brings a fresh, clean look and feel. This is for our customers who have asked for a friendlier interface and improved readability. We've included more social networking features, to enhance the sense of community in your wiki. For example, author photos are now shown in comments and in the 'Recent Updates' on the Dashboard. And the Social Bookmarking plugin is now shipped with Confluence, allowing you to share bookmarks with your team.

Other popular new features include default content for spaces, labels on templates, the ability to backdate or rename news items (blog posts), official MySQL 5.0 support and PDF export of images.

Upgrade Notes

Draft Pages

All draft pages will be discarded during the upgrade process. Confluence administrators should warn users of the Confluence site that drafts will not survive the upgrade.

Plugins

If you are using any third-party plugins, please test them thoroughly before rolling 2.6 into production.

For plugin developers: the Joda-time library has been upgraded from 0.98 to 1.4 in Confluence 2.6. Plugins that use the date formatting or parsing functionality of Joda-time will need to be recompiled to work with Confluence 2.6.

Custom Themes

Custom Confluence 2.5.x themes are expected to be compatible with 2.6 without authors needing to make any change to their existing themes. This is because Confluence will, by default, include all Confluence 2.5.x specific styles automatically. But if you would like to upgrade your theme to use the latest Confluence 2.6 style and typography, you will need to update the way you include stylesheets in your theme. You can read full instructions here.

Custom Page Layout

1. If a space uses a custom decorator page layout, the new Confluence 2.6 decorator is not applied. This may causes GUI oddities, such as:
    - On a page, the View, Edit, Info and Attachments tabs are shown as a vertical bulleted list instead of tabs.
    - Comments do not show properly.

    Fix: Apply the Default Page Layout, then re-insert the custom code.
2. In addition, if you are using pagetree navigation to form a table of contents on the left, you may find that your wiki text becomes *italic* after upgrading to Confluence 2.6.
    - Cause: To create the left-hand panel, you have probably inserted a chunk of HTML/CSS in the space's page layout. The chunk of code may use a `<blockquote>` element to align the body of the page and draw a blue line on its left. Confluence

2.6 stylesheets apply the italic style to blockquotes.
- Fix: Use a new format for your left-hand panel. One possible example is given here.

Steps in detail:

1. Go to the Space Admin screen and click 'Edit' to view your customised Page Layout.
2. Copy the customised code.
3. Cancel the edit.
4. Click 'Reset Default' to apply the new Confluence 2.6 default page layout.
5. If you are using the `pagetree` navigation panel, edit your customised code as described above.
6. On the Space Admin screen, click 'Create Custom' to create a custom page layout.
7. Reinsert your customised code and click 'Save'.

Upgrade Procedure

Upgrading Confluence should be pretty easy. *We strongly recommend that you backup your `confluence.home` directory and database before upgrading!*

You can get the latest version of Confluence here.

If you are upgrading from Confluence 2.2.x or a later version, you can find instructions here.

If you are upgrading directly from 2.1 or earlier, you should also read the 2.2 Release Notes for warnings about the 2.1 -> 2.2 upgrade.

Issues resolved in Confluence 2.6

Below is the full list of issues resolved by Confluence 2.6. You can read the release notes here.

Errors were reported by the JIRA trusted connection.

- APP_UNKNOWN; Unknown Application: {0}; ["confluence:4557196"]

JIRA Issues (176 issues)

Type	Key	Summary	Status	Resolution
	CONF-1576	Rename news (blog posts)	Resolved	Fixed
	CONF-1580	Allow user to specify a non-standard SMTP mail server port	Resolved	Fixed
	CONF-1711	Backdate blog entries	Resolved	Fixed
	CONF-1728	PDF Export error - No meaningful layout in block after many attempts	Closed	Fixed
	CONF-2162	Export Page includes unneeded icons/emoticons	Closed	Fixed
	CONF-2661	addGroup converts specified groupname to lowercase and causes error while retrieving	Resolved	Fixed
	CONF-3050	Links inside html quotes get listed on the "Undefined pages" page	Resolved	Fixed
	CONF-3355	[foo] inside {noformat} appears on the "Undefined Pages" page	Resolved	Fixed
	CONF-3373	Userlister icons absent in HTML export	Resolved	Fixed
	CONF-4186	Heading styles for "Printable View" and "Regular View" are vastly different	Closed	Fixed
	CONF-4419	Order by name	Closed	Fixed
	CONF-4567	Add "display description" parameter to the Macro that lists the spaces on the Dashboard.	Closed	Fixed
	CONF-4739	Multi-page PDF exports get confused by attachments on different pages but with identical file names.	Resolved	Fixed
	CONF-4743	HTML space export does not use either global or space layout for index page	Resolved	Fixed
		DataIntegrityViolationException on MySQL 5.0 during Confluence configuration at the admin		

	CONF-4755	account creation step	Resolved	Fixed
	CONF-4964	Improve consistency of spacing of lists and paragraphs	Closed	Fixed
	CONF-5453	Space with no administration rights won't allow addition of user/group	Resolved	Fixed
	CONF-5608	Add default label(s) to template creation	Resolved	Fixed
	CONF-5849	PDF export does not honor image width property	Resolved	Fixed
	CONF-5954	ConfluenceSoapServiceImpl.changeMyPassword() incorrectly delegates to SoapServiceDelegator.changeUserPassword	Resolved	Fixed
	CONF-6049	Export of Documentation Space to PDF and XML restoration are broken	Resolved	Fixed
	CONF-6284	Pages served over HTTPS that embed Flash movies display a security warning on IE	Resolved	Fixed
	CONF-6361	Use relative font sizes for page text	Resolved	Fixed
	CONF-6528	Images with explicit height and width of 0 are showing up in PDF exports	Resolved	Fixed
	CONF-6544	Images exported into a PDF file are not in a high resolution	Resolved	Fixed
	CONF-6587	Enable browser font resizing	Resolved	Fixed
	CONF-6606	Images created by Gallery Macro are not correctly created into PDF-exports	Resolved	Fixed
	CONF-6738	Long words fail to wrap in PDF exports	Closed	Fixed
	CONF-6829	Allow sales links to remain localized even if user switches language	Resolved	Fixed
	CONF-6951	Second access to same attachment stored in database is not found	Resolved	Fixed
	CONF-6966	Add ability to disable/customize external link icon.	Resolved	Fixed
	CONF-6987	Simplify Rich Text Editor localisation	Resolved	Fixed
	CONF-7399	conflicting messages in cac and confluence internal documentation	Closed	Fixed
	CONF-7414	Officially support mySQL 5.0	Resolved	Fixed
	CONF-7529	Move page operation should not add space keys to links inside noformat or code blocks	Resolved	Fixed
	CONF-7643	Searching for pages with a certain label in ONE space returns the pages related to the label in ALL spaces (when clicking on "Next>>" link on first result page)	Resolved	Fixed
	CONF-7739	Some content migrated from 2.0 to 2.3 fails to render (due to:- system error: java.lang.String java.lang.ClassClastException: java.langString)	Resolved	Fixed
	CONF-7857	Authenticating LDAP users doesn't use the userSearchFilter for its test search	Closed	Fixed
	CONF-7863	blog-posts macro does not work from an included page macro	Resolved	Fixed
	CONF-7864	Remove trailing and prefacing empty character in SPACE name	Resolved	Fixed
	CONF-7875	flowchart or graphviz macro images are not visible in HTML export space		Fixed

			Resolved	
	CONF-7946	Querying Bandana Context during Restore from Backup returns null	Resolved	Fixed
	CONF-7970	Labels that are no longer associated with any content should not be displayed	Resolved	Fixed
	CONF-7985	Rich Text Editor - Improper handling of Line Feed in {code} parts	Resolved	Fixed
	CONF-7987	Daily notification mail contains unsubstituted term "${baseUrl} "	Resolved	Fixed
	CONF-7997	The code issue still occours.	Resolved	Fixed
	CONF-7998	Code macro fails to hide brackets	Resolved	Fixed
	CONF-8133	Upgrade to Tomcat 5.5.23 in standalone	Resolved	Fixed
	CONF-8152	Link to results in other spaces given when searching all spaces	Resolved	Fixed
	CONF-8163	Space creation broken when running Confluence in Shared Mode and not having space groups	Resolved	Fixed
	CONF-8170	A link within a comment breaks when the linked page is renamed	Resolved	Fixed
	CONF-8203	Renderer javadoc wrong	Closed	Fixed
	CONF-8249	Rename 'Maximum Attachments per Form' setting to "... per Upload"	Resolved	Fixed
	CONF-8297	Attachments macro doesn't list previous versions of attachments though "old' property is set to true	Resolved	Fixed
	CONF-8339	Image and Link editing in Rich Text is not fully translated when Foreign Language pack is applied.	Closed	Fixed
	CONF-8343	Poor quality of Thumbnails	Resolved	Fixed
	CONF-8348	Confusion regarding enabling "External User Management" option in General Configuration menu when integrating with LDAP	Resolved	Fixed
	CONF-8392	Sort favourite spaces alphabetically in search drop-down	Resolved	Fixed
	CONF-8397	Bundle Social Bookmarking 1.0 plugin	Resolved	Fixed
	CONF-8402	README.txt contains out of date information about support	Resolved	Fixed
	CONF-8407	"Restore backup" does not detect invalid ZIP files, exceptions are rendered	Resolved	Fixed
	CONF-8417	Missing plugin-info knocks Confluence over	Resolved	Fixed
	CONF-8422	Plugin respository exception	Closed	Fixed
	CONF-8423	NullPointerException when editing a group from Manage Groups	Closed	Fixed
	CONF-8434	Exporting spaces with pages containing a ContentPermission may throw an exception	Resolved	Fixed
	CONF-8450	On the user management page, if you click "Search" without entering results you get a java exception	Resolved	Fixed
	CONF-8451	Duplicate javamail / mail JARs		Fixed

			Resolved	
	CONF-8454	Hyperlinks containing round brackets are broken	Resolved	Fixed
	CONF-8459	Export Layouts don't work for Spaces, only for Site	Resolved	Fixed
	CONF-8468	Members of groups in Group Management are printed in ugly technical way	Resolved	Fixed
	CONF-8489	Dynamic Tasklist does not work with a "-" sign in the Title of the Tasklist	Resolved	Fixed
	CONF-8491	A small typo on the mailto link on the Administration view user profile.	Resolved	Fixed
	CONF-8503	Upgrade joda-time dependency from version 0.98 to version 1.4	Resolved	Fixed
	CONF-8506	Preformatted text from the Rich Text Editor removes link properties	Resolved	Fixed
	CONF-8523	Edit Space Permissions failure in IE7 (works in FireFox 2.0.0.3)	Resolved	Fixed
	CONF-8533	Rich-Text-Editor failed to load for some users	Resolved	Fixed
	CONF-8543	Junit macro broken	Resolved	Fixed
	CONF-8544	Underscores used to work in template values but now don't	Resolved	Fixed
	CONF-8554	Misleading errormessage when trying to edit a nonexisting users group (via direct URL-access)	Resolved	Fixed
	CONF-8573	No warning of page currently being edited in Confluence Clustered Environment	Resolved	Fixed
	CONF-8580	Indexing unprintable/encrypted PDFs fails	Resolved	Fixed
	CONF-8586	Creating a page on an not authorized space	Resolved	Fixed
	CONF-8593	DefaultBreadcrumbsManager generates invalid HTML	Resolved	Fixed
	CONF-8600	Exporting a Space sometimes fails, throwing a Nullpointer-Exception	Resolved	Fixed
	CONF-8608	Improve Indexing Error Handling	Resolved	Fixed
	CONF-8623	Improve and speed up plugin resource loading	Resolved	Fixed
	CONF-8630	errors.jsp should return HTTP 500 Server Error instead of HTTP 200 OK when errors are present	Resolved	Fixed
	CONF-8631	Exports should be done asynchronously, not to give the impression that Confluence has hung	Resolved	Fixed
	CONF-8634	When 'AND' is used to search, the label 'and' is matched	Resolved	Fixed
	CONF-8644	It is not possible to globally disable mail archiving any more (Regression)	Resolved	Fixed
	CONF-8657	Rich Text Editor broken with IE6 and French language pack	Resolved	Fixed
	CONF-8658	In rare cases new users don't get added to the proper group and therefore can't use Confluence	Resolved	Fixed
	CONF-8665	exporting a space to XML incorrectly includes comments, even when 'Include comments' is		Fixed

		deselected.	Resolved	
	CONF-8682	"Not Permitted" error when I try to use the time sheet template	Resolved	Fixed
	CONF-8690	REGRESSION - Copy Page not permitted (extranet)	Resolved	Fixed
	CONF-8742	Hyperlinks do not wrap	Closed	Fixed
	CONF-8746	Make threaded comments the default for new installations	Resolved	Fixed
	CONF-8754	Insert Link Popup Page Icon links are incorrect	Closed	Fixed
	CONF-8758	New theme does not have permalink for comments	Resolved	Fixed
	CONF-8762	Multithreaded access to HashMap, can cause infinite loop	Resolved	Fixed
	CONF-8767	The create page templates page does not display action errors	Resolved	Fixed
	CONF-8770	Email address exposure - email hidding option is ignored in user lookup	Resolved	Fixed
	CONF-8771	View template bugs: Labels are not displayed, content is (wrongly) aligned to the right	Resolved	Fixed
	CONF-8786	confluence-administrators no longer fits into its cell in Global Permissions	Resolved	Fixed
	CONF-8787	Clickr theme's fonts get overridden by new stylesheet in wiki content and RTE	Resolved	Fixed
	CONF-8791	Set-up wizard theme broken	Closed	Fixed
	CONF-8802	New quoted typography doesn't un-italicise markup	Resolved	Fixed
	CONF-8804	Typography Fixes	Closed	Fixed
	CONF-8813	New theme breaks Calendar plugin completely	Resolved	Fixed
	CONF-8814	Foldernav styles have disappeared in new stylesheet	Resolved	Fixed
	CONF-8815	Tabnav ID appears multiple times in markup; should be a class	Resolved	Fixed
	CONF-8827	Viewing a historical version of a page has weird note styles at the top	Resolved	Fixed
	CONF-8831	Rich text editor 'Insert Link' button triggers pop-up blocker	Resolved	Fixed
	CONF-8854	Modifying the title of page removes the page formatting of pages linking to it	Resolved	Fixed
	CONF-8855	Page title length is not being validated, leading to errorpage for titles longer than 255 characters	Resolved	Fixed
	CONF-8856	code macro looks bad in firefox on linux	Resolved	Won't Fix
	CONF-8857	Comment UI (clickr theme) - blue box in the left top corner of comment editing panel	Resolved	Fixed
	CONF-8858	Incoming Links to Blog Posts broken	Resolved	Fixed
	CONF-8869	JavaScript error occurs when editing a page without having permission to restrict it		Fixed

			Resolved	
	CONF-8904	Footer broken in Clickr theme on extranet	Resolved	Fixed
	CONF-8906	Merge Charles' 2.5.5 branch commits into trunk	Resolved	Fixed
	CONF-8909	'Preview End' box doesn't join up with sides of preview box on Safari	Closed	Fixed
	CONF-8912	Panel macro shows last line grey on EAC	Closed	Fixed
	CONF-8917	XSS vulnerability: space name and key not validated nor escaped	Resolved	Fixed
	CONF-8944	Profile settings are lost when using Crowd with Confluence	Closed	Fixed
	CONF-8950	XSS vulnerability in app/spaces/listattachmentforspace.action	Resolved	Fixed
	CONF-8956	stored XSS vulnerability in app/themes/leftnavigation/configuretheme.action	Resolved	Fixed
	CONF-8962	Stylesheet not loading on setup	Resolved	Fixed
	CONF-8978	Vulnerability against DoS attack via labels	Resolved	Fixed
	CONF-8979	Vulnerability against DoS attack at permission setting	Resolved	Fixed
	CONF-8980	XSS vulnerability at "Edit Space Permissions"	Resolved	Fixed
	CONF-8986	Tab spacing is wrong on Classic and Left Navigation theme	Resolved	Fixed
	CONF-8987	Tabs are missing using the Theme Builder theme	Resolved	Fixed
	CONF-8988	"Robert Paulsen" misspelled in notation guide	Resolved	Fixed
	CONF-8993	Reflected XSS Vulnerability in the Feed Builder	Resolved	Fixed
	CONF-8994	Change renewal hyperlink to use new website redirects	Closed	Fixed
	CONF-9000	OutOfMemoryError's during indexing	Resolved	Fixed
	CONF-9007	Moving an attachment does not update the respective links	Resolved	Fixed
	CONF-9024	Update professional German and French translations in Confluence	Closed	Fixed
	CONF-9054	NPE thrown when attempting to export space as PDF	Resolved	Fixed
	CONF-9060	Missing text in breadcrumbs when viewing changes since last login	Resolved	Fixed
	CONF-9067	Division by zero in SnipSnapImporter	Resolved	Fixed
	CONF-9111	Add hints encouraging users to provide a profile picture	Resolved	Fixed
	CONF-9117	Confluence API supports adding user with null password, but users will null passwords produce NullPointerException when using the osuser to atlassian-user migration utility jsp	Resolved	Fixed
	CONF-9124	Documentation update: MySQL 5.0 is supported when used with Confluence 2.5 and above		Fixed

			Resolved	
	CONF-9167	the '{code}' and '{noformat}' macros truncate long lines of code	Closed	Fixed
	CONF-9174	Space and global permissions show $user.fullName for LDAP users which have been deactivated	Resolved	Fixed
	CONF-9179	Left nav theme icon missing in select theme page	Resolved	Fixed
	CONF-9180	Can't comment in Clickr theme	Resolved	Fixed
	CONF-9187	Note joda-time incompatibility in 2.6 release notes	Closed	Fixed
	CONF-9188	Remove page removal not permitted text from edit page	Resolved	Fixed
	CONF-9195	Confluence 2.5.6 ldap configuration failing on osuser2atluser.jsp migration	Resolved	Fixed
	CONF-9201	Redundant information in new RSS feed format	Resolved	Fixed
	CONF-9214	News posts are always shown as having been edited	Resolved	Fixed
	CONF-9215	Internet Explorer issues specific to 2.6	Resolved	Fixed
	CONF-9218	Themes can override admin layout	Resolved	Fixed
	CONF-9231	Entering invalid page labels hides label edit section	Resolved	Fixed
	CONF-9237	AuthenticatedUserThreadLocal does not clear on Logut Action	Resolved	Fixed
	CONF-9249	Duplicate dependencies in WEB-INF/lib/	Resolved	Fixed
	CONF-9251	"Too many open files" error during index operations	Resolved	Fixed
	CONF-9276	Recompile SocialBookmarking plugin against 2.6 final	Closed	Fixed
	CONF-9303	Caching is not enabled by default for the hibernate repository	Resolved	Fixed
	CONF-9346	Include parent Comment in comment notifications	Resolved	Fixed
	CONF-9401	Weird string appears above the edit box when *editting* an existing comment	Resolved	Fixed
	CONF-9434	Enabling Caching for Hibernate Repository causes net.sf.hibernate.LazyInitializationException error	Resolved	Fixed
	CONF-9450	Don't put the "Cancel" button so close to the "Save" button when editing pages	Resolved	Fixed
	CONF-9452	Accessing an invalid attachment download URL causes exception	Resolved	Fixed
	CONF-9465	Confluence jar's pom contains unsubstituted property for artifactid	Resolved	Fixed
	CONF-9475	Make simple modifications to default layout to allow page-by-page customization of the UI	Resolved	Fixed
	CONF-9509	java.lang.RuntimeException: Unable to delete working directory	Resolved	Fixed
	CONF-9521	TinyMceServlet has a memory leak		Fixed

			Closed	
	CONF-9555	Improve macros.vm backwards compatibility	Resolved	Fixed
	CONF-9589	Link to an attachment combined with an image renders incorrectly	Closed	Fixed
	CONF-9643	Canceling of "Remove Page" should take user back to where he/she came from	Closed	Fixed
	CONF-9810	Space is left with no "admin" and users being locked-out of administering a space	Resolved	Fixed
	CONF-9819	In edit mode (RTE) code style text shows with variable font instead of fixed font	Closed	Fixed

Release Notes 1.0

Confluence 1.0

Ring the bells. Break out the champagne. Paint the town red. Dance naked in the streets. Or at least that's what we'll be up to tonight, because we're finally releasing Confluence 1.0!

Confluence 1.0 represents the hard work and dedication of quite a few people. Obviously there's the development team: Ara, Armond, Charles, Dave, Mike and Ross, but there's also the JIRA team and the rest of the guys at Altlassian who have offered support, advice, and loud music. A huge thanks also goes out to everyone involved in the beta-testing program. Your real-world use of Confluence gave us invaluable suggestions and bug-reports, and we apologise for the times we've messed up your databases on the way.

We're pretty proud of all the cool stuff we've managed to fit into Confluence already, and we're looking forward to making it even better in forthcoming versions (Remember, a license entitles you to a year of upgrades).

The only sad moment for the day is that for the 1.0 release, we had to lose emoticons. The regular expression responsible for turning into a smiley-face was causing pages to take *ten seconds* to render. Commenting out the filter lowers the rendering time of even highly complex pages to a few hundred milliseconds. We promise, the smileys will return in a future version, faster and stronger than ever! Or if you don't mind the cost, you can re-enable them by uncommenting the emoticonFilter line in `wikiSubsystemContext.xml`

Users of late Release Candidate builds will notice a substantial improvement in performance as a result.

Contents

1. New Features
2. Notable Bug-fixes since RC6
3. Outstanding Issues
4. Upgrading from RC6
5. Notable Features from Previous Releases

See also: Issues Resolved for 1.0

New Features

For once, we managed to resist the urge to add any new features this week, since we were busy cleaning everything up for today's release.

Notable Bug-fixes

- We have verified that you can now use Confluence's internal datasources with MySQL (for users of RC5 or earlier, see the upgrade section below for information on how to change to the production-ready datasources)
- Fixed the database transaction problems that were causing problems with user/group management, page renaming and commenting
- [~user] links no longer place the context path in the link twice
- Fixed a rendering regression that caused text to be spaced incorrectly within panels
- The 404 error page no longer causes a NullPointerException
- The recent comments RSS feed returns a valid RSS version number
- The 'diff' link now works in HTML edit notification emails when Confluence is installed in a non-root context.

Outstanding Issues

- The Emoticon filter is currently shipped disabled, because it was causing page-load times to go through the roof CONF-963
- You can not complete the Confluence setup process if you have external user-management enabled. Set up Confluence *before* setting up external user-management. CONF-950
- It is possible to make a "create page" link to a page with invalid characters in its title (which will thus fail to be created) CONF-810
- If you are using PostgreSQL, please make sure you use the JDBC3 version of the Postgres JDBC drivers. Some users have encountered problems with the JDBC2 drivers.
- Links to images that are page-attachments are broken in email notifications CONF-878
- There continue to be problems with high-bit characters in pages when Confluence is deployed on Resin CONF-569

Upgrading from RC6

We haven't made any changes to the database schema, so you should just be able to install the new version of Confluence, point it at your existing confluenceHome directory, and carry on as before. I would, however, recommend making a backup before you upgrade, just in case.

If you:

1. have not gone through the Confluence database setup steps since Confluence 1.0RC5
2. are running the embedded HSQL database, or are connecting to a database directly (not through an Application Server Datasource)

Then you will need to add a few connection pooling properties to confluence.cfg.xml file in your confluenceHome directory. Edit the file, and insert the following somewhere in the <Properties> section:

If you do not fall into the category described above, or if you find these lines are already in the file, you *do not* need to do this!

```
1.<property name="hibernate.c3p0.max_statements"><![CDATA[0]]></property>
2.<property name="hibernate.c3p0.min_size"><![CDATA[0]]></property>
3.<property name="hibernate.c3p0.timeout"><![CDATA[30]]></property>
4.<property name="hibernate.c3p0.max_size"><![CDATA[4]]></property>
```

This will enable c3p0 connection-pooling, which is far more reliable than the default pooling that we were using previously.

Notable Features from Previous Releases

Here's a quick retrospective of some of the cool things we have added to Confluence during the beta- and RC- releases. It is by no means comprehensive, but you can follow the links to the original release notes for an idea of just how busy we've been in recent months.

Confluence 1.0a2 (November 6, 2003)

- Shortcut links allow you to create special URL Shortcuts to point to existing web services: i.e. CONF-195
- Undefined and Orphaned Pages reports
- Revert page to previous version

Confluence 1.0a3 (December 5, 2003)

- Comments for pages
- Links in exported PDFs are internal links if the destination page has also been exported, external links otherwise
- Personal History popup keeps track of the pages you have visited
- Page Information screen lists all the incoming and outgoing links on a page
- Printable version of every page

Confluence 1.0b1 (December 19, 2003)

- Pages that have changed since your last login are highlighted in the 'recent changes' list
- Pages can be organised in parent-child hierarchies, allowing for clearer organisation of pages within a space
- Page renaming refactors links to the page, wherever they may appear in the site
- Search works across all content in the site, including comments, space descriptions and user information
- A new permission was added for creating comments
- {color} macro for coloured text
- Parameters allowed on image links

Confluence 1.0b3 (January 19, 2004)

- Beta2 was released a few days earlier, but we'd prefer not to talk about it.
- Searching works across PDF, HTML, XML and Word attachments
- The colour-scheme and page decorators can be configured across the site
- The site homepage is configurable
- Emoticons 
- New macros: {include}, {include-html}, {rss}, {search}, {jiraissues}, {junitreport}
- User profiles can be longer than 255 characters
- Much better handling of anonymous contributions

Confluence 1.0b4 (January 26, 2004)

- FatCow, our acceptance-testing framework was introduced to the world
- Confluence now supports chronologically organised content: blog posts
- You can now move pages between spaces, and have their links updated accordingly
- en – and em — dashes are supported
- You can link directly to attachments
- New macros: {blog-posts}, {anchor}, and an improved {rss} macro

Confluence 1.0rc1 (February 6, 2004)

- XML-RPC and SOAP APIs allow you to programmatically interact with Confluence
- Trackback allows Confluence to notify other sites of links, and be notified by them
- User notifications on page and space editing were overhauled
- Look and Feel configuration was enabled individually for each space
- Linking to space descriptions and user profile pages directly is now possible: Confluence User Community, Mike Cannon-Brookes
- History popup tracks more than just pages

Confluence 1.0rc2 (February 13, 2004)

- Improvements to the remote API
- RSS and HTML macros now use HTTP proxies if configured to

Confluence 1.0rc5 (February 20, 2004)

- RC3, RC4 and RC5 were released within days of each other, in a flurry of mad bug-fixing
- Page templates can now be filled in in-line
- If a page is moved or renamed and you go to the URL it used to inhabit, Confluence will try to direct you to its new location
- Exported PDFs now incorporate the site's stylesheet

Confluence 1.0rc6 (March 5, 2004)

- You can prevent people signing up, for private Confluence installations
- You can configure Confluence to mask user email addresses
- User management can now be shared with JIRA
- Page templates can now contain drop-down menus and text areas
- New macro: {html}

Demonstration Site

Atlassian have set up a **demonstration space** called 'the Confluence Test Space' so that you can try out Confluence for yourself.

Click the link above to go to the demo space

The demonstration space has been configured so that anyone can create or edit pages within it. When using it, you should keep the following in mind:

- This server may not be running exactly the same version of Confluence as is available for purchase. While we try to keep the two in sync, there may be differences between the demo site and the downloadable version of Confluence. The version that the site is running can be found at the bottom of each page.
- Because the space is open for anyone to edit, Atlassian cannot be responsible for the content on it. Do not rely on anything you read in the test space.
- The contents of the test space may be edited or deleted at any time. Periodically, Atlassian may restore the test space from backup, deleting everything that has been added to it by visitors.
- For a full demonstration of Confluence, including its administrative features, you should download an evaluation instead.
- If you have any questions, contact us and we'll be happy to answer them.

Want to try Confluence yourself?
Experiment with Confluence in our demonstration space >>

Issues Resolved for 1.0

Errors were reported by the JIRA trusted connection.

- APP_UNKNOWN; Unknown Application: {0}; ["confluence:4557196"]

JIRA Issues (36 issues)

Type	Key	Summary	Status	Resolution
	CONF-797	Provide documentation about shared user database (Jira+Confluence)	Closed	Fixed
	CONF-219	Import from other wiki's	Closed	Fixed
	CONF-1060	PDF Indexing	Closed	Invalid
	CONF-764	Create jspwiki importer	Closed	Fixed
	CONF-936	Page save and preview is slow for large pages with lots of markup	Closed	Fixed
	CONF-944	Linking Images (thumbnailing / popups)	Closed	Duplicate
	CONF-916	Joined Numbered Bullet Points	Closed	Fixed
	CONF-551	export page is unusable for large spaces		Fixed

			Closed	
	CONF-876	Possibility to hide email addresses	Closed	Fixed
	CONF-913	Blog post results in the search page don't have edit and remove icons next to them	Closed	Fixed
	CONF-943	Can not delete users under Postgres	Closed	Fixed
	CONF-833	rss feed cannot read a confluence feed	Closed	Fixed
	CONF-941	Access Administration.action results in a Page Not Found	Closed	Fixed
	CONF-883	Security Management vs Group confluence-users	Closed	Cannot Reproduce
	CONF-937	User to group assignments don't stick	Closed	Duplicate
	CONF-904	java.lang.RuntimeException: Caught an Ognl exception while getting property space	Closed	Cannot Reproduce
	CONF-933	Delegating user management to JIRA causes Exception	Closed	Fixed
	CONF-919	commentrss.action returning rss 0.92	Closed	Fixed
	CONF-1008	Cannot add page comments	Closed	Fixed
	CONF-1057	Daily backup doesn't seem to do anything	Closed	Invalid
	CONF-896	Batch update row count wrong: 0	Closed	Cannot Reproduce
	CONF-806	Email report shows incorrect # of comments on a topic	Closed	Fixed
	CONF-946	NPE trying to rename the space home page	Closed	Fixed
	CONF-940	Page edited email does not use full URL for 'View Changes' link	Closed	Fixed
	CONF-931	apostrophe followed by closing parenthesis rendered as wink emoticon	Closed	Fixed
	CONF-932	exporting table to PDF fails	Closed	Fixed
	CONF-924	[~username] links are wrong if Confluence is run through forwarding	Closed	Duplicate
	CONF-920	Paragrap spacing issues, Panel and after headings	Closed	Fixed
	CONF-923	can't assign users to groups	Closed	Fixed
	CONF-1161	Must remove correct email adress from profile to prevent spam	Closed	Duplicate
	CONF-917	Tables not getting formatted correctly within the lists!	Closed	Fixed
	CONF-918	Sticky "Add Comment" Textfield	Closed	Won't Fix
	CONF-643	soap wsdl is broken (and you need to add a remote api component)	Closed	Fixed
	CONF-897	[~profile] links place context name twice in link		Fixed

			Closed	
	CONF-899	Image attributes don't work when embedded in links	Closed	Fixed
	CONF-921	Regression: Spacing within a panel.	Closed	Fixed

Release Notes 1.0.1

Confluence 1.0.1

Over the course of the last month of supporting Confluence 1.0, there are a number of patches that we have been distributing to fix specific problems our users have encountered. Confluence 1.0.1 is a maintenance release into which all these patches have been integrated.

Who Should Upgrade?

All the issues that were resolved for this release are listed below. We have not snuck in any other changes: what you see here in the release notes is precisely what you will get. If you find something on the list that directly affects you, or that you feel justifies the effort of an upgrade, then do so. Otherwise, feel free to stick with 1.0.

Upgrade Procedure

If You Have Customised `osuser.xml`

If you have customised Confluence's user-management, for example to integrate it with LDAP or JIRA, you will have to integrate your changes to account for the caching OSUser providers we introduced in Confluence 1.0.1. See this document for more details: [Confluence Caching OSUser Provider]. Updated instructions for integrating with JIRA user management are here: Delegate user management to use JIRA logins

Otherwise

To avoid the possibility of data-loss, you should back up your ConfluenceHome directory and your database before upgrading, and perform a full backup from within the application.

Changes in 1.0.1

See also Issues Resolved for 1.0.1

Direct Attachment Links

In response to a loud demand from users, links to attachments using the [^attachment.jpg] syntax will download the attachment file directly, instead of linking to an anchor in the destination page.

Sybase ASE Support

Confluence 1.0.1 resolves the following problems that were causing Confluence not to work with Sybase ASE 12.5.1:

- ORDER-BY mappings failing on Sybase (CONF-1021)
- Can't add users under Sybase (CONF-1022)
- Weird datatype error under Sybase (CONF-1024)
- Sybase doesn't like complicated distinct selects (404 page fails) (CONF-1025)
- Backup import fails under Sybase (CONF-1063)

These bug-fixes may also improve Confluence's compatibility with other untested databases. They will not, however, have any effect on Confluence's operation against PostgreSQL, MySQL or HSQL.

Microsoft SQL Server Dialect in Setup Page

The Microsoft SQL Server database dialect was missing from the database setup page. It has now been added to the list. Confluence has not yet been tested on Microsoft SQL Server, and the usefulness of this option is not yet guaranteed. However, since the Sybase issues above are now resolved and SQL Server belongs to the same family as Sybase, it would be well worth a try.

JIRA User Provider Caching

Users who were linking their user management to JIRA's using the supplied provider were experiencing significant performance problems as a result. 1.0.1 introduces caching to the user provider, which should speed up these installations significantly.

Global Reports Visibility

Under Confluence 1.0, the global "undefined pages" and "orphaned pages" reports did not properly filter out pages that the user could not see. The user could not see the content of any page they did not have access to, but they could learn of the existence of (and names of) pages and spaces they were not permitted to see. This bug is fixed in 1.0.1

Locale-Independent Dates in Backup/Restore

In Confluence 1.0, dates were written into backup files using a localised representation of the month. As such, if you exported Confluence data from a server in one locale it might not import successfully into a server with a different Locale setting. Confluence 1.0.1 still recognises the 1.0 export format, but its own exports will write out dates in a locale-independent format.

As noted above, this means that data exported from Confluence 1.0.1 *can not* be imported successfully into Confluence 1.0.

Fix Browser Crash on Viewing Some Templates

Previously, if you created a template containing no variables, then anyone attempting to preview or use that template would have their browser hang in an infinite Javascript loop. Confluence 1.0.1 fixes this bug.

Typo Fixed on User Group Editing Page

A single-character change from "privilage" to "privilege".

Issues Resolved for 1.0.1

Errors were reported by the JIRA trusted connection.

- APP_UNKNOWN; Unknown Application: {0}; ["confluence:4557196"]

JIRA Issues (16 issues)				
Type	**Key**	**Summary**	**Status**	**Resolution**
	CONF-1007	Ability to specify attachment link behavior	Closed	Fixed
	CONF-1065	Add Caching to JIRA user providers	Closed	Fixed
	CONF-1020	Dates are localised on backup, lost on restore	Closed	Fixed
	CONF-1021	ORDER-BY mappings failing on Sybase	Closed	Fixed
	CONF-1043	Missing Dialect Class Name for Microsoft Sql Server	Closed	Fixed
	CONF-1063	Backup import fails under Sybase	Closed	Fixed
	CONF-1003	typo on edit user groups	Closed	Fixed
	CONF-1028	JavaScript runs bananas when viewing Templates	Closed	Duplicate
	CONF-1114	The emoticons don't seem to work....	Closed	Duplicate
	CONF-1038	Template with no form fields can't be previewed	Closed	Fixed
	CONF-1070	Users without permissions can see undefined pages	Closed	Fixed
	CONF-1005	Script causing browser to hang	Closed	Duplicate
	CONF-1055	"Global reports" visibility bug	Closed	Fixed
	CONF-1024	Weird datatype error under Sybase	Closed	Fixed
	CONF-1025	Sybase doesn't like complicated distinct selects	Closed	Fixed
	CONF-1022	Can't add a user under Sybase	Closed	Fixed

Release Notes 1.0.3

Confluence 1.0.3

Confluence 1.0.3 is another maintenance release, hopefully the last maintenance release we will need on the 1.0 branch before moving full steam ahead towards version 1.1. It fixes some bugs regarding the remote SOAP/XML-RPC API, the display of the 404 error page under Postgresql, and the display of the site's stylesheet under certain versions of Internet Explorer.

The Confluence development team are now working hard on 1.1, which will contain significant improvements and new features. Remember, a Confluence license entitles you to a year of upgrades, so if you buy 1.0.3 today, you'll be able to upgrade to 1.1 for no extra charge when the time comes.

See also: Issues Resolved for 1.0.3

Who should upgrade?

Confluence users should be running at least Confluence 1.0.2. Versions prior to 1.0.2 contain a bug that will cause their licenses to expire a month after the build date listed in the page footer, regardless of the true expiry date of the license. Versions 1.0.2 and later do not contain this bug. (Note: In the absence of this bug, only trial licenses ever expire. Purchased licenses for Confluence *do not expire*.)

All the issues that were resolved for this release are listed below. We have not snuck in any other changes: what you see here in the release notes is precisely what you will get. If you find something on the list that directly affects you, or that you feel justifies the effort of an upgrade, then do so. Otherwise, feel free to stick with 1.0.2

Upgrade Procedure

Confluence 1.0.3 contains no database or configuration file changes, so you should just be able to unpack it on top of your existing Confluence installation:

1. Shut down the Confluence server
2. Back up `confluence/WEB-INF/classes/confluence-init.properties`, and if you have customised it, `confluence/WEB-INF/classes/osuser.xml`
3. Unpack Confluence 1.0.3 in the same location as your existing Confluence installation
4. Restore the two files you backed up in step 1 to the 1.0.3 installation.
5. Start Confluence

To avoid the possibility of data-loss, you should back up your ConfluenceHome directory and your database before upgrading, and perform a full backup from within the application.

Changes in 1.0.3

Global RSS Feed Fixes

The global RSS feeds found on the Dashboard page were not being generated correctly: an extra ')' was being added to the end of links. This has been fixed in 1.0.3

Page Not Found Fix for Postgresql

Users running Confluence against a Postgresql database were encountering a system error whenever someone attempted to visit a page that did not exist. This has been fixed.

Remote API Fixes

Several outstanding issues with the remote API that were holding back the [TimTam|] client have been fixed. In addition a `getVersion` method has been added to allow client authors to determine which version of Confluence a server is running, and adjust their features to match. The full description of Confluence's remote API is here: Remote API Specification

- The WSDL file now respects the server's configured base URL, so SOAP can be used on servers that are behind a proxy.
- `getPermissions` now recognises when a user is in the `confluence-administrators` super-user group.
- `storePage` handles re-parenting a page correctly
- `storePage` will now allow a page to be renamed (all links to the page are automatically redirected)

Other Fixes

- You can now comment on a blog post when the title contains non-US-ASCII characters
- The bug that was causing some versions of Internet Explorer 6 to not display the site's stylesheet has been fixed

Issues Resolved for 1.0.3

Errors were reported by the JIRA trusted connection.

- APP_UNKNOWN; Unknown Application: {0}; ["confluence:4557196"]

JIRA Issues (14 issues)

Type	Key	Summary	Status	Resolution
	CONF-1123	Add getVersion to remote api	Closed	Fixed
	CONF-1100	404 page dies under Postgresql	Closed	Fixed
	CONF-1101	Global RSS feed links have extra)	Closed	Fixed
	CONF-1091	storePage does not handle re-parenting correctly	Closed	Fixed
	CONF-974	Store Page fails when page is renamed.	Closed	Fixed
	CONF-1140	Remote API permissions not respecting "superuser" group	Closed	Fixed
	CONF-1068	Can not comment on blog posts which title contain special characters	Closed	Fixed
	CONF-1088	SOAP Service broken on confluence.at	Closed	Fixed
	CONF-1186	Some versions of IE6 can't see the stylesheet	Closed	Fixed
	CONF-1166	SOAP interface not using Base URL	Closed	Fixed
	CONF-1083	Cannot Delete User from Group confluence-adminstrators	Closed	Won't Fix

	CONF-1200	Permissions check for commenting fails	Closed	Fixed
	CONF-982	Page not found blows up with oracle dialect	Closed	Fixed
	CONF-1118	Some SQL Exception raise randamly	Closed	Duplicate

Release Notes 1.0a2

Confluence 1.0a2

Confluence 1.0a2 was released November 6, 2003.

New Features

- Shortcut links (CONF-195 & CONF-247) - shortcut links enable you to easily make links to any 'web service' from within the syntax of a Confluence page. Simply specify your shortcut and the URL to link to (eg `"google", "http://www.google.com/search?q=$1"`) to add your own links. You could use this to link directly to JIRA installations, Google, intranets, searches or anything that has a 'URI' based interface.
- Revert page to a previous version (CONF-32) - revert to a previous version of any given page from the list of previous versions.
- Forgot password function (CONF-88) - retrieve your password if you happen to have forgotten it
- Report: undefined pages (CONF-197) - list of all pages which are linked to from other pages, but not defined.
- Report: orphaned pages (CONF-198) - shows all the pages which are never linked to (ie have content but no way to get to them).
- Interstitial help page after creating a space (CONF-251)
- 'Notify me', and HTML/text email preferences, with HTML email content
- Links to non-exported pages are now maintained in exported PDFs (the links point back to the online Confluence installation)
- Custom Radeox macros can now be used within Confluence - this enables you to write your own macros to produce custom content
- XML import/export of a single space (CONF-199) - useful for porting spaces between installations
- XML import/export of a single page (CONF-200) - useful for creating pages automatically, or sync'ing pages between installations.
- Many other minor fixes, improvements and performance tweaks

Release Notes 1.0a3

Confluence 1.0a3

Confluence 1.0a3 was released December 5 2003

New Features

- Completely rearchitected around the spring framework
- Page comments (CONF-235) - users can now attach comments to a page.
- Email notification for new comments.
- Supports non-ASCII characters in page names (CONF-297)
- Improved handling of internal and external links in exported PDFs
- Editing popup to help with inserting links (internal and external) into pages (CONF-129)
- Personal history (CONF-196) - a popup window containing your recently visited pages.
- Page information screen (CONF-270) lists all the incoming and outgoing links for a page, and all attachments
- User profile UI has been enhanced (CONF-172)
- Dashboard and Space summary pages have had a big UI overhaul)
- Brief help on Wiki notation now available on create and edit pages. (CONF-295)
- Improved available help on various pages.
- Enhanced mail-server administration.
- All pages have a printable version accessible from the icon at the top-right of the page (CONF-338)
- HTML and text email notifications have been greatly improved and are now much more useful. Links in HTML emails work seamlessly with online Confluence version.
- Many other minor fixes, improvements and performance tweaks

Release Notes 1.0b1

Confluence 1.0b1

Confluence 1.0b1 was released 19th December 2003.

That's right - we finally got it out the door! *cheers from around the room erupt as developers across in the Atlassian offices crack open a beer* .

New Features

- It's soooo much faster as we've implemented caching at multiple levels - trust me, it now *flies* in development. Here's hoping it flies for you too. (For those with a development bent, turn on profiling via the URL to see abeautiful profiling stack trace on stdout)
- The last login date is now tracked, pages and spaces changed since your last login are highlighted in green (as these are presumably pages you want to view), you can see a diff of the currently viewed page against the date you last logged in and there are various related reports of spaces/pages changed since your last login.
- Parent/child page hierarchies have been implemented. You can specify the parent for any pages and the ancestry is reflected in the breadcrumbs list. The full hierarchy is shown (including the current page's position) on the Page Information screen, and there is a

related macro to show the children in various different types of trees (see the Notation Guide for details)
- Search has been completely overhauled - search now works across pages, comments and space descriptions and the results now display descriptions with highlighted search terms (this is really cool!)
- Page renaming and link refactoring now works across comments, space descriptions and pages - neat!
- The Notation Guide has been significantly improved with more useful examples, a list of the available macros and a tabbed interface. Also the edit/create pages now have a 'quick notation guide' to the most commonly used operations
- There is a new 'Create Comment' permission - so that you can make spaces which are publicly viewable/commentable, but not editable.
- The UI of the 'create link' and 'external link' links within a rendered page have been improved with subtle icons
- Page comments now have a sexy icon of their own, and can be shown/hidden at will. You can also permalink to any given comment and comments are fully indexed.
- Undefined and orphaned pages reports have been tightened significantly so that they now report truly undefined/orphaned pages
- You can create a list of URLs ignored from external referrers - this is useful where your server is known by multiple IP / domain name combinations, or you want to ignore all internal referrers (Admin / General configuration)
- {color} macro to colour blocks of text
- Logging in to Confluence now takes you to your intended destination, not the Dashboard
- There is an option to choose whether the default link for a space goes to the space homepage or the space summary (people internally and externally seem to be divided about which it should go to - so you can choose for yourselves)
- You can now edit and add SMTP servers through the web interface
- The administration interface has been completely overhauled and tightened
- Added a Mail Queue administration page
- Added a View System Information page
- Added a pretty error page, and the ability to submit but reports (including exports) directly to us
- The permissions viewing and editing pages have been greatly simplified and improved across the application
- Confluence now has a licensing system, so you will need a license key to evaluate (these can be generated online as with JIRA).
- Image embedding can now include parameters (such as height, width, alignment etc) - see the Notation Guide for details.
- https:// links are now treated just like ftp:// and http:// external links
- Various rendering bugs have been fixed.
- Many other minor fixes, improvements and performance tweaks

Release Notes 1.0b2

I see dead pages!

Confluence 1.0b2 was released on Friday January 16, 2004 with a major bug in its search indexing that caused it to index pages that no longer existed. It was replaced by 1.0b3 the next Monday.

See also: Release Notes 1.0b3 and Issues Resolved for 1.0b3

Feature List

The features listed on this page refer to Confluence version 1.0 beta1.

Content Management

- Content is organised into discrete spaces in which users can create and edit pages.
- Easy-to-learn, easy-to-read but powerful notation for designing pages and linking between them, based on the popular textile markup.
- Arbitrary files can be attached to pages.
- Comments can be left on pages.
- Page templating allows rapid creation of boiler-plate pages.
- Pages can be organised into hierarchies.

Update Tracking

- Each page has a full change history, accessible as coloured diffs highlighting each change made.
- Simple Reports which pages have been updated since you last logged in.
- Email subscription informs you of pages that have been added or updated, or comments left.
- RSS feeds available for new and updated pages across the site, or in individual spaces.

Searching

- Full-text searching of all content, including pages, comments and space descriptions

Site Management

- Smart page-renaming updates links across the entire site, even in comments or space descriptions.
- Pages or spaces can be imported from plain text files, or exported to PDF or HTML.
- Simple backup and restore to XML files.
- Automatic daily backup feature.
- Reports of "orphaned" pages that exist but are not linked to, and of links that point to pages that do not yet exist.

User Management and Security

- Flexible security, from a public site, to a personal space.
- User- and group-based permissions.
- Separate viewing, commenting, editing and administration permissions for each space.

Deployment and Compatibility

- Available on a wide variety of platforms, either as a stand-alone server, or as a web-archive deployable into a compatible Java application server.
- Compatible with a large number of databases

Support

- Confluence is under full-time development, with licensees entitled to a year of free updates.
- Atlassian's Legendary Service.

Release Notes 1.0b3

Confluence 1.0b3

1.0b3 was released on January 19th, 2004. It replaces the short-lived 1.0b2 release, after 1.0b2 managed to escape into the wild with a bug in the search index that made it see pages that weren't there any more. Since 1.0b2 only existed for one weekend, the release-notes for both versions have been folded into a single document.

This is a massive release. Over 90 issues (new features, improvements and bug-fixes) were resolved between beta-1 and beta-3, with even more improvements made below the issue-tracking radar.

Contents

1. Unknown macro: {link}
2. Unknown macro: {link}
3. Unknown macro: {link}
4. Unknown macro: {link}
5. Unknown macro: {link}

See also: Issues Resolved for 1.0b3

New Features

Search Attachments

Attachments are now indexed for searching just like any other content in the Confluence space. Alongside text, HTML and XML attachments, Confluence will also index Word, RTF and PDF documents.

Configurable Look and Feel

The site administrators can now customise the colour-scheme of the Confluence installation. More advanced configuration can be done by editing the site's templates through a web interface, changing the appearance of the whole site.

Site Homepage

The site administrator can now set the site homepage for users who have not logged in: choosing from either the dashboard (the default behaviour in beta1 and before), or any of the Space homepages within the site.

Users who have logged in can choose their own site homepage in their user profile.

RSS Feed for New Comments

You can subscribe to all new comments added to a space using an RSS newsreader. (The link to the RSS feed can be found on the Space Summary page)

Emoticons



New Macros

- {include} – include the contents of one Confluence page within another
- {include-html} – include the contents of an HTML document within a Confluence page (turned off by default. See Enabling the html-include Macro for more information)
- {rss} – include an external RSS feed
- {search} – include the results of a Confluence search
- {jiraissues} – integrate Jira issue reports with your Confluence site
- {junitreport} – include JUnit test result data

Improvements

- Users' login names and full names are indexed for searching
- Users' profiles may now be longer than 255 characters, and are also indexed for searching
- Usernames are no longer case-sensitive
- Anonymous contributions are clearly labeled
- Users are warned if they are editing or commenting without having logged in
- Notification emails now link to the appropriate "diff" page, so you can quickly see what has changed

- Page diffs now highlight precisely what changed within each line: very useful when just one or two words change in a long paragraph
- Lists of child pages and links are sorted alphabetically
- Long lists of pages or search results are paginated
- Shortcut links can be given different link text in the same way as other links (e.g. [Search for Confluence on Google|confluence@google])
- Incoming links and "hot referrers" are listed in the sidebar of the page view
- Child pages are listed below the page contents in the page view

Notable Bug-fixes

- You no longer get the old page (or don't see the comment) immediately after editing or commenting on a page CONF-453
- Persistent login cookies no longer conflict with a JIRA installation on the same server CONF-440
- Persistent login cookies no longer fail for users with certain characters in their username CONF-387
- Notification emails no longer send out garbage for anonymous changes. CONF-421
- Users who are not logged in no longer see strange table titles CONF-422
- And many more....

Outstanding Issues

- You should restart Confluence immediately after finishing the initial setup steps, to avoid data loss CONF-493
- New-lines may not be drawn if the next line starts with whitespace CONF-475
- Emoticons are rendered inside {noformat} blocks CONF-502
- If you put a {children} macro after an {include} macro, it will list the children of the included page. CONF-504

Issues Resolved for 1.0b3

Errors were reported by the JIRA trusted connection.

- APP_UNKNOWN; Unknown Application: {0}; ["confluence:4557196"]

JIRA Issues (3 issues)				
Type	**Key**	**Summary**	**Status**	**Resolution**
	CONF-497	Search includes old versions of pages	Closed	Fixed
	CONF-495	Problem with incoming links	Closed	Fixed
	CONF-496	Problem with rename	Closed	Fixed

Issues resolved for 1.0b2

1.0b3 was a quick bug-fix release for 1.0b2, so here are the issues resolved in 1.0b2:

Errors were reported by the JIRA trusted connection.

- APP_UNKNOWN; Unknown Application: {0}; ["confluence:4557196"]

JIRA Issues (81 issues)				
Type	**Key**	**Summary**	**Status**	**Resolution**
	CONF-304	Sitemesh/Velocity Integration	Closed	Fixed
	CONF-425	Upgrade to CVS HEAD of XWork / WebWork 2	Closed	Fixed
	CONF-429	Upgrade SiteMesh and use Velocity decorators	Closed	Fixed
	CONF-410	Update confluence features list	Closed	Fixed
	CONF-438	JUnit XML displaying macro !	Closed	Fixed
	CONF-483	{include:page} macro	Closed	Fixed
	CONF-307	Ability to search attachments	Closed	Fixed
	CONF-437	External RSS macro!	Closed	Fixed
	CONF-456	Global colour-scheme configuration	Closed	Fixed
	CONF-418	Better handling of anonymous contributions	Closed	Fixed
	CONF-173	Remote editable space decorators	Closed	Fixed
	CONF-407	RSS feed for recently added comments	Closed	Fixed
	CONF-417	One-click bug submission	Closed	Duplicate

	CONF-420	Set "Site Homepage"	Closed	Fixed
	CONF-449	Macro for fetching/importing JIRA issues	Closed	Fixed
	CONF-290	Configurable Look & Feel	Closed	Fixed
	CONF-412	Make usernames and user full names searchable	Closed	Fixed
	CONF-458	JIRA macro column selection	Closed	Fixed
	CONF-462	Documentation for decorator editing	Closed	Fixed
	CONF-394	Update links from other spaces when renaming pages	Closed	Fixed
	CONF-431	Per word diffing	Closed	Fixed
	CONF-443	Shortcuts should allow "my link name" like other links do	Closed	Fixed
	CONF-457	Put nice colour-picker on colour customisation screen	Closed	Fixed
	CONF-401	Parent child relationships should be thought about more	Closed	Fixed
	CONF-371	Child pages invisible by default	Closed	Fixed
	CONF-413	Add diff link to "page edited" email	Closed	Fixed
	CONF-415	Make diffs highlight changes within a line.	Closed	Duplicate
	CONF-397	Fix user profile so it can hold > 255 chars	Closed	Fixed
	CONF-477	Add emoticons to confluence	Closed	Fixed
	CONF-471	Pagination for long lists (search is the first!)	Closed	Fixed
	CONF-419	Warn user if they're commenting/editing anonymously	Closed	Fixed
	CONF-373	Upgrade to the Spring SessionInView filter	Closed	Fixed
	CONF-454	Improve the JUNIT macro	Closed	Fixed
	CONF-388	Include user in "Recently Updated Pages"	Closed	Fixed
	CONF-387	Move "incoming links" back to the page	Closed	Fixed
	CONF-482	Page templates should be editable on the web interface and saved in exports	Closed	Fixed
	CONF-408	Escaped characters don't work as they should	Closed	Fixed
	CONF-450	BackupJob does not have a Hibernate session	Closed	Fixed
	CONF-392	Mistyped link syntax gets rendered weirdly	Closed	Fixed
	CONF-446	ampersand in links breaks them	Closed	Fixed
	CONF-467	"Display Default Decorator" shows edited template, not default	Closed	Fixed
	CONF-488	Too much white space generated?	Closed	Fixed
	CONF-424	{children} macro barfs on removed child page	Closed	Fixed
	CONF-414	Replace all webwork.urlEncode with generalUtil.urlEncode	Closed	Fixed
	CONF-487	HTML emails contain #emailUserLink (mike)	Closed	Fixed
	CONF-481	Newly added comment doesn't appear when posted	Closed	Duplicate
	CONF-463	Path admin page still thinks it's a setup step	Closed	Fixed
	CONF-468	Link extraction should exclude {code} contents	Closed	Fixed
	CONF-432	Exception when diffing added line	Closed	Fixed
	CONF-433	No security checking in FileServerServlet!	Closed	Fixed
	CONF-411	Cancel button not working on Add Comment Dialog	Closed	Fixed
	CONF-442	line breaks (\\) don't work	Closed	Fixed

	CONF-435	Macros still run inside {noformat} block	Closed	Fixed
	CONF-439	Email password doesn't seem to work	Closed	Fixed
	CONF-441	"my link name" links in tables don't work right	Closed	Fixed
	CONF-404	Undefined pages report misrenders when link source is a comment	Closed	Fixed
	CONF-405	Usernames (while logging in or setting up a new account) shouldn't be considered as case sensitive	Closed	Fixed
	CONF-384	Login cookie appears to be broken on confluence.atlassian.com:8080	Closed	Fixed
	CONF-452	Permissions check for /signup.action always fails	Closed	Fixed
	CONF-393	(*boldme* does not work.	Closed	Fixed
	CONF-485	Incoming links from space descriptions buggy	Closed	Fixed
	CONF-472	adding a comment does not display it right away	Resolved	Fixed
	CONF-494	Export from data originating in beta1 may be unimportable in beta2	Closed	Fixed
	CONF-440	Confluence Login cookies conflict with JIRA	Closed	Fixed
	CONF-451	Setting a homepage makes dashboard inaccessible.	Closed	Fixed
	CONF-470	Apostrophe double-encoded inside {code} block	Closed	Fixed
	CONF-476	Editing personal information in user profile stops working	Closed	Fixed
	CONF-423	Should not be able to add ANYONE group to ADMINISTRATE Conlfuence/Space permissions	Closed	Fixed
	CONF-453	"Add Page" permission checking is not consistent when logging in as anonymous	Closed	Fixed
	CONF-409	Logging in annonymously & the UI display keys!	Closed	Fixed
	CONF-466	{color} macro messed up by surrounding {{monospace}} markup	Closed	Fixed
	CONF-422	i18n text inside a webwork 2 component do not display (only their keys do) for anonymous users	Closed	Duplicate
	CONF-406	Could not parse : Dec 19, 2003 : java.text.ParseException	Closed	Fixed
	CONF-354	Can not use LDAPCredentialsProvider	Closed	Won't Fix
	CONF-421	NPE when sending notifcations for anonymous user activity	Closed	Fixed
	CONF-479	Errors invisible on "rename page" form	Closed	Fixed
	CONF-447	Error automatically creating pages with illegal names	Closed	Fixed
	CONF-383	Export space fails	Closed	Won't Fix
	CONF-486	Space descriptions can't contain links in their space	Closed	Fixed
	CONF-459	' is converted into &#8217	Closed	Cannot Reproduce
	CONF-464	Standalone tomcat version is not working!	Closed	Fixed

Release Notes 1.0b4

Confluence 1.0b4

To belatedly ring in the Year of the Code-Monkey, it's time for Confluence 1.0 beta 4. We've had some great feedback on the last few betas, thanks to everyone for submitting bugs, and contributing to the forums and discussion space.

Contents

1. Unknown macro: {link}
2. Unknown macro: {link}
3. Unknown macro: {link}
4. Unknown macro: {link}
5. Unknown macro: {link}

6. Unknown macro: {link}

See also: Issues Resolved for 1.0b4

Atlassian FatCow

Along with 1.0b4, we're also releasing FatCow: Functional Acceptance Testing for the Confluence Wiki. Styled after Ward Cunningham's FIT and Bob Martin's Fitnesse, FatCow allows you to define web-based acceptance tests in wiki notation, and then run them from inside Confluence. Here's a tutorial showing how to write a quick FatCow test suite that makes sure Confluence shows up on Google. 

FatCow is Open Source, and also serves as example code for anyone who wants to extend Confluence by writing their own macros (something that we'll be looking to make easier in future releases).

New Features

Blog Posts

Each space can now host a "blog" of pages that are organised by date, rather than just by title. This is pretty useful if you want to attach some kind of updating news to a space. The "create blog post" button can be found on the right-hand toolbar.

Blogging support is pretty basic right now: you can create posts, and you can include the most recent posts in a page using the {blog-posts} macro. Rest assured, we'll be piling on the features in the next few releases.

Move Pages Between Spaces

A much requested feature, our newest refactoring lets you move pages cleanly from one space to another. We're using this already to maintain a private space where we stage documentation waiting to be transferred to the main documentation space.

Configurable Site Description

You can now change the text in the site description that appears on the user's dashboard: somewhere to put welcome messages or MOTDs. And, of course, it understands Confluence markup.

#includePage("Page") Velocimacro

For the decorator-editors, you can use this render the contents of a page anywhere inside a decorator. The page has to be in the space that the user is looking at, and if the page does not exist, nothing will be rendered.

New Macros

- {blog-posts} displays the most recent blog posts for a space.
- {rss} macro now has maxEntries and titleBar parameters.
- {anchor} macro allows you to create named anchors in a page (link to them with [SPACE:page#anchor])

Improvements

- You can now draw en – and em — dashes.
- mailto: links are now drawn as just the email address, like so user@example.com
- the {search} macro now excludes the page it was included in from the search results
- you can also link to attachments using #-anchors [SPACE:page#attachment.pdf]
- you are given the opportunity to pick a template when creating a page from a link, and any entered page title survives picking a page template

Notable Bug-fixes

- Spurious error message about editing a stale version of a page have been squashed.
- No longer crashes when you add a user to certain groups.
- Some database queries have been rewritten to work around the fact that MySQL doesn't understand sub-selects.
- Diffs more reliably highlight changed words
- Several minor rendering problems to do with deeply nested lists have been fixed.
- and many more...

Outstanding Issues

- You should restart Confluence immediately after finishing the initial setup steps, to avoid data loss CONF-493
- New-lines may not be drawn if the next line starts with whitespace CONF-475
- Emoticons are rendered inside {noformat} blocks CONF-502

Issues Resolved for 1.0b4

Errors were reported by the JIRA trusted connection.

- APP_UNKNOWN; Unknown Application: {0}; ["confluence:4557196"]

JIRA Issues (39 issues)				
Type	**Key**	**Summary**	**Status**	**Resolution**

	CONF-460	FIT macros and integration	Closed	Fixed
	CONF-501	— and – filters	Closed	Fixed
	CONF-291	Ability to move a pages between spaces	Closed	Fixed
	CONF-192	Ability to link to attachments	Closed	Fixed
	CONF-234	Blog posts	Closed	Fixed
	CONF-509	Site description	Closed	Fixed
	CONF-531	Add Max Entries parameter to rss macro	Closed	Fixed
	CONF-534	Search macro should exclude page it's included on	Closed	Fixed
	CONF-554	Support emdash like Textile	Closed	Fixed
	CONF-544	add a titleBar=true|false option to rss macro	Closed	Fixed
	CONF-555	mailto links should look better	Closed	Fixed
	CONF-513	Form rules on 'Create Admin Account' during setup!	Closed	Fixed
	CONF-543	Preview & Confirm is always telling me the page is outdated	Closed	Fixed
	CONF-517	Back button "rollbacks" changes when previewing a page	Closed	Won't Fix
	CONF-505	Internal anchor links have the external link icon	Closed	Fixed
	CONF-537	HTML export is broken	Closed	Fixed
	CONF-528	jiraissues macro does not display due column	Closed	Fixed
	CONF-516	Cannot add user to additional groups	Closed	Fixed
	CONF-524	Edit Page Conflicts is buggy - often seems to detect conflicts which aren't there	Closed	Fixed
	CONF-519	Lists only work to 4 levels	Closed	Fixed
	CONF-504	{include} macro confuses {children} macro.	Closed	Fixed
	CONF-535	Removing a page gives exception on MySQL	Closed	Fixed
	CONF-510	'Browse Pages' does not show the recently modified pages (in green)...	Closed	Fixed
	CONF-508	Space description links don't take you to the space description	Closed	Fixed
	CONF-474	Don't get the chance to fill out variables	Closed	Fixed
	CONF-444	JDBC error accessing orphaned pages on mysql	Closed	Fixed
	CONF-529	template variables with an underscore in the name don't highlight properly	Closed	Fixed
	CONF-547	MySQL doesn't support sub-selects	Closed	Fixed
	CONF-526	Moving page doesn't break parent/child relationships	Closed	Fixed
	CONF-567	Home link should take users to the home page	Closed	Fixed
	CONF-445	JDBC error accessing undefined pages on mysql	Closed	Fixed
	CONF-523	Single word diff rendering is buggy	Closed	Fixed
	CONF-566	Previewing multiple times while editing confuses the versioning system	Closed	Duplicate
	CONF-546	Hot Referrers includes editing links	Closed	Won't Fix
	CONF-507	SnipSnap import does not add users to confluence-users	Closed	Fixed
	CONF-518	Deeper item in list can't be bold	Closed	Fixed
	CONF-515	Home link on summary page goes to... summary.	Closed	Fixed
	CONF-512	Confluence doesn't like it when I re-add a deleted child page	Closed	Fixed
	CONF-522	Page-Space relation is corrupted after latest upgrade (blog related table modifications)	Closed	Fixed

Release Notes 1.0rc1

Confluence 1.0rc1

It's Friday again, and that means it's time for us to release Murray; Confluence 1.0 Release Candidate 1.

RC1 marks the start of the feature-freeze as we plunge headfirst towards a stable 1.0 release. From now until 1.0-final we will be concentrating on fixing up the remaining bugs and polishing the interface. We are hoping that we won't need to send out a second release candidate, but we can't really make any promises until we've had a chance to go through the inevitable deluge of new issues that will follow this release.

Of course, this means that there are features you want to see in Confluence that won't make it into 1.0-final. Keep those ideas and suggestions coming. Confluence isn't going to stop at 1.0 – we just had to draw a line *somewhere* or we'd end up perpetually in beta. We have no plans to slow down development after the first major release. Your Confluence licence includes a year of upgrades, and if you think Confluence is way cool now, it's just going to get better.

As always, before upgrading an existing Confluence installation, be sure to back up your data.

As you can see below, we've been busy the last two weeks. A hard-earned thirst needs a good cold beer.

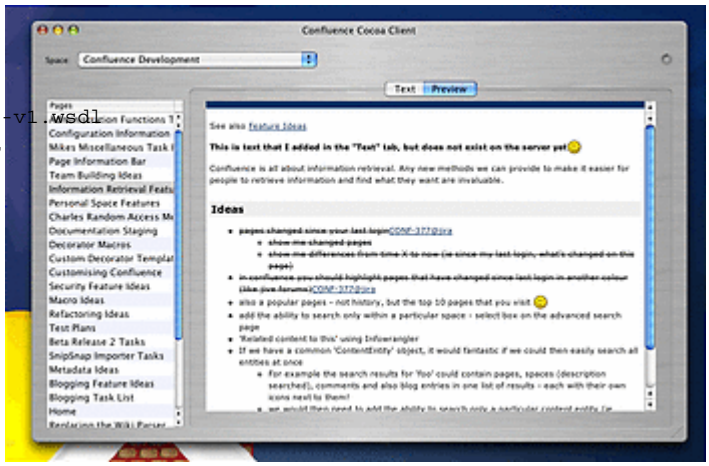
Contents

1. Unknown macro: {link}
2. Unknown macro: {link}
3. Unknown macro: {link}
4. Unknown macro: {link}
5. Unknown macro: {link}
6. Unknown macro: {link}

See also: Issues Resolved for 1.0rc1

New Features

Remote API



Confluence can now be browsed and edited remotely via SOAP or XML-RPC. There are more details of the XML-RPC API here, while the SOAP WSDL can be downloaded from `$your_confluence_root$/rpc/soap/confluenceservice-v1.wsdl`. Throwing together a GUI client for Confluence is now pretty easy, so feel free to embed Confluence support into your favourite editor, web browser or IDE.

Blogging Enhancements

The sketchy blogging support from Beta4 has been upgraded and overhauled. Enhancements include:

- The look and feel of blog pages has been greatly improved
- A "recent blog posts" list available from the Space Summary page
- Daily and monthly views for blog posts
- Blog posts can now be edited and deleted
- You can now link to blog posts using the internal link syntax. Blog posts are addressed by their posting-date and title, like so: [SPACEKEY:/2004/01/03/Blog Post Title]
- You can link to the daily views, too: [SPACEKEY:/2004/01/03/]
- RSS feeds are available for new blog posts, both globally (from the dashboard), and for each space (from the space summary page)

Also, anonymous users are no longer allowed to post blog entries. We couldn't really think of a use-case where anonymous blog-posts would be wanted, but if you can, let us know!

Configurable User Notification

With any knowledge-sharing application, it's important to keep informed as to what has changed. One way to do this is by subscribing to any of the RSS feeds offered by Confluence. Another is to have Confluence email you directly whenever there has been a change. Up 'til now, users have had to make do with the ad-hoc regular-expression based notifications that were put in for developers. WIth RC1, we now have a more fine-grained and user-friendly way to choose how you find out about changes to the site.

1. Each day, Confluence can email you a summary of any changes that have been made to the site in the past 24 hours. You can activate this feature in your user preferences.
2. You can subscribe to "watch" a page from the option in the right-hand operations menu. Whenever the page is modified, commented on, or a new attachment is added, you will receive an email. When you no longer want to watch the page, you can click the icon again to unsubscribe.
3. You can subscribe to watch an entire space from the option on the space summary page. This subscribes you to all pages in that space, and also notifies you when new pages are created.

Each user can now also choose whether or not they will be notified of changes they make themselves.

Trackbacks

While Confluence has always tracked links between pages within the Confluence site, it now has a way to track links to and from external sites: the Trackback API. In this way, a Confluence page can be informed when another site has mentioned it, and inform other sites that it has linked to them.

We have implemented the Trackback and Trackback auto-discovery APIs across pages and blog-posts. Trackback is a widely implemented API that allows web pages to notify each other of links. You can enable (or disable) trackbacks from the Confluence general administration page.

When Trackbacks are enabled, each Page and Blog Post within Confluence is set up to receive trackback pings, and contains the auto-discovery code to allow clients to automatically find out how to send those pings. Trackbacks that are received are listed in the right-hand sidebar of the page.

Also when Trackbacks are enabled, Confluence will perform auto-discovery on each outgoing link from a page to see if the destination is equipped to receive trackbacks, and send its own trackback ping.

Look forward to future versions of Confluence making *very* interesting use of this feature. 

Per-Space Look and Feel

The template-- and colour-scheme editing features that were available on a global basis in previous versions of Confluence can now be configured separately for each space, allowing you to apply a different look and feel to the various spaces within your Confluence site.

New Macros

- {panel} macro allows you to draw a shaded box containing some content

Improvements

- Import and Export should now be much faster, and use significantly less memory
- The user browser now has a search function, to make it possible to manage large numbers of users effectively
- Incoming links are no longer displayed on a page if the user is unable to view the page being linked from
- The {anchor}, {blog-posts}, {index} and {quote} macros are now documented properly.
- You can escape smileys by putting a backslash before their last character, to cause them not to be rendered as images :)
- You can now link to peoples' user profile pages with [~username] (e.g. Charles Miller)
- You can now link to space homepages with [SPACEKEY:] (e.g. [TEST:])
- mailto: links have a nifty icon: user@example.com
- The History popup now tracks your visits to blog posts, space summaries and user profile pages
- Recent changes listings on the dashboard, user profile and space summary pages now includes changes to all content, not just pages.
- New Emoticons! (well, icons really)

(y)	(n)	(i)	(/)	(x)	(!)
					

Notable Bug-fixes

- Added HTML headers to detect and prevent RSS, HTML-Include and FatCow macros from being made to loop in on themselves. CONF-525
- The {anchor} macro, and [DOCPRIV:null] links now work as advertised. CONF-616, CONF-605
- Fixed a divide-by-zero error in page diffs CONF-584
- Included `javax.transaction` libraries with release, to allow the WAR to run under Tomcat 5 CONF-613
- And many more rendering and stability fixes...

Outstanding Issues

- Still some problems with character encoding in page titles CONF-569
- HTML include macro interacts badly with other Radeox filters CONF-549

Database Changes from Beta4

New Tables

TRACKBACK and NOTIFICATION tables were introduced. These tables should be generated automatically when you first start RC1.

New column in the EXTRNLINKS table

- alter table EXTRNLINKS add column CONTENTTYPE varchar(255);

Constraint Change on LINKS Table

The "not null" constraint was removed from the DESTPAGETITLE column of the LINKS table. Consult your database documentation on how

to alter your database for this change. For example:

- *MySQL 3.23:* alter table LINKS modify DESTPAGETITLE VARCHAR(255);
- *PostgreSQL 7.3.2* alter table LINKS alter DESTPAGETITLE drop not null;

Issues Resolved for 1.0rc1

Errors were reported by the JIRA trusted connection.

- APP_UNKNOWN; Unknown Application: {0}; ["confluence:4557196"]

JIRA Issues (63 issues)				
Type	**Key**	**Summary**	**Status**	**Resolution**
	CONF-579	Daily Notification Report	Closed	Fixed
	CONF-540	Recent Blog Posts RSS feed	Closed	Fixed
	CONF-559	Add Trackback support	Closed	Fixed
	CONF-280	Implement VP Wiki API	Closed	Won't Fix
	CONF-545	create a panel macro	Closed	Fixed
	CONF-581	Notify me for this page	Closed	Fixed
	CONF-591	Remote XML-RPC API	Closed	Fixed
	CONF-615	Create SOAP API	Closed	Fixed
	CONF-541	Blog Post Daily View	Closed	Fixed
	CONF-465	Per-space colour schemes	Closed	Fixed
	CONF-538	Internal links to blog-posts	Closed	Fixed
	CONF-542	Blog Post Monthly View	Closed	Fixed
	CONF-580	Notify me for this space	Closed	Fixed
	CONF-553	Link directly to user profile	Closed	Fixed
	CONF-626	Recent Blog Posts page	Resolved	Fixed
	CONF-561	Space specific decorators	Closed	Fixed
	CONF-617	History popup now tracks viewing user info pages	Closed	Fixed
	CONF-568	documentation for {anchor} macro missing	Closed	Fixed
	CONF-601	I18n for execution threads that aren't triggered by web requests	Closed	Fixed
	CONF-562	"New file attached" notification email	Closed	Fixed
	CONF-558	nice icon for mailto: links	Closed	Fixed
	CONF-506	Remove blurb from {jiraissues} header.	Closed	Fixed
	CONF-557	Improve user browser	Closed	Fixed
	CONF-578	Display of comment section remembered by page	Closed	Fixed
	CONF-539	Edit/delete blog posts	Closed	Fixed
	CONF-576	Blog page L&F needs to be more blog like	Closed	Fixed
	CONF-583	Do not notify user of his own actions	Closed	Fixed
	CONF-625	"Recent Changes" lists should include chages to all content types	Closed	Fixed
	CONF-502	Need a way to escape smileys	Resolved	Fixed
	CONF-600	listpages.action should show a page count	Closed	Fixed
	CONF-564	Error deleting template	Closed	Cannot Reproduce
	CONF-624	Prevent trackback autodiscovery from downloading just anything	Closed	Fixed

	CONF-610	Pages with a parent are, by definition, not orphaned	Closed	Fixed
	CONF-563	Incoming links displayed for pages you can't see	Closed	Fixed
	CONF-588	Links in noformat macro broken	Closed	Fixed
	CONF-572	ImportExportException	Closed	Fixed
	CONF-570	Moving page doesn't update internal links	Closed	Fixed
	CONF-612	Bug on "Moving Page"	Closed	Fixed
	CONF-3286	CLONE -New line disappears	Closed	Cannot Reproduce
	CONF-622	Errors with new links	Closed	Fixed
	CONF-621	Stack trace editing issue	Closed	Fixed
	CONF-525	RSS macro can loop on itself	Closed	Fixed
	CONF-592	Changing parent page doesn't work	Closed	Fixed
	CONF-620	Can't edit personal profile	Closed	Fixed
	CONF-536	Renaming backup files is dangerous	Closed	Fixed
	CONF-574	missing velocity file for blog post preview	Closed	Fixed
	CONF-552	search macro output is not updated	Closed	Fixed
	CONF-582	exception when removing page	Closed	Cannot Reproduce
	CONF-475	New line disappears	Closed	Won't Fix
	CONF-609	Deny blog posting for anonymous users	Closed	Fixed
	CONF-623	children Hierarchy macro fails if page is assigned to heirarchy later	Closed	Fixed
	CONF-571	Links in PDF export broken (contain www.atlassiannull.com/blah)	Closed	Fixed
	CONF-618	Can not setup with hsqdl	Closed	Fixed
	CONF-616	Links containing &#xxxx; entities are malformed	Closed	Fixed
	CONF-613	javax.transaction not included in lib jars	Closed	Fixed
	CONF-605	Page anchors showing in links	Closed	Fixed
	CONF-602	No notification email is sent!	Closed	Fixed
	CONF-573	Insert Link is broken	Closed	Fixed
	CONF-585	faulty rendering of anchor links	Closed	Fixed
	CONF-584	page diff results in ArithmeticException	Closed	Fixed
	CONF-628	Can not choose custom backup path	Closed	Fixed
	CONF-597	Templates feature doesnt seem to work in B3	Closed	Duplicate
	CONF-590	When renaming a page name truncalates first letter of new link	Closed	Cannot Reproduce

Release Notes 1.0rc2

Confluence 1.0rc2

After a week of frantic bug-fixing, we have released Confluence 1.0 Release Candidate 2. Thanks everyone for reporting all the bugs they found in RC1. While we really hope you'll keep reporting anything you find that doesn't work, we're also hoping that there's a lot less to report now.

We have released RC2 because it's a *lot* more stable than RC1. We know a lot of you are already using Confluence in some kind of production capacity, and we don't want to keep you waiting for the important fixes that we put in this week.

Since we're in the middle of a feature freeze, there obviously aren't any new major features in this release, but there are a number of incremental improvements, and a *lot* of fixes.

But first, a...

WARNING

If you are upgrading from Confluence 1.0rc1, **do not** perform a backup of your data from within Confluence. There was a serious bug with RC1 that caused user permissions to be corrupted during data exports.

There have been no database changes between RC1 and RC2, so you should simply be able to run RC2 against your existing data without performing a backup/restore.

If you wish to back your data up safely before the upgrade, you will need to do so manually: shut down Confluence, make a copy of your Confluence home directory (as defined in `confluence-init.properties`, and if you are storing your data in something other than HSQL, use your database's native backup procedure.

Contents

1. Unknown macro: {link}
2. Unknown macro: {link}
3. Unknown macro: {link}

See also: Issues Resolved for 1.0rc2

Improvements

- Remote XML-RPC and SOAP APIs now have a `getPermissions()` method
- {noformat} macro accepts a title parameter
- Page view and create links are no longer displayed if the user does not have permission to view or create the destination page.
- When you create a page from viewing a previous page, you are given another chance to make that page the new page's parent.
- RSS and HTML include macros use HTML proxies if they are defined using the standard Java `http.proxyHost` and `http.proxyPort` system properties.
- Default session timeout is now 60 minutes
- Improvements to user browser and user profile page UI
- Shortcut link definitions are now backed up and restored

Notable Bug-fixes

Stability

We've done a lot of work tracking down the source of any exception and page error that has been reported to us, and fixing their causes. We've also made one or two improvements to the error reporting page, but we're hoping you won't see it nearly as often (if at all) any more.

Page Templates

The creation, editing and deletion of page templates should now be a lot more reliable and provide a smoother user experience.

Also

- You no longer have to restart Confluence after its initial setup CONF-493
- Performing a full data export no longer corrupts users' group membership data. CONF-645
- You can now link to profiles of users with an @ symbol in their usernames CONF-639
- Trying to create a page with illegal characters in its name no longer loses your page content on some browsers. CONF-713
- Trackback pings are now sent for URLs that are not surrounded by square brackets CONF-708
- And, of course, many more...

Outstanding Issues

The two major areas we still need to work on are the PDF export and the use of non-ASCII characters in pages (especially page titles). Handling of both are much better than they were a week ago, but there's still some work to do before they're completely reliable.

Issues Resolved for 1.0rc2

Errors were reported by the JIRA trusted connection.

- APP_UNKNOWN; Unknown Application: {0}; ["confluence:4557196"]

JIRA Issues (74 issues)

Type	Key	Summary	Status	Resolution
	CONF-589	Way to prevent a http URL from being rendered as a link	Closed	Fixed
	CONF-255	Please please please. Support external (LDAP) groups.	Closed	Won't Fix
	CONF-532	Add title parameter to noformat macro (and you can rename it at the same time to 'block')	Closed	Fixed

	CONF-642	Pages that have more than one version have different icon	Closed	Fixed
	CONF-668	Up default session timeout to 60 minutes	Closed	Fixed
	CONF-640	Fix user browser UI	Closed	Fixed
	CONF-641	Improve user profile UI	Closed	Fixed
	CONF-636	anchor links to local page anchors classed as incoming link	Closed	Fixed
	CONF-670	Javascript 'Make previous page into parent' link	Closed	Fixed
	CONF-700	Create a page from a sub page should implicit include the parent page in the Create Page Dialogue	Closed	Fixed
	CONF-646	Move "new blog post" away from "add child page"	Closed	Fixed
	CONF-697	Edit my profile page help information is inconsistent	Closed	Fixed
	CONF-654	Make double-encoding smarter	Closed	Fixed
	CONF-666	Remote APIs need a getPermissions() method	Closed	Fixed
	CONF-714	RSS and HTML include macros should use proxies if defined	Closed	Fixed
	CONF-596	User friendly options for exporting a space	Closed	Duplicate
	CONF-656	emoticons path is wrong	Closed	Fixed
	CONF-631	SOAP service can't get started	Closed	Fixed
	CONF-689	pdf export of page containing ndash (--) fails	Closed	Fixed
	CONF-690	italic text effect in link text broken	Closed	Fixed
	CONF-708	External links that are not bracketed "[ ]" are not pinged for trackbacks	Closed	Fixed
	CONF-659	Snipsnap Import Fails.	Closed	Fixed
	CONF-662	Login required to browse spaces	Closed	Won't Fix
	CONF-703	Unable to create template of same name, after deleting origional	Closed	Cannot Reproduce
	CONF-704	Template stops working once you've edited it once	Closed	Fixed
	CONF-673	isUserWatchingSpace throws null pointer sometimes	Closed	Fixed
	CONF-706	Recent updates lists gets page modifier wrong if anonymously edited.	Closed	Fixed
	CONF-652	RSS feeds throw NullPointerExceptions	Closed	Fixed
	CONF-644	Cannot backup data	Closed	Fixed

	CONF-658	Panel, first list item not recognized, and not closing on multiple list items	Closed	Fixed
	CONF-669	Exception clicking "New Blogs" link	Closed	Fixed
	CONF-681	Rename problems (CONF-496) persist in pages made with previous versions	Closed	Fixed
	CONF-685	Attempting to export a non-perfect page as PDF breaks	Closed	Fixed
	CONF-619	Upgrade from B3 to B4 gives InfrastructureException	Closed	Won't Fix
	CONF-493	Confluence not fully set up until first restart	Closed	Fixed
	CONF-484	Some links in the documentation site are "create new page" links	Closed	Fixed
	CONF-627	Link icons don't show up in PDFs	Closed	Fixed
	CONF-360	Edit Profile corrupts user record	Closed	Duplicate
	CONF-687	Comments to blog posts that appear on the dashboard recently updated list are broken	Closed	Fixed
	CONF-688	SOAP service has $Proxy name	Closed	Fixed
	CONF-691	Can't delete page templates that have been edited	Closed	Fixed
	CONF-696	Wrong type of date returned incall to getPage and getPageHistory	Closed	Fixed
	CONF-699	Renaming page and only changing the case gives error message	Closed	Fixed
	CONF-713	Creating a page with bad characters in the title loses page content	Closed	Fixed
	CONF-711	NPE in Global "spaces report"	Closed	Fixed
	CONF-716	SQL error removing user on Postgres	Closed	Fixed
	CONF-679	Blog RSS DTD gives 404	Closed	Fixed
	CONF-683	PDF export fails on {children} macro	Closed	Fixed
	CONF-633	Can't restore the extranet data locally	Closed	Won't Fix
	CONF-675	When running confluence under a non-default context, exporting PDF breaks	Closed	Duplicate
	CONF-648	Unable to set the name of a page using templates once a template is selected	Closed	Fixed
	CONF-695	Inconsistencies in remote soap api	Closed	Fixed
	CONF-677	Exporting a page as PDF with macros results in "null"	Closed	Fixed
	CONF-705	When creating a document with the template Java error occurs	Closed	Cannot Reproduce
	CONF-667	admin login doesn't work in Safari	Closed	Fixed

	CONF-645	groups disappearing for users	Closed	Fixed
	CONF-635	Can not view the users which have their username start with capital !	Closed	Fixed
	CONF-678	Exporting entire space gives ImportExportException	Closed	Fixed
	CONF-684	Clicking on Previous Version -> Version gives NPE	Closed	Fixed
	CONF-698	Anonymous users cannot access the RSS feeds. Throws a NPE	Closed	Duplicate
	CONF-639	[~username] link fails when username contains @	Closed	Fixed
	CONF-651	Viewing previos version diff "to previous" throws NullPointerExceptions	Closed	Fixed
	CONF-647	Templates feature doesnt seem to work in RC1	Closed	Fixed
	CONF-655	Upload attachment without specifying file should have nicer error	Closed	Fixed
	CONF-674	Notifications not updated when a page is deleted	Closed	Fixed
	CONF-672	NPE with GlobalRSSFeed	Closed	Fixed
	CONF-632	ClassCastException in BackupJob	Closed	Fixed
	CONF-634	Shortcut/Interwiki links are not backed up and restored!	Closed	Fixed
	CONF-702	Export to PDF doesn't render all text correctly	Closed	Fixed
	CONF-701	Export page to PDF, click on generated link for page that doesn't exists...generates a NPE	Closed	Duplicate
	CONF-657	Various crashes with space-less links in user profiles	Closed	Fixed
	CONF-349	Create page from template	Closed	Fixed
	CONF-686	Servlet context added for export causing problem with mod_jk	Closed	Fixed
	CONF-649	"Create Page" link displayed even when user can't create a page	Closed	Fixed

Release Notes 1.0rc5

Confluence 1.0rc5

It's time to release Confluence 1.0rc5 into the wild: it's the fourth Release Candidate since last Friday, and we're definitely getting close here. The last two release-candidates did not have their own release-notes, so these notes will also include things that were fixed for those versions.

Contents

1. New Features
2. New Macros
3. Improvements
4. Notable Bug-fixes
5. Outstanding Issues

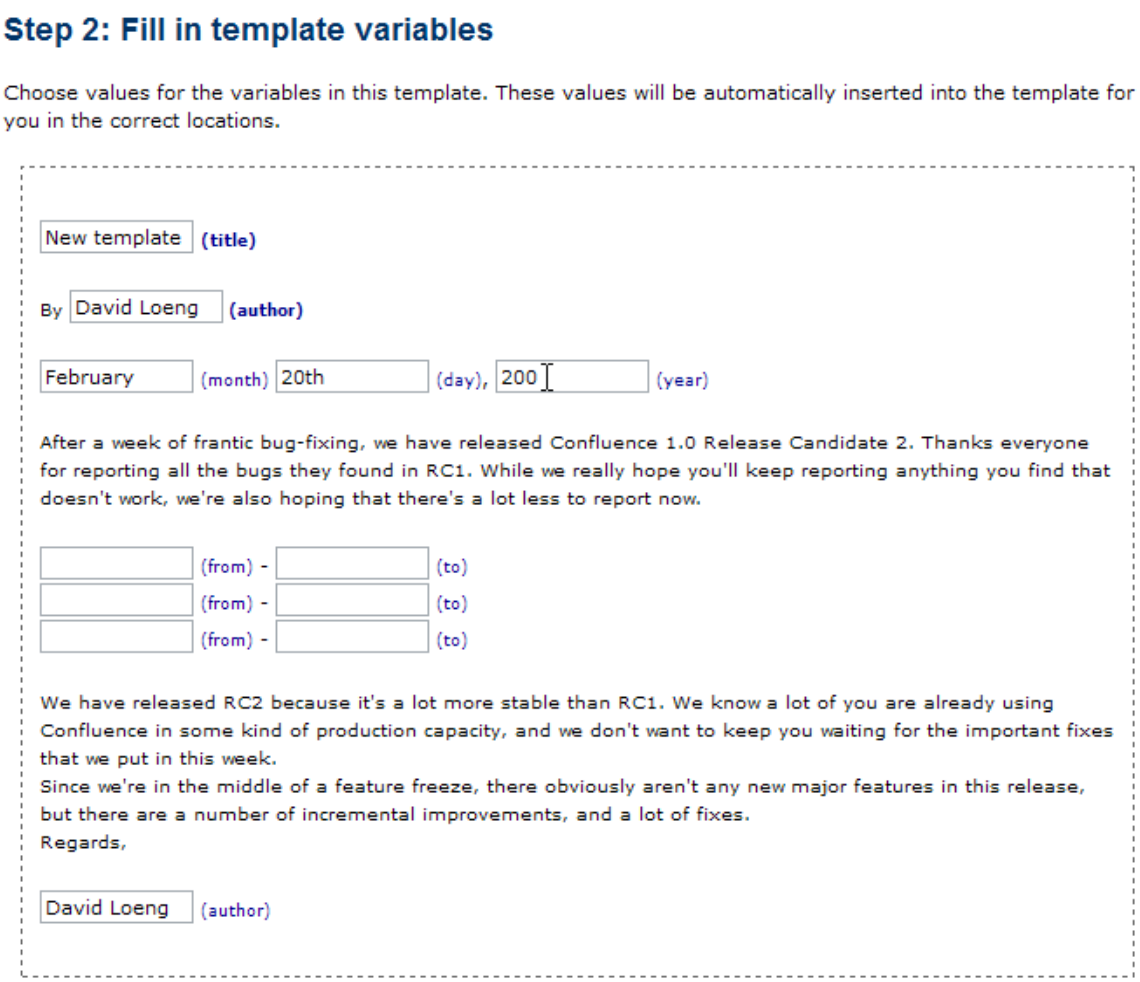
See also: Issues Resolved for 1.0rc5

New Features

In keeping with our promise to not add any features during the Release Candidate process, we have been very restrained, and only added two new features

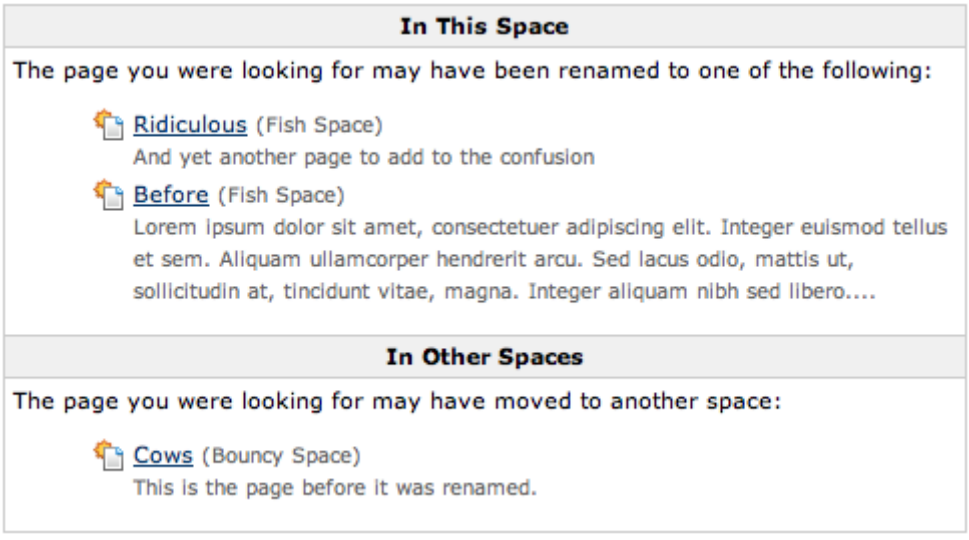
Template Preview

Page templates have been given an overhaul: when filling in your template variables, you are presented with a template of the rendered page, with input fields where the variables will be inserted. If you have the same variable used in several places, the form will take care of keeping them all in synch.



Page Redirection

When you rename or move a page, the page's old URL will attempt to give users some clue as to where they should be looking: checking which pages have had this name in the past, or redirecting users to pages in another space with the same name.



New Macros

- {weblogs} is a synonym for {blog-posts}

Improvements

- The user signup screen has a more friendly UI
- Exported PDFs now incorporate the site's stylesheet, and thus look a lot better.
- Confluence now has a useful "404" error page that gives the user some suggestions of how to find what they were looking for.
- An Administrators page, linked from the footer of each page, lists those users in the "confluence-admin" group. (If you don't want your address to be made public on this page, create an administrative group with some other name)
- The error you receive when you exceed your licensed user limit is more informative, and no longer directs users to email Atlassian. 
- The SnipSnap importer converts SnipSnap blog posts to Confluence blog posts.
- Several methods added to the XML-RPC and SOAP APIs for managing users and spaces. (See Conf Remote API Additions)
- The Daily Report email is more informative and more readable
- Page comments and Parent/child relationships are maintained during HTML and PDF exports
- news: and nntp: URL schemes are now recognised in bracketed links

Notable Bug-fixes

Since RC2

- Links between pages now work when you have a context path other than / (oops!)
- The daily report email is no longer sent every minute (oops!)
- You can no longer download attachments without being logged in (OOPS!)
- Search results no longer highlight words that were terms in previous searches

Since RC3

- Page redirects now work properly under Orion 2.0.2

Since RC4

- Removing a user no longer causes their notifications to hang around and crash things
- Removing a blog post works, too
- Many fixes to PDF rendering and exporting
- Many fixes to handling of Latin character set (See outstanding issues below)
- Invalid macros now give a more informative error than "no group 3"
- The 404 error page no longer requires you to log in to view it
- Macros that generate HTML (page include, HTML include, JIRA, RSS, FatCow) bypass the remainder of the page processing, and so should have far fewer formatting errors now.
- User browser filter remembers your search across pages

Plus, of course, innumerable fixed to annoying crashing bugs across all three releases.

Outstanding Issues

Latin Characters under Resin

We've tested creating pages with non-ASCII titles and content across several different browsers and several different operating systems, and they seem to be working reliably now... except on Resin. This will be quite noticeable since confluence.atlassian.com itself is running under Resin, but we have so far been unable to come up with a solution that works on this application server.

For users who need this functionality, we suggest running under Tomcat.

- On MySQL, the "orphaned pages" report may include the space's homepage CONF-766
- Combining block macros with lists is dangerous CONF-756
- Under some circumstances, paragraph tags will not be closed CONF-746

Issues Resolved for 1.0rc5

Errors were reported by the JIRA trusted connection.

- APP_UNKNOWN; Unknown Application: {0}; ["confluence:4557196"]

JIRA Issues (35 issues)				
Type	**Key**	**Summary**	**Status**	**Resolution**
	CONF-775	Add user management methods to remote API for administrators	Closed	Fixed
	CONF-787	Allow news: and nntp: URL schemes in [links]	Closed	Fixed
	CONF-788	Forced newline // should give <br clear="all"/>	Closed	Fixed

	CONF-750	Improve the UI of the signup screen	Closed	Fixed
	CONF-773	Make daily email a bit more readable	Closed	Fixed
	CONF-499	Renaming a page should leave behind an HTTP redirect	Closed	Fixed
	CONF-755	create a ConfluenceException and throw when we have a user correcatble error	Closed	Fixed
	CONF-723	The parent/child relationships are not reflected in HTML exports!	Closed	Fixed
	CONF-763	space scoped templates not showing up on Browse Templates page	Closed	Fixed
	CONF-719	PDF Export Bookmark browser doesn't reflex parent/child page relationship	Closed	Fixed
	CONF-781	NPE on 500 error response that contains no exception under Orion	Closed	Fixed
	CONF-751	Moving page edits space descriptions?	Closed	Fixed
	CONF-739	User browser filter not maintained in session	Closed	Fixed
	CONF-774	Daily Email not picking up all of the changes in a day	Closed	Fixed
	CONF-676	Exporting as a PDF, document with list items generates garbage in PDF	Closed	Fixed
	CONF-744	Exporting a page as PDF with {fatcontroller} breaks	Closed	Fixed
	CONF-772	putting in bogus URL doesnt show 404, later 404 looks unwell	Closed	Fixed
	CONF-767	html export zip appears empty as XP compressed folder	Closed	Won't Fix
	CONF-770	NPE in space look and feel action	Closed	Fixed
	CONF-769	Getting a lot of these traces running RC4 under tomcat 4.1.27	Closed	Won't Fix
	CONF-777	Remote API does not incidate space comment or admin perms	Closed	Fixed
	CONF-745	Panel, extra white space at top and extra extra when included	Closed	Fixed
	CONF-556	Import fails if export is from different database	Closed	Fixed
	CONF-758	content-by-user macro broken	Closed	Fixed
	CONF-760	Error message for an unknown macro is misleading	Closed	Fixed
	CONF-754	weblog macro broken in rc3 after a snipsnap import	Closed	Fixed
	CONF-782	Page doesn't get listed instantly	Closed	Cannot Reproduce
	CONF-761	Watch this Space didn't work	Closed	Fixed
	CONF-771	NPE on "recently updated pages report" on dashboard	Closed	Fixed

	CONF-757	PDF's show question marks instead of bulletin points	Closed	Duplicate
	CONF-779	Removing a user doesn't remove notifications	Closed	Fixed
	CONF-729	renderContent generates bad html for images from blog macro & for attached images	Closed	Fixed
	CONF-710	Can not export a page (as PDF) while there are two anchors with the same name!	Closed	Fixed
	CONF-768	Logging in occasionally redirects user straight to the fourohfour page :)	Closed	Fixed
	CONF-752	ImmutableException when trying to reset password	Closed	Cannot Reproduce

Release Notes 1.0rc6

Confluence 1.0rc6

Confluence 1.0rc6 fixes more bugs. Lots of them. There aren't many new and interesting features to report, although we have implemented the much-requested ability to mask email addresses and prevent public signup. Mostly, we've been toiling away fixing the lots of little problems that are getting in the way of us calling Confluence truly worthy of the 1.0 label.

Contents

1. New Features
2. New Macros
3. Improvements
4. Notable Bug-fixes
5. Outstanding Issues

See also: Issues Resolved for 1.0rc6

New Features

Private Confluence Configuration

In private mode, the only way to add users to Confluence is through the administrative interface: users can not sign up on their own. This way, you can prevent random passers-by from signing on to your Confluence installation. The site administrator can configure this in the general administration settings.

Email Address Privacy

Another much-requested feature, the site administrator has three options for email-address privacy:

- Public: just like pre-rc6, users 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.
- Private: only site administrators can see users' email addresses.

Administrators should be aware that even in private mode, anyone in the `confluence-admin` group will still be listed (with their email address) on the administrators page, although their addresses will be masked.

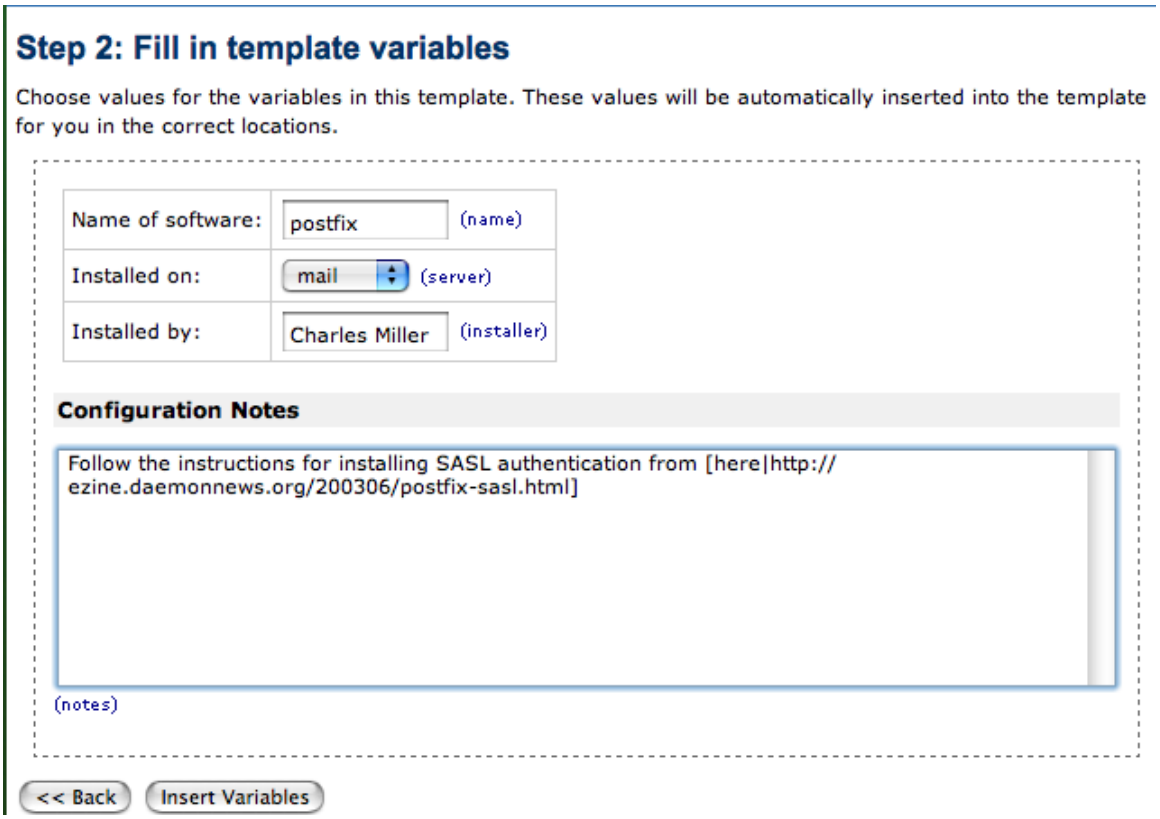
Share User Management with JIRA

For sites with both a Confluence and a JIRA installation, Confluence can be set up to delegate its user-management to JIRA.

Template Field Types

Templates now support drop-down menus and text-input areas.

`@variablename|textarea(5,10)@` will give you a 5 x 10 text-area called 'variablename'.
`@variablename|list(one,two,three,four)@` will give you a drop-down list called 'variablename', with options one, two, three and four.



New Macros

- The {html} will allow you to insert arbitrary HTML code into a page. This macro is turned off by default, as it is a security-risk on public systems. We only recommend you enable it on private or intranet-based Confluence installations where you trust your users not to write malicious HTML code. Here's a trivial example of its use:

```
1.\{html\}
2.<p>This HTML is <b>inserted</b> into the page</p>
3.\{html\}
```

- The {junitreport} macro has an option: "reportdetails=failuresonly" that will cause it to only show details of tests that have failed. For example: {junitreport:directory=file:///var/tests|reportdetail=failuresonly}

Improvements

- Users are now taken to their preferred homepage rather than the dashboard after logging in
- Text-fields in a template are disabled during preview and viewing
- Orphaned and Undefined page reports are now paginated
- Macros such as {code} or {noformat} are now made part of a list if they are placed adjacent to them

Notable Bug-fixes

- User->Group relationships should no longer disappear during manual or scheduled backups
- Search results no longer allow people to see content they might not be allowed to see
- Comments should no longer appear out of order on a page
- Moving a page between spaces no longer breaks #anchor links
- Many rendering fixes, some subtle, some not so
- Many stability fixes, especially regarding import and export

Outstanding Issues

Random MySQL Disconnections

If you are running Confluence against MySQL using Confluence's built-in datasource, the connection with the database server may be lost after long periods of inactivity. We are pretty sure we have a fix for this, but since the bug takes several hours to manifest, we were not able to test the fix before the release of rc6.

The good news is that if we are right, it will only require a quick edit of your configuration file to implement the fix. Further announcements will be made both on confluence.atlassian.com, and the confluence-user mailing-list.

In the meantime, the workaround is to not use Confluence's built-in datasource, but to configure Confluence to use your application-server's JNDI datasources instead. Instructions for doing this with Tomcat can be found here, and if you need more help, don't hesitate to contact us at confluence-support@atlassian.com.

Latin Characters under Resin

We've tested creating pages with non-ASCII titles and content across several different browsers and several different operating systems, and they seem to be working reliably now... except on Resin. This will be quite noticeable since confluence.atlassian.com itself is running under Resin, but we have so far been unable to come up with a solution that works on this application server.

For users who need this functionality, we suggest running under Tomcat.

Also..

- On MySQL, the "orphaned pages" report may include the space's homepage CONF-766
- You can create a link to a page with an illegal title: prompting the user to create a page that can not exist CONF-810
- Trackbacks are not sent for shortcut links CONF-888

Issues Resolved for 1.0rc6

Errors were reported by the JIRA trusted connection.

- APP_UNKNOWN; Unknown Application: {0}; ["confluence:4557196"]

JIRA Issues (52 issues)

Type	Key	Summary	Status	Resolution
	CONF-786	Document running JIRA and Confluence on one standalone	Closed	Fixed
	CONF-849	Delegate confluence user management to JIRA	Closed	Fixed
	CONF-813	Mask/hide email preference	Closed	Fixed
	CONF-792	Private setup	Closed	Fixed
	CONF-868	Inline HTML Macro	Closed	Fixed
	CONF-815	Disable textfields in template on viewing and preview	Closed	Fixed
	CONF-808	Undefined and Orphanned pages are not getting paginated!	Closed	Fixed
	CONF-812	Wording on signup page	Closed	Fixed
	CONF-837	Have Junit report show failures only	Closed	Fixed
	CONF-903	Some final really quick UI fixes?	Closed	Won't Fix
	CONF-905	It would be nice to be able to break up a list over several lines	Closed	Fixed
	CONF-789	Allow Template variables to have types	Closed	Fixed
	CONF-906	Users should be taken to their specified home page after successful login	Closed	Fixed
	CONF-756	Lists, and noformat blocks combination not working	Closed	Fixed
	CONF-912	Templates barf on anchors with no context	Closed	Fixed
	CONF-855	Change notification email links do not use full server path	Closed	Fixed
	CONF-882	Browse pages does not show newly added pages until a page is deleted	Closed	Fixed
	CONF-872	list items with russian symbols break list	Closed	Fixed
	CONF-909	MySQL disconnects if configured for direct connection	Closed	Fixed

	CONF-846	{noformat} renders spurious semicolon	Closed	Fixed
	CONF-860	Date / Time display in 12 hour format, without AM/PM	Closed	Fixed
	CONF-828	Search results include restricted spaces	Closed	Fixed
	CONF-804	Daily Email says all edits are by Anonymous	Closed	Fixed
	CONF-798	Inter-page links in exported PDF wrong	Closed	Fixed
	CONF-869	Crash on viewing own profile	Closed	Fixed
	CONF-870	Remove space fails with a SQL Integrity Contraint Violation	Closed	Fixed
	CONF-873	Need to restart tomcat if confluence unused for a while	Closed	Duplicate
	CONF-871	faulty rendering of sequenced text effects	Closed	Fixed
	CONF-820	hibernate lazy instantiation problem on getRecentlyUpdatedContent() in ViewSpaceAction	Closed	Cannot Reproduce
	CONF-836	Primary key conflicts after an import	Closed	Fixed
	CONF-838	Template of email notifications for 'text' format doesn't exist.	Closed	Fixed
	CONF-908	Input area to enter info in user profile is very small when using IE 6.0	Closed	Fixed
	CONF-854	ParseException when trying to use rss / jiraissues macro	Closed	Won't Fix
	CONF-857	Blank Space on Enclosed {panel}	Closed	Fixed
	CONF-746	HTML Paragraph not closed if starts with number colon	Closed	Fixed
	CONF-915	Bold not working inside a sub-numbered Bullet	Closed	Fixed
	CONF-875	Unexpected end of input stream	Closed	Duplicate
	CONF-839	Group bases permissions doesn't works properly	Closed	Fixed
	CONF-891	Comments appearing out-of-order	Closed	Fixed
	CONF-682	Single letter phrase notation doesn't work	Closed	Fixed
	CONF-549	html-include macro broken	Closed	Fixed
	CONF-863	Inline images get garbled when exporting to PDF	Closed	Won't Fix
	CONF-829	Licensing error appears on 'enter licence' screen	Closed	Fixed
	CONF-821	{blog-posts} macro gives weird black/grey/black/grey stripes	Closed	Fixed
	CONF-841	Moving page between space borks #anchor links	Closed	Fixed

	CONF-902	Blog post macro recursion problem	Closed	Fixed
	CONF-881	User/group relations vanish	Closed	Fixed
	CONF-874	Auto backup is not working	Closed	Cannot Reproduce
	CONF-894	RSS feeds	Closed	Cannot Reproduce
	CONF-826	Jira issues macro doesn't render correct title if it includes html tag	Closed	Fixed
	CONF-910	Image attributes allow malicious javascript	Closed	Fixed
	CONF-790	{noformat} inside bullets	Closed	Duplicate

Release Notes 1.1

Nymboida, Nym to his friends, is Confluence 1.1. He wants to be your friend.

Confluence 1.1 is the first major update to Confluence. It's faster, more reliable, and packed with new features. Thanks to our policy of a year's free upgrades, any current Confluence customer will be able to upgrade to 1.1 at no cost.

Current customers, or new users who wish to try out Confluence for 30 days can download either the standalone or WAR distributions from the Atlassian website: http://www.atlassian.com/software/confluence

Upgrading from 1.0.3a

Upgrading Confluence should be pretty easy: you can find instructions here. *We strongly recommend that you backup your `confluence.home` directory and database before upgrading!*

- Users who have enabled external user-management by customising their `osuser.xml` file will need to read this document also.
- Users who have used MySQL or Postgresql with Confluence 1.0 should read this document which explains how to get rid of any extraneous triggers or indexes that might have been created.

Contents

1. New Features
2. New Macros
3. Improvements
4. Notable Bug-fixes
5. Outstanding Issues

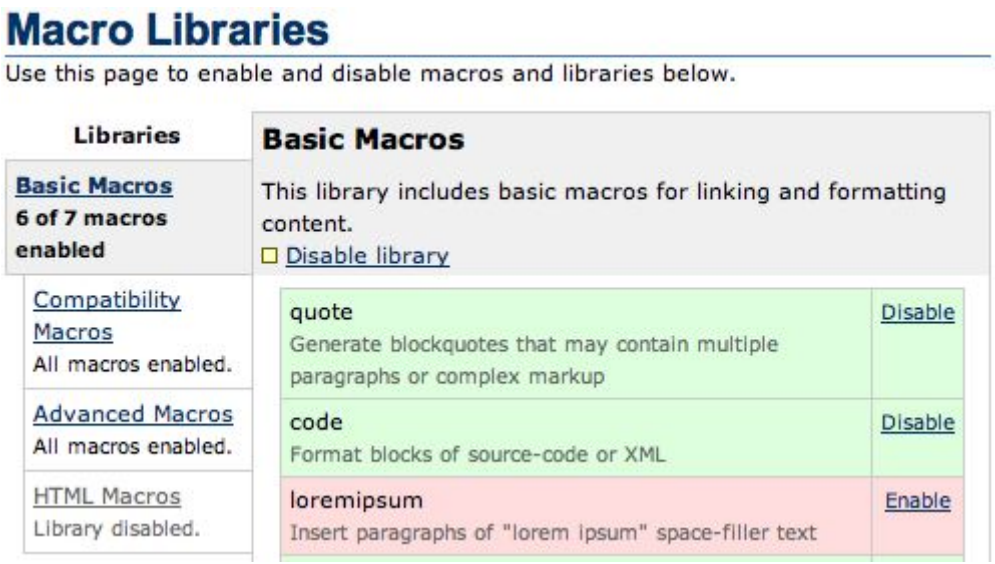
See also: Issues Resolved for 1.1

New Features

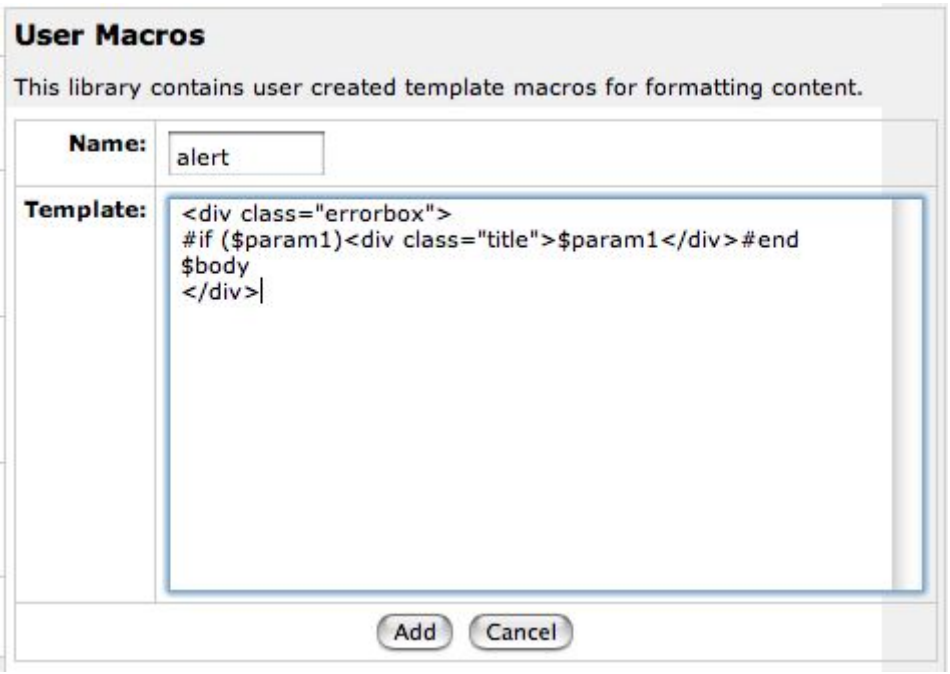
Macro Management

The biggest Nymboida new feature from an end user point of view is going to be the rewritten macro support.

Management - You can now enable and disable macros from this convenient (and very attractive) console. Macros are now grouped into libraries to make the management and creation of related macros simpler. Libraries can be installed, activated or deactivated as a single entity.



User Macros - We've also had a lot of requests from users for simple formatting macros: people who wanted their code snippets or notes to be formated in a certain way. Confluence 1.1 contains a simple way for administrators to create new macros from within the application without a restart: defining a template that the macro will apply to its arguments or content. User macros are very handy for providing consistent formatting and layouts across your pages.



Custom Macros - Installation of new macro libraries is a breeze: simply drop the library `.jar` file inside the web application, restart Confluence, and your new macros will appear in the management console. Confluence 1.1 also includes a custom macro toolkit (with Task List example shown below) to help users build their own complex macros easily using Java.



More information:

- Guide to Confluence Macros - a guide to the Macro Management console
- User Macros - how to write User Macros
- Custom Java Macros - obsolete - a guide (and worked example) for writing and distributing your own Java macros for Confluence

Attachments: versioning, comments and WebDAV support

Versioning - The number one feature request for Confluence 1.0 was versioned attachments. You ask, we deliver! Confluence can now have multiple versions of the same file attached to a page. It will keep a history of different versions of attachments, expandable dynamically.

Comments - Each attachment can now be accompanied with a comment describing why it is there, what it's about or the reason for it being attached. This is useful for tracking the differences between attachment versions, as well as for informing users as to why they should be interested in a particular file.



WebDAV - You can configure Confluence to store attachments in an external WebDAV server. This allows companies with an existing WebDAV infrastructure to give users alternative ways of accessing attachments and attaching files to Confluence pages.

Improvements to Page and Space Exports

- Exporting pages and spaces to PDF should be a great deal more reliable now. Whereas in Confluence 1.0, the export was likely to fail if the pages contained questionable markup, we now dynamically clean up each page before we export it.
- You can choose whether or not page comments are exported. Often you want to exclude comments if you want to send out a PDF of a page or publish a space as a website.
- The templates used to publish pages and spaces to HTML can be customised on a global and per-space basis. This makes it easy to publish a static website from the contents of a Confluence space: customise your templates, choose which pages to export, and voila.

Customise Display of Blog Posts

The templates used to display blog posts (both individually and in lists) can be edited either globally, or on a per-space basis, in the same way as the site's main decorators are editable. Also, the interface for editing templates has improved and will continue to do so in future releases.

Page Locking

If you want to prevent another user from editing, deleting or renaming a page, you can now lock it. You may want to do this because certain pages in an otherwise public space (for example, front pages, disclaimers or copyright notices) must remain un-edited, or just because you're working on the page and don't want anyone else to interfere just yet.

Locking a page restricts editing to a single user (yourself), or members of a particular group. Anyone with the space administrative permission can override or delete a lock.

Oracle and Weblogic Support

We have made a concerted attack on the various issues that were making Confluence unreliable on Oracle and Weblogic, and we're pleased to report that both have been running just fine in testing. Most importantly, the bug that was preventing pages being saved to Oracle when they were over 4Kb in length has been fixed.

For details of what precisely was fixed, see:

- The Weblogic support super-issue in JIRA
- The Oracle support super-issue in JIRA

Search Powerpoint and Excel Attachments

Confluence already searches across attached Word, PDF and RTF documents, XML, HTML and any plain text file. This search support has now been expanded in Nymboida to include searching and indexing of all text within Microsoft Powerpoint presentations and Excel spreadsheets.

In addition, the new attachment comments are also searchable.

TinyURL

Every page has a unique, short URL displayed at the top of its Page Information page. This makes it easier to send colleague's the URLs for pages with long titles via email, instant message or IRC.

New and Improved Macros

New

- {excerpt} allows you to mark a portion of the page as its "excerpt". This has no effect on the page itself, but other macros (such as {blog-posts} and {children}) can use the excerpt as a short summary of the content of the page.
- {excerpt-include} includes one page's excerpt in another page.

Improved

- {code} macro supports coloured highlighting for several more languages: JavaScript, ActionScript, XML and SQL.
- {code} macro can have its title and border customised in the same way as the {panel} macro
- {blog-posts} takes an optional `time` parameter to indicate how far back it should look for blog posts. For example, {blog-posts:time=7d} will show all blog-posts within the last seven days
- {blog-posts} takes an optional `content` parameter to change the way the blog-posts are displayed. `content=excerpts` displays excerpts instead of the full content of the blog entry (using the {excerpt} macro if available, otherwise extracting the first few hundred characters of the post). `content=titles` displays the entries as a list of titles.
- {children} takes an optional `excerpts=true` parameter: if any of the children have an excerpt available, the first line will be displayed in the list.

Improvements

Improvements to the Markup Engine

Many improvements have been made to the Confluence markup parser, fixing niggling inconsistencies, and allowing many more combinations of effects. If you want to produce something like the following, you can:

Name	Value

Some code	`1.public static void main(String[] args)`
	`2.{`
	`3.  System.out.println("Hello World");`
	`4.}`

A list	• Item 1 • Item 2 • Item 3 • Item 4
A panel	**I like cheese** • Brie • Camembert • Gruyere • Cheddar

You can also center an image using !image.gif|align=center!

Improvements to Linking

- You can specify a link title (which appears in the mouse-over tooltip) by adding another section to the link: [Link Body Text|Page Name|Link Title]
- [//foo/bar/baz.html] will create a relative URL link to `/foo/bar/baz.html` on the same server as Confluence is running
- `file:` URL links work
- UNC-style links: [\\SERVER\share\directory\file.doc] will create a `file://` link to the file on an external share

Note: Some web browsers (specifically Mozilla) consider `file:` URLs to be a security hazard, and do not follow them.

And a whole lot more...

Here's the quick fire version of some of the other improvements we've made in this release:

- You can resize the recent changes lists on the dashboard and space summary pages. This change is persisted in your user preferences.
- UI state (whether comments or attachments are open, size of recent changes lists) persists between sessions for registered users.
- The maximum attachment size can be configured from the General Configuration administration page.
- Going to http://yoursite.example.com/display now brings up a list of spaces instead of an error page.
- The display of the space summary page is significantly faster.
- Newly created users don't show up on the dashboard recent changes list unless they edit their profiles.
- The word "Confluence" has been moved to the end of page titles instead of the beginning, making them easier to distinguish in tabs

and bookmarks.
- There are more ways to navigate to the "recent blog posts" page for a space.
- Removing a user is significantly faster.
- Headings in pages are automatically turned into anchors with the same name.
- Creating a space now creates an index page as well as a home-page.

Notable Bug-fixes

It's hard to know where to start. We've fixed a lot of bugs across the whole application. If you want to know what's been fixed, you're probably best off looking for yourself.

Two areas, however, have been improved enough to deserve special mention.

- **PDF Export** - as mentioned above, we've made the PDF export much, much more reliable than it once was. Where before a page or space may have confused the PDF converter into not working, it should now be able to handle any markup you throw up at it.
- **International characters** - Many issues related to the use of non-ASCII characters in page titles, links, page contents and RSS feeds have been resolved since Confluence 1.0. Our users in non-English-speaking countries should find Confluence a much more pleasant and seamless experience now than they may have before.

Issues Resolved for 1.1

Error formatting macro: jiraissues: java.lang.IllegalArgumentException: You are not allowed to get a result set of more than 200 results. Current search returns 208 results

Release Notes 1.1.1

Confluence 1.1.1 is a maintenance release that fixes some bugs regarding attachments, page links and notifications. Remember, a Confluence license entitles you to a year of upgrades, and this upgrade along with future updates will be free of charge.

Who should upgrade?

This release mainly fixes the bugs discovered in our recent 1.1 release. The major issues resolved are listed below, or you can see the full list here. As some of the issues resolved have a significant impact on the user experience of Confluence, we recommend that anyone currently running 1.1 upgrade to 1.1.1.

Upgrade Procedure

1. Shut down the Confluence server
2. Back up `confluence/WEB-INF/classes/confluence-init.properties`, and if you have customised it, `confluence/WEB-INF/classes/osuser.xml`
3. Unpack Confluence 1.1.1 in the same location as your existing Confluence installation
4. Restore the two files you backed up in step 1 to the 1.1.1 installation.
5. Start Confluence

To avoid the possibility of data-loss, you should back up your ConfluenceHome directory and your database before upgrading, and perform a full backup from within the application.

Changes in 1.1.1

Attachment fixes

- Attachment versioning introduced a new table in Confluence that had a column whose name conflicted with some databases like Sybase
- Links to older version of attachments could not be accessed if the webapp was deployed with a context path
- Attaching files that contained '+' and '&' symbols caused certain pages in Confluence to break

These have been fixed in 1.1.1.

Notification fixes

Users opting to receive html formatted email received emails with the correct subject but no content. In addition, the "View Changes" link in the daily change email was broken. These are now fixed.

Email Server fixes

An upgrade to an email component used in Confluence in 1.1, caused an error to be displayed when users attempted to modify their email server settings. An upgrade has been added in 1.1.1 that will fix this problem to save users from having to perform the fix manually.

Export to PDF fix

Spaces labelled with names containing an '&' symbol could not be exported to PDF. This has been fixed.

Wiki Notation fix

The notation for a horizontal ruler was changed to five dashes (up from 4) in 1.1. This stopped the ruler from showing. This has been changed back in 1.1.1.

Issues Resolved for 1.1.1

Errors were reported by the JIRA trusted connection.

- APP_UNKNOWN; Unknown Application: {0}; ["confluence:4557196"]

JIRA Issues (25 issues)				
Type	**Key**	**Summary**	**Status**	**Resolution**
	CONF-1510	Mail Queued problem	Closed	Fixed
	CONF-1395	Upgrade {children} macro to optionally specify a page	Closed	Fixed
	CONF-1391	Old versions of pages appear when you reindex	Closed	Fixed
	CONF-1388	.sh files in the standalone release aren't executable	Closed	Fixed
	CONF-1387	URL for the non-tip version of an attachment is incorrect	Closed	Duplicate
	CONF-1386	Problem afer Upgrade from 1.0.3 to 1.1	Closed	Fixed
	CONF-1384	Redirect sometimes faster than previous txn commit	Closed	Fixed
	CONF-1382	Notifications are missing content	Closed	Fixed
	CONF-1380	Crash at creating new group	Closed	Fixed
	CONF-1377	SEVERE Ognl exception in server logs	Closed	Fixed
	CONF-1372	Front page should have greyed out icons instead of missing icons	Closed	Fixed
	CONF-1368	Links to older versions of attachments don't include the webapp context	Closed	Fixed
	CONF-1365	Wrong interaction between include and children macros	Closed	Fixed
	CONF-1362	Restore just spins when file access is denied and attachment directories cannot be restored	Closed	Fixed
	CONF-1361	Using Ampersand (&) in Space names breaks PDF export	Closed	Fixed
	CONF-1358	"VERSION" is a reserved word in a lot of databases (Attachment table	Closed	Fixed
	CONF-1357	Incorrect title is passed when page is created from the link	Closed	Fixed
	CONF-1356	Reset Default Colour Scheme	Closed	Fixed
	CONF-1353	Links to pages with dashes in title to not render correctly	Closed	Fixed
	CONF-1352	Attachments with plus sign in file name cause an NPE when loaded	Closed	Fixed
	CONF-1349	Links to attachments includes the ^ symbol	Closed	Fixed
	CONF-1346	No such setting:webwork.ui.templateDir - When confluence is root web app	Closed	Fixed
	CONF-1345	Moving page with linked attachment	Closed	Fixed
	CONF-1344	Confluence 1.1 Installation wizard - step 7 System Error		Fixed

			Closed	
	CONF-1343	WAR release has wrong readme.txt etc	Closed	Fixed

Release Notes 1.1.2

This fixes an upgrade problem in 1.1.1 where users attemtping to upgrade from 1.0.x directly to 1.1.1 encountered an "Upgrade Failed" error. Therefore, if you are still using a version older than 1.1, and am planning to upgrade, please upgrade to 1.1.2 to avoid this problem.

If you have already upgraded to 1.1.1 from 1.0.3a and cannot start Confluence due to the upgrade error, simply download 1.1.2 and upgrade to it. This should fix the problem.

If you have already upgraded successfully to 1.1.1 from 1.1 then you may safely ignore this upgrade.

Useful tips when upgrading from 1.0.3a

Upgrading Confluence should be pretty easy: you can find instructions here. *We strongly recommend that you backup your `confluence.home` directory and database before upgrading!*

- Users who have enabled external user-management by customising their `osuser.xml` file will need to read this document also.
- Users who have used MySQL or Postgresql with Confluence 1.0 should read this document which explains how to get rid of any extraneous triggers or indexes that might have been created.

Release Notes 1.2

Atlassian is proud to present Confluence 1.2, otherwise known as Swan, to the world. Existing customers who wish to upgrade, or new users who wish to try out Confluence for 30 days can download either the standalone or WAR distributions from the Atlassian website: http://www.atlassian.com/software/confluence

Swan is the second major update to Confluence, and once again all existing customers can upgrade for free, thanks to the provision for one year of free upgrades in your license. You will find that Swan contains significant new features in the areas of user management, search and space browsing, as well as the usual raft of enhancements, bug fixes and things we just couldn't resist throwing in at the last minute.

In all, 109 issues were resolved between 1.1.2 and 1.2. You can see the full list here: Issues Resolved for 1.2

Looking towards the future, we are going to attempt to increase the frequency of releases. While having one release every three months makes for impressively long release notes, it also means that customers are often left waiting longer than might be necessary for important enhancements or bug-fixes.

Upgrading from 1.1.2

Upgrading Confluence should be pretty easy: you can find instructions here. *We strongly recommend that you backup your `confluence.home` directory and database before upgrading!*

You will need to reindex your site after upgrading to enable some of the new site search features. You can do this from the "Rebuild Search Index" link on the Administration Console.

Site administrators upgrading from 1.1.2 or earlier should take particular note of the changes to global and space permissions. Confluence should automatically upgrade old permissions to the new scheme when upgrading from or importing data from earlier versions, but it is a good idea to check that your spaces are still secure and accessible after the upgrade, just in case.

Upgrading from 1.0.3a

Users upgrading directly from 1.0.3a should also read the Nymboida Release Notes for caveats regarding the 1.0 -> 1.1 upgrade.

Contents

1. New Features
2. Improvements
3. Notable Bug-fixes

See also: Issues Resolved for 1.1

New Features

Page List Views



The old "list pages" screen has been enhanced, almost to the point of being unrecogniseable. In its place, we now offer three ways to browse

the contents of a space.

The **Alphabetical View** looks like the "list pages" screen from 1.1 and earlier, but it has been enhanced to allow you to find pages quickly by the first letter of their title.

The **Directory View** presents the contents of a space as a tree, allowing you to drill down the hierarchy of parent and child pages (You will need Javascript turned on to use this view).

The **Search View** presents Swan's enhanced search interface, already focused on the space you are looking at.

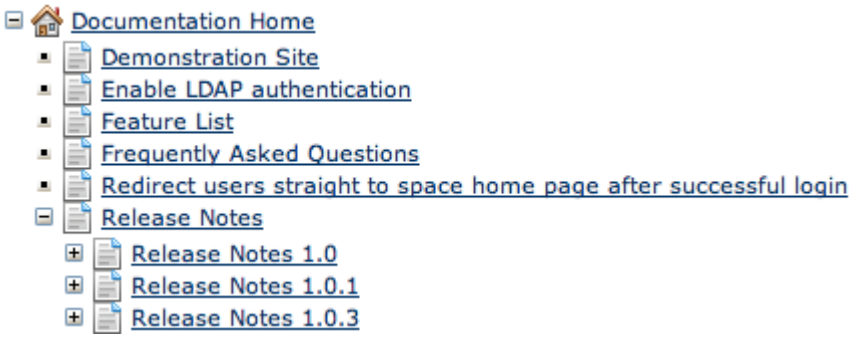


Image Thumbnails and Thumbnail Galleries

You can now display a thumbnail of any image attached to a page. Clicking on the thumbnail will pop up a window containing the full-sized image. For example, the nice picture of the waterfall from my holiday in California was generated by the following wiki-markup.

```
1.!waterfall.jpg|align=right thumbnail!
```



In addition, the {gallery} macro will create a thumbnail gallery of *all* the images that are attached to the page the macro is included in. You can see an example of this macro in action here: Thumbnail and Gallery Example.

The maximum sizes for thumbnails are configured in the General Configuration section of the Administrative Console.

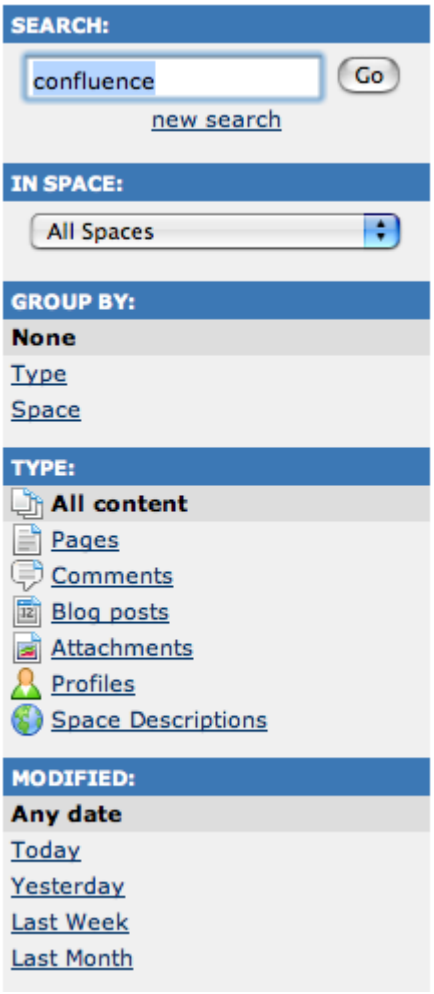
Threaded Comments

Bowing to public demand, we have added the ability to arrange comments by thread: users may respond to a particular comment, and the comments will be arranged in accordance with who responded to what.

The default is still for comments to be presented as a flat list, but you can enable comment-threading in the General Configuration section of the Administrative Console.

Improvements

Enhanced Search



The site-search functionality of Confluence has been enhanced, both subtly behind the scenes, and quite obviously in front of them.

Behind the scenes, we now index more information such as attachment comments and filenames. We have improved the indexing of Powerpoint presentations so that more text is extracted from the Powerpoint file. We have also tweaked the ranking algorithms behind the search so that you are more likely to find the page (space, comment, blog post, attachment, user...) you are looking for.

More obviously, we have given the user a lot more control over searching with an intuitive user interface that allows you to limit searches by space, type of content, or date last modified, and to group search results by type and space.

The {search} macro has also been updated to take advantage of the new search features (the options are described more fully in the notation guide linked from Confluence's edit page):

```
1.{search:query=Confluence|maxLimit=5|spacekey=DISC|type=page}
```

New Permissions Interface



Two of the biggest issues users have been having with permissions in Confluence were:

- The interface for setting permissions was confusing.
- It was hard to keep sites that allowed anonymous access below the workgroup license 25-user limit, especially if they used external user management.

Both of these issues have been fixed in Swan, but it has involved a slight change in the way permissions are checked. This should not be a problem for new Confluence users: the new system is easier to use than the old,. However, if you are upgrading from 1.1.2 or earlier, you should read this document to see what's different.

And a whole lot more...

Here's the quick fire version of some of the other improvements we've made in this release:

- The "double-click to edit" feature has been removed.
- Blog views now come with a monthly calendar showing nearby posts
- Macro libraries can now include images, or any other content the macro needs to serve over the web: look out for new example code in the [macro documentation] soon.
- Users can set their profile page to be their site homepage
- You can delete all referrer links that match your exclusion list (in General Configuration), to rid the site of referrer spam
- The site's administrative, search and editing pages are served with robots exclusion meta tags so that only the site's main content will be indexed by well-behaved search engines
- Restoring backups should be significantly faster for installations with many users
- Administrators can choose to restore backups from the filesystem instead of uploading them (important if your backup is quite large)
- Confluence now serves files with a proper Content-Length header, so browsers can display the download accurately
- CSS stylesheets are hidden in emails so they don't mess up Lotus Notes
- Export and backup filenames use a neater yyyymmdd date format

Notable Bug-fixes

- Exports created on Windows can now be imported on Unix-like operating systems without having to fix the path separators
- Hyphens in page headings or anchors will no longer render strangely
- Restoring a backup during setup now properly generates the search indexes
- Page "short links" now respect the configured site URL, regardless of which URL the user is accessing the site from
- Strange Powerpoint files no longer choke the indexer
- Users with Admin privileges can now see the link to the Administration Console without being in the 'confluence-administrators' superuser group
- Space administrators can modify space templates without having global administrator privileges.
- User macros are now lower-cased by default
- It is now possible to change the case of a page title by renaming the page

Issues Resolved for 1.2

Errors were reported by the JIRA trusted connection.

- APP_UNKNOWN; Unknown Application: {0}; ["confluence:4557196"]

JIRA Issues (117 issues)				
Type	**Key**	**Summary**	**Status**	**Resolution**
	CONF-1264	BOB: Run functional tests against Weblogic	Closed	Fixed
	CONF-1442	Use atlassian core's thread appender instead of confluence's	Closed	Fixed
	CONF-1393	Upgrade libraries	Closed	Fixed
	CONF-1130	Allow import/restore from server local filesystem	Closed	Fixed
	CONF-785	Gallery Page Macro	Closed	Fixed
	CONF-784	Inline image attachment as thumbnail	Closed	Fixed
	CONF-765	Calendar for blog posts	Closed	Fixed
	CONF-1497	Option to scope searching	Closed	Duplicate
	CONF-1539	"View in hierarchy" link from a page	Closed	Fixed
	CONF-1695	Improve threaded comments L&F	Closed	Fixed
	CONF-1668	Improve Administration Setup Paths page	Closed	Fixed
	CONF-1529	Incorrect number of licensed users displayed	Closed	Fixed
	CONF-1490	Add "Hide Comments" link when comments are shown	Closed	Fixed
	CONF-817	Revisit UI for adding permissions	Closed	Fixed
	CONF-1676	Email template fix for Lotus Notes 6.5	Closed	Fixed
	CONF-1477	Change backup filenames to 2004-06-29 format	Closed	Fixed
	CONF-1658	Zipped exports created on a Windows Confluence instance are not platform-independent	Closed	Fixed

	CONF-1163	Put NOINDEX NOARCHIVE tags on adminstrative and search actions	Closed	Fixed
	CONF-665	Warn if adding a space permission to user without USECONFLUENCE permission	Closed	Fixed
	CONF-886	Hierarchy view as proper treeview	Closed	Fixed
	CONF-809	Scope-base searches (on-site and via {search} macro)!	Closed	Fixed
	CONF-1196	Add App Specific Exceptions to Remote API	Closed	Fixed
	CONF-1239	Ability to clear Referrers	Closed	Fixed
	CONF-1408	Restoration process take a long time updating indexes	Closed	Fixed
	CONF-1376	Allow macro libraries to include web-served resources	Closed	Fixed
	CONF-1487	Alphabetical, Directory and Search views for page listing	Closed	Fixed
	CONF-1491	Double-click to edit page feature unintuitive	Closed	Duplicate
	CONF-957	Quick Search should notice if you've typed in a page title	Closed	Fixed
	CONF-985	Threaded Comments	Closed	Fixed
	CONF-480	Alphabet links in paged search results / page links	Closed	Fixed
	CONF-1260	DHTML spaces control	Closed	Duplicate
	CONF-961	Add "my profile" to the list of a user's available homepages	Closed	Fixed
	CONF-1359	Improve search	Closed	Fixed
	CONF-1447	Index attachment names	Closed	Fixed
	CONF-1588	FileServerServlet doesn't serve file sizes	Closed	Fixed
	CONF-1564	ClobStringType requires active transaction synchronization	Closed	Fixed
	CONF-1682	Space keys are alphanumeric, not ASCII	Closed	Fixed
	CONF-1692	Update Page fails in MailNotificationQueueItem.createFromTemplateFile	Closed	Duplicate
	CONF-1617	Couldn't restore directory from backup error	Closed	Fixed
	CONF-1486	i18n title shows up as hashes when exporting to PDF	Closed	Fixed
	CONF-1392	Import on initial setup doesn't index	Closed	Fixed
	CONF-1524	Hyphen in Blog text causes corruption.	Closed	Fixed
	CONF-1423	View mail servers doesn't show From address, but Edit operation does.	Closed	Fixed
	CONF-1501	Rebuilding Search Index take forever	Closed	Fixed
	CONF-1612	Brackets "()" break rendering of headings	Closed	Fixed
	CONF-1508	Very long headings render bad anchor tag	Closed	Fixed
	CONF-1292	Mail queue not updated with queued notification items	Closed	Fixed
	CONF-1536	template.not.found - i8n message missing	Closed	Fixed
	CONF-1548	User History page with velocity bug?	Closed	Fixed
	CONF-1419	Dates in a heading are not rendered correctly	Closed	Fixed
	CONF-1045	Exports performed on Windows may have \ as path separator in zip	Closed	Duplicate
	CONF-1659	Anchor links don't work in exported PDF's	Closed	Fixed
	CONF-1666	Deleting a group should delete that group's permissions	Closed	Fixed
	CONF-1426	Please remove the double-click-starts-editing feature	Closed	Fixed
	CONF-1471	Broken links in 'Confluence Notation Guide'	Closed	Fixed
	CONF-1430	Children macro choking on dashes in page names	Closed	Fixed

	CONF-1701	Gallery macro prints $thumb.attachment.comment under all thumbnails	Closed	Fixed
	CONF-1696	Must-fixes for search	Closed	Fixed
	CONF-1690	Bad URL for "edit space homepage" in space created screen	Closed	Fixed
	CONF-1121	Changing a word in a page shows wrong diff	Closed	Fixed
	CONF-1684	Attachment search fails when limited by space	Closed	Fixed
	CONF-1582	Links not rendered correctly using Remote render	Closed	Fixed
	CONF-1440	Confluence still doesn't clean the temp directory	Closed	Won't Fix
	CONF-1454	h3 element rendering junk	Closed	Fixed
	CONF-1441	Special Caracters in headers do not render properly	Closed	Fixed
	CONF-1407	Headers with Certain Characters	Closed	Fixed
	CONF-1599	Children macro chokes on minus characters	Closed	Fixed
	CONF-1630	View space permission description wrongly says that a user can edit a page	Closed	Fixed
	CONF-1519	page anchor links break	Closed	Fixed
	CONF-1459	Special characters in "search" result in various problems	Closed	Won't Fix
	CONF-1550	Page Information: Incorrect Short URL if webapp context is not "/"	Closed	Cannot Reproduce
	CONF-1515	Special german characters (ä ö ü) not being exported to PDF properly	Closed	Fixed
	CONF-1146	Use of quotes in page name brings issues with it	Closed	Fixed
	CONF-1326	Can't change mail format of daily summary messages.	Closed	Fixed
	CONF-1269	Export space (html) fails, if attached image missing	Closed	Fixed
	CONF-1699	Entering "*" as search string causes error.	Closed	Fixed
	CONF-1710	Upgrade failed: Can't call commit when autocommit=true	Closed	Fixed
	CONF-1663	A small error in the	Closed	Fixed
	CONF-1350	Random PermissionCheckDispatcher crashes when viewing pages	Closed	Fixed
	CONF-1418	Spaces in File links not preserved on Wiki Export	Closed	Fixed
	CONF-1417	jiraissues macro ignores the first entry in the colums parameter	Closed	Cannot Reproduce
	CONF-1401	Wrong document tree in exported PDF files	Closed	Fixed
	CONF-1403	Date formats of jiraissues macro	Closed	Fixed
	CONF-1439	TXT daily summary generate raw HTML	Closed	Duplicate
	CONF-1432	InfrastructureException after update	Closed	Resolved Locally
	CONF-1445	Icons don't get exported in PDFs	Resolved	Fixed
	CONF-1452	rename page bug	Closed	Fixed
	CONF-1468	Display of paragraphs in comment blocks not consistent	Closed	Fixed
	CONF-1480	URL returned from search is incorrect for attachements	Closed	Fixed
	CONF-1492	"short link" not respecting site URL	Closed	Fixed
	CONF-1507	Non-administrator users can perfrom some admin tasks	Closed	Fixed
	CONF-1533	PowerPoint search not working (example included)	Closed	Fixed
	CONF-1535	Nullpointer exception on updating a page that is being watched by a new user	Closed	Fixed
	CONF-1549	"Restore Page" restores entire confluence?	Closed	Fixed

	CONF-1593	java.lang.ClassNotFoundException: weblogic.jdbc.extensions.WLConnection	Closed	Fixed
	CONF-1615	user suddenly gets system error messages	Closed	Fixed
	CONF-1662	NPE in Mail notification	Closed	Duplicate
	CONF-1681	Only confluence admins can add and modify templates (space or global)	Closed	Fixed
	CONF-1677	java.lang.IllegalAccessError on PDF Export	Closed	Fixed
	CONF-1369	Error when export page as PDF	Closed	Fixed
	CONF-1228	Granting "anyone" use permission breaks with JIRA user management	Closed	Duplicate
	CONF-1691	$generalUtil.formatDateTime($page.lastModificationDage) showing up all over the place	Closed	Fixed
	CONF-1252	(Confluence Changes in the last 24 hours) has incorrect link to edit profile	Closed	Fixed
	CONF-1249	Zip file creation problems	Closed	Fixed
	CONF-1205	Search results for velocity get truncated on wiki.opensymphony.com	Closed	Fixed
	CONF-1420	Bulleted list items from JSPWiki not converted when no space after *	Closed	Fixed
	CONF-1443	On doubleclick inside Add Comment editor new page is loaded and comment is lost	Closed	Fixed
	CONF-1431	Error commenting Blog	Closed	Resolved Locally
	CONF-1428	cannot set page with unicode characters in title as parent page	Closed	Fixed
	CONF-1673	Importing into Swan seems to wipe space permissions	Closed	Fixed
	CONF-1404	Empty error queue message incorrect	Closed	Fixed
	CONF-1499	Wrong number of licensed users reported when anonymous access enabled	Closed	Duplicate
	CONF-1512	User Macros with upper-case characters ignored	Closed	Fixed
	CONF-1579	Exception when attaching a file	Closed	Fixed
	CONF-1574	Using % in page title causes error	Closed	Fixed
	CONF-1650	Administrate tab only shows if user is in confluence-admininstrators	Closed	Fixed
	CONF-1645	{excerpt-include} macro is documented with the wrong syntax	Closed	Fixed

Permissions Changes in 1.2

For Swan, changes were made to the way Confluence manages and checks permissions. This document is a guide for anyone migrating from Confluence 1.1.2 or earlier describing why the changes were made, and what this means to existing Confluence installations.

Why Change?

There were two goals behind changing permissions in Confluence:

1. Fix a significant problem whereby users with external user management enabled could not also enable Confluence anonymous access without blowing out their 25-user workgroup license.
2. Make the user interface for assigning and viewing permissions easier to use and understand.

What Changed?

"Anyone" User Removed

In Confluence 1.1.3, there was an "Anyone" user, who represented anyone using the system, whether logged in or not.

This user was the source of the licensing problem, and no longer exists in Confluence 1.2

Anonymous Permissions Added

Confluence 1.2 has explicit permissions for anonymous users. These permissions are only granted to users who are not logged in.

There is nothing stopping an administrator granting some permission to anonymous users, but not granting it to logged-in users. This results in the slightly bizarre possibility that a user might be able to do more before logging in than they can after. Luckily, this is a pretty easy situation for administrators to avoid.

Enabling anonymous access has no effect on Confluence's user count for licensing purposes.

Guard Permissions Added

The roles of the "Use Confluence" and "View Space" permissions have been expanded so that now they are *required* permissions before a user or group can be granted any more rights.

Before a user has access to *anything* in the Confluence server, they must first have "Use Confluence" permission, and likewise before a user has access of any kind to a space, they must first have the "View Space" permission.

For licensing purposes, your number of users is equivalent to the number of non-anonymous users with the "Use Confluence" guard permission.

Migrating from 1.1 to 1.2

Migrating Automatically

When you upgrade from Confluence 1.1, or when you restore a backup created in 1.1 into Confluence 1.2, an upgrade task will run to automatically migrate your permissions to the new scheme, while keeping them consistent with your 1.1 security settings. The task will make the following changes:

- All "Anyone" permissions will be converted into two separate permissions: one for Anonymous access, and one for the `confluence-users` group. (If the `confluence-users` group does not exist, this step will be skipped)
- Any user or group with some global or space permission will also be granted the equivalent guard permission.

After starting up with the new version of Confluence, we suggest that you check that the permissions have migrated successfully. While we have tested the migration code, maintaining your site's security is important enough to warrant a double-check, just in case.

Migrating Manually

If the automatic migration does not complete successfully, which would most likely happen if you have removed the "confluence-users" group, you will need to perform the above steps manually, through the user administration interface.

Release Notes 1.2.1

Confluence 1.2.1 is a maintenance release that fixes some bugs that users may have encountered using Confluence 1.2. It incorporates improvements to performance for large Confluence installations, and fixes bugs related to the remote API, over-use of disk space, and a few annoying errors users were experiencing when setting up a new Confluence instance.

1.2.1 is a free upgrade for all existing Confluence customers.

Who should upgrade?

The issues resolved below are all either fixes to problems that have effected small numbers of users, or improve areas of Confluence that may not be used in your installation. As such, we recommend you read through the release notes and decide whether this upgrade is necessary.

If 1.2 is working fine, and none of the issues below are bothering you, there is no need to upgrade.

Upgrade Procedure

Upgrading Confluence should be pretty easy. If you are upgrading from Confluence 1.2, you can find instructions here. *We strongly recommend that you backup your `confluence.home` directory and database before upgrading!*

If you are upgrading from Confluence 1.1.2 or earlier, be sure to read the upgrade instructions in the Confluence 1.2 release notes.

Changes in 1.2.1

See also: Issues Resolved for 1.2.1

Remote API Fixes

- getBlogEntries() will no longer erroneously return SQL Timestamps instead of date objects — CONF-1756)
- getPage() will now successfully retrieve a previous page version — CONF-1721
- renderContent() can now successfully render a page containing the {blog-posts} macro — CONF-1717

Setup Fixes

- Some users were encountering the following error during setup: "Unable to store Trigger with name: 'backupTrigger' and group: 'DEFAULT', because one already exists with this identification." This should no longer occur — CONF-1760
- Some users were encountering a NullPointerException in org.apache.lucene.store.FSDirectory.create() during setup when connecting to an external datasource. This should no longer occur — CONF-1767

Performance and Efficiency Fixes

- Installations with large search indexes will no longer experience a performance degredation when saving content — CONF-1759
- Backups will no longer leave a redundant exploded copy of the files being backed up in the temp directory — CONF-1752
- Deleting a space was not deleting its attachments from the confluence home directory. This is now fixed — CONF-1765
- Thumbnails are now stored in their own directory, so they won't be included unnecessarily in backups — CONF-1785

Other Issues Resolved

- Confluence now generates significantly higher-quality image thumbnails (Thanks to Mike Aizatsky for the tip) — CONF-1725

> **Note**
>
> Unless you are running JDK 5.0 (which we do not recommend as it is still in pre-release), this improvement will only be visible if you are running Confluence on Mac OS X

- URLs longer than 255 characters will no longer cause an exception when saving a page — CONF-1743
- Pages with titles containing quotes no longer break PDF exports — CONF-1719
- Internet Explorer 6 SP2 will no longer corrupt zip-files downloaded from Confluence — CONF-1669

Issues Resolved for 1.2.1

Errors were reported by the JIRA trusted connection.

- APP_UNKNOWN; Unknown Application: {0}; ["confluence:4557196"]

JIRA Issues (19 issues)

Type	Key	Summary	Status	Resolution
	CONF-1785	Move thumbnails out of attachments directory	Closed	Fixed
	CONF-1752	Be more proactive about cleaning temp directory after backup	Closed	Fixed
	CONF-1825	"Log In" link should not be relative to server base URL	Closed	Fixed
	CONF-1759	Don't optimise lucene index every time we add to it	Closed	Fixed
	CONF-1669	Downloads corrupt in IE 6 XP SP2	Closed	Fixed
	CONF-1775	Strange behaviour with deactivated users & JIRA user delegation	Closed	Fixed
	CONF-1476	Attachments seem to randomly fail	Resolved	Fixed
	CONF-1743	URLs > 255 chars in a page kill page	Closed	Fixed
	CONF-1719	Exporting a page with a quote in the title breaks space exporting	Closed	Fixed
	CONF-1765	Deleting a space does not delete its attachments	Closed	Fixed
	CONF-1760	Unable to store Trigger with name: 'backupTrigger' and group: 'DEFAULT', because one already exists with this identification.	Closed	Fixed
	CONF-1717	{blog-posts} macro isn't rendered during RPC renderContent call	Closed	Fixed
	CONF-1592	SnipSnap import renames pages because of illegal title characters - but doesnt update the links	Closed	Fixed
	CONF-1721	getPage for a prev version fails	Closed	Fixed
	CONF-1767	NullPointerException in Lucene FSDirectory.create() during setup	Closed	Fixed
	CONF-1729	Error purging referrers	Closed	Cannot Reproduce
	CONF-1725	Poor quality of thumbnails	Closed	Fixed
	CONF-1756	getBlogEntires marshalling bug	Closed	Fixed
	CONF-1494	Quartz ObjectAlreadyExistsException exception on MS SQL Server		Duplicate

			Closed	

Release Notes 1.2.2

Confluence 1.2.2 is a maintenance release that fixes some bugs that users may have encountered using Confluence 1.2 and 1.2.1. It fixes problems with attachment downloading, text file imports and text-only emails.

1.2.2 is a free upgrade for all existing Confluence customers.

Who should upgrade?

Confluence 1.2.2 includes a fix for CONF-1810. This bug will cause serious problems for anyone who attaches multiple versions of the same file to a Confluence page, corrupting subsequent downloads.

As such, Confluence 1.2.2 is a *recommended upgrade* for anyone running 1.2 or 1.2.1.

Upgrade Procedure

Upgrading Confluence should be pretty easy. If you are upgrading from Confluence 1.2 or 1.2.1, you can find instructions here. *We strongly recommend that you backup your `confluence.home` directory and database before upgrading!*

If you are upgrading from Confluence 1.1.2 or earlier, be sure to read the upgrade instructions in the Confluence 1.2 release notes.

Changes in 1.2.2

See also: Issues Resolved for 1.2.2

- The correct file size will now be stored for multiple versions of the same attachment, fixing a serious download corruption bug – CONF-1810
- Long {excerpt} blocks no longer overrun their database field – CONF-1705
- The {code} macro will no longer garble XML – CONF-1829
- Page diffs now escape HTML tags correctly – CONF-1830
- Fixed Postgresql error when you try to import text files that contain the nul (\0) character – CONF-1739
- "New Blog Post" icon in Page Operations URL now goes to the correct URL – CONF-1848
- "Next" link at the bottom of the alphabetical page listing now works – CONF-1797
- Text-formatted Confluence daily emails are now sent as text – CONF-1724
- The search input box on the "404 Not Found" page has been fixed – CONF-1800

Issues Resolved for 1.2.2

Errors were reported by the JIRA trusted connection.

- APP_UNKNOWN; Unknown Application: {0}; ["confluence:4557196"]

JIRA Issues (11 issues)

Type	Key	Summary	Status	Resolution
	CONF-1812	Modified user list was being confused with new user list for daily email notifications	Closed	Fixed
	CONF-1731	Phrase searches are case-sensitive	Closed	Cannot Reproduce
	CONF-1739	java.lang.IllegalArgumentException: \0 not allowed during text file import	Closed	Fixed
	CONF-1724	Confluence daily mails are HTML but wrong Content-type	Closed	Fixed
	CONF-1705	Long excerpts (>255 ch) kills page	Closed	Fixed
	CONF-1797	"Next" link at bottom of alphabetical listing page is broken.	Closed	Fixed
	CONF-1830	Diff output is not HTML escaped	Closed	Fixed
	CONF-1810	Uploading new version of attachment does not update file size in database	Closed	Fixed
	CONF-1848	"New Blog Post" icon in Page Operations (page.operations-icons) has incorrect URL.	Closed	Fixed
	CONF-1800	Search form on 404 page not working	Closed	Fixed
	CONF-1670	NullPointerException in setup restore step	Resolved	Fixed

Release Notes 1.2.3

Confluence 1.2.3 is a maintenance release that fixes some bugs that users may have encountered using the 1.2 series of Confluence. It fixes problems with Oracle support on Weblogic, thumbnail display, errors moving pages between spaces.

1.2.2 is a free upgrade for all existing Confluence customers.

Who should upgrade?

The issues resolved below are all either fixes to problems that have effected small numbers of users, or improve areas of Confluence that may not be used in your installation. As such, we recommend you read through the release notes and decide whether this upgrade is necessary.

Of particular note should be CONF-1911 which might effect customers trying to use Confluence with Oracle under Weblogic, and CONF-1914, a security issue where users may determine the names of attachments that they can not access.

If Confluence 1.2.2 is working fine for you, feel free to stick with it.

Upgrade Procedure

Upgrading Confluence should be pretty easy. If you are upgrading from Confluence 1.2 or higher, you can find instructions here. *We strongly recommend that you backup your `confluence.home` directory and database before upgrading!*

If you are upgrading from Confluence 1.1.2 or earlier, be sure to read the upgrade instructions in the Confluence 1.2 release notes.

Changes in 1.2.3

See also: Issues Resolved for 1.2.3

- Oracle LOB handler now works correctly under Weblogic – CONF-1911
- Attachments are no longer listed in search results if the user is not permitted to download them – CONF-1914
- Tomcat no longer hijacks Confluence's logging – CONF-1934
- Incoming page links no longer listed multiple times – CONF-1928
- Comments in daily update report are now linked properly – CONF-1904
- Thumbnails and the gallery macro no longer draw images with width and height of 0 – CONF-1861
- Fixed a system error moving pages between spaces when they are linked to from a comment – CONF-1851
- Fixed a system error when removing an attachment – CONF-1861

Issues Resolved for 1.2.3

Errors were reported by the JIRA trusted connection.

- APP_UNKNOWN; Unknown Application: {0}; ["confluence:4557196"]

JIRA Issues (28 issues)				
Type	**Key**	**Summary**	**Status**	**Resolution**
	CONF-1817	Datasource issues	Closed	Fixed
	CONF-1751	Add an "archive this backup" option to the manual backups	Resolved	Fixed
	CONF-1928	Only list a referring page once	Closed	Fixed
	CONF-1626	Page title is changed to Error and Page Operations is blank	Closed	Incomplete
	CONF-1730	Incoming links are duplicated if the links are to different anchors	Resolved	Duplicate
	CONF-1622	Top re-edit button for Templates does not work	Resolved	Cannot Reproduce
	CONF-1774	Renaming Home page causes system error	Closed	Handled by Support
	CONF-1913	Invalid key param throws NPE	Closed	Fixed
	CONF-1581	RPC call to render does not render blog posts	Resolved	Cannot Reproduce
	CONF-1794	"Incoming Links" only works for links in correct case	Resolved	Duplicate
	CONF-1911	Could not create Oracle LOB	Closed	Fixed
	CONF-1869	The {gallery} macro generates 1 pixel thumbnails	Closed	Fixed
	CONF-1904	New comments in daily change report not linked	Closed	Fixed
	CONF-1933	NullPointerException when removing attachment	Closed	Fixed

	CONF-1757	RPC excption returns html (500 page)	Resolved	Cannot Reproduce
	CONF-1851	Exception moving page between spaces	Closed	Fixed
	CONF-1600	NullPointerException	Resolved	Cannot Reproduce
	CONF-1737	Gallery macro: no JDK image support	Closed	Won't Fix
	CONF-1861	Thumbnails are referenced that can not be drawn	Closed	Fixed
	CONF-1472	Repository corruption	Resolved	Incomplete
	CONF-1527	ClobStringType requires active transaction synchronization	Closed	Cannot Reproduce
	CONF-1603	Log4J configuration is ignored in stand-alone Confluence - and seems to be using jdk1.4 logging	Resolved	Won't Fix
	CONF-1934	Replace commons-logging usages with log4j	Closed	Fixed
	CONF-1905	Watch spaces does not send notifications unless "Notify on my actions" is checked	Closed	Cannot Reproduce
	CONF-1914	Unpermitted attachments are displayed in search results	Closed	Fixed
	CONF-1778	Directory view listing only shows space link	Resolved	Won't Fix
	CONF-1627	x Identical Incoming links are displayed x times	Resolved	Duplicate
	CONF-1920	Importing from JSPWiki breaks some links	Closed	Fixed

Release Notes 1.3

Atlassian Software is proud to present Confluence 1.3 (otherwise known as Murrumbidgee). Existing customers who wish to upgrade, or new users who wish to try out Confluence for 30 days can download either the standalone or WAR distributions from the Atlassian website: http://www.atlassian.com/software/confluence

Murrumbidgee is the third major update to Confluence (in less than nine months!), and once again we've raised the bar of what you should expect from a professional wiki. Confluence 1.3 looks better, performs better, installs more easily and does far more than it ever has before.

A big thanks to everyone who reported bugs, and offered suggestions over the last few months, especially everyone who helped by trying out our Development Releases.

Upgrading from 1.2.3

Upgrading Confluence should be pretty easy. *We strongly recommend that you backup your confluence.home directory and database before upgrading!*

You will need to reindex your site after upgrading to enable some of the new site search features. You can do this from the "Rebuild Search Index" link on the Administration Console.

Also, we've added a bunch of space-level permissions since 1.2. While our upgrade process should make sure everyone has the same permissions after the upgrade as they did before, it's a good idea to check to make sure nothing has been missed.

Upgrading from 1.1.2 or Earlier

Users upgrading from an earlier version of Confluence should check the release-notes of the other major Confluence releases:

- Release Notes 1.2
- Release Notes 1.1

Contents

- New Features
- New Macros
- Improvements
- Notable Bug-fixes

See also: Issues Resolved for 1.3

New Features

Many of the features added to Confluence 1.3 are of interest only to site administrators. To get an idea of what's changed from the perspective of a user, you should read What's New in 1.3

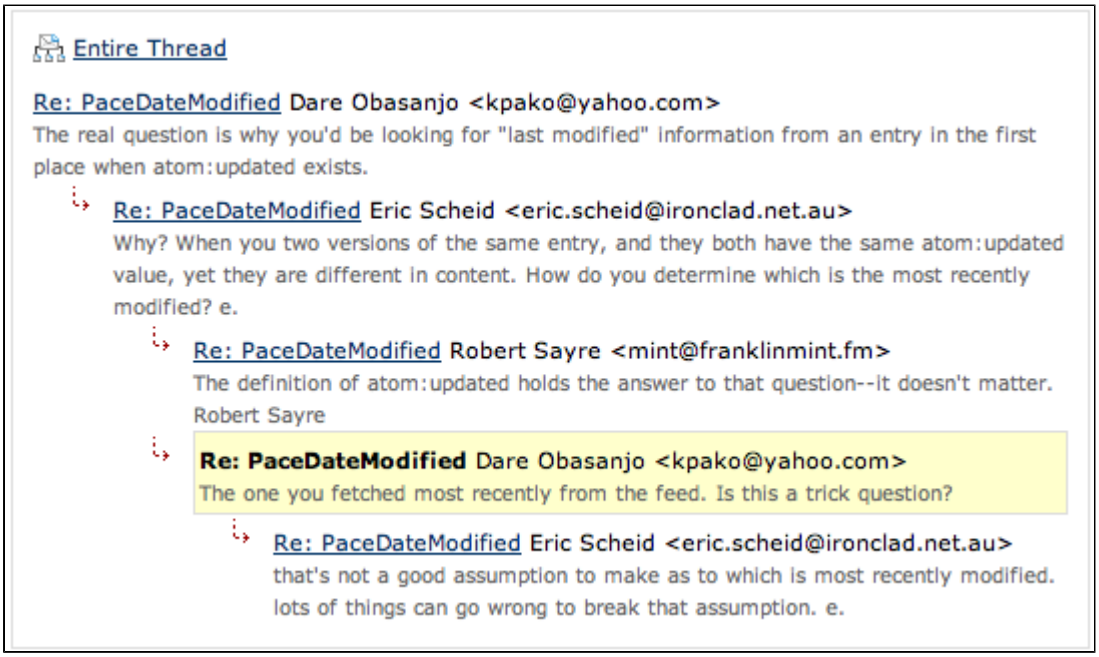
Mail Archiving



Confluence is a collaboration tool. When communication happens through Confluence, it gets archived, indexed and interlinked so the whole team can benefit from the information. But what happens to communication that happens via email? At best it gets hidden away in the corner of one team member's inbox, at worst it just falls into nowhere.

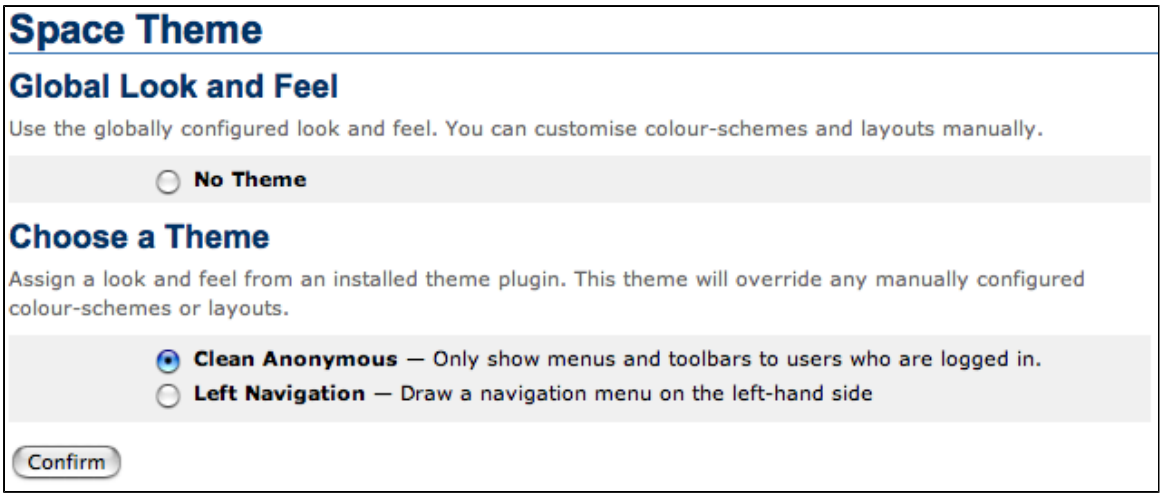
Now, you can put that mail into Confluence as well.

Each Confluence space has a mail archive. You can import mail directly by uploading an mbox file, or you can have Confluence regularly poll a POP mailbox for new mail. Once mail is imported into Confluence it can be browsed chronologically or by thread, and searched using Confluence's powerful full-text search functionality.



The threading and searching functionalities within Confluence are more powerful, and more useable than most dedicated mail archives that you will find online! For more information about Confluence's mail support, read the Mail Archiving FAQ.

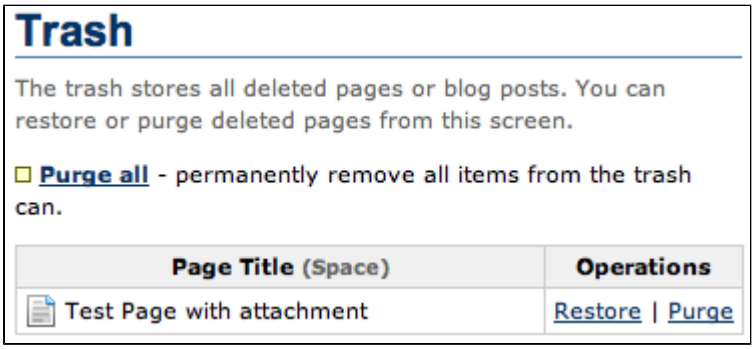
Themes



It is now possible to package a particular Confluence look and feel into a theme. Themes can be dropped into a Confluence installation, after which they will be available for global or space administrators to customise the look and feel of the site.

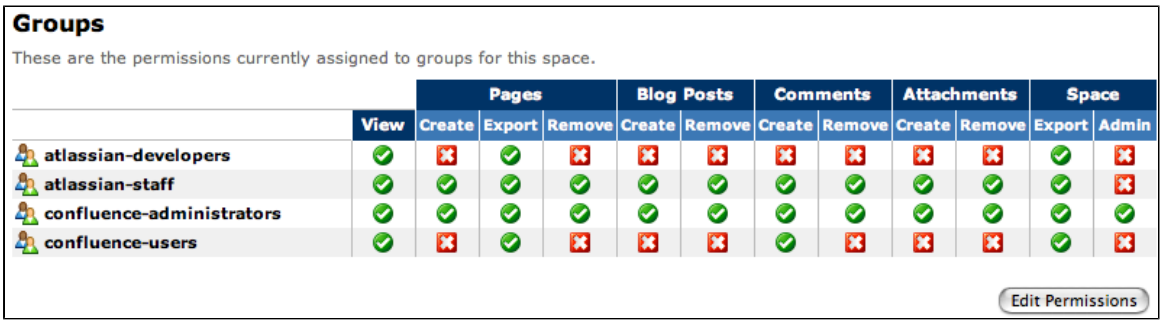
Themes are built using the new plugin architecture that has been built into Confluence: for more details on how to create your own theme, see Theme Module.

The Trash Can



One of the most popular feature-requests, it is now possible to undelete pages, mail and blog-posts. When content is deleted from a space it is moved to that space's trash can. Space administrators can restore deleted pages, or consign them to oblivion.

More Granular Space Permissions



Another of our most popular feature-requests, we've divided up the space permissions so administrators have more control over what users can and can not do.

New Emoticons

Emoticon	Name	Syntax
	add	(+)
	forbidden	(-)
	help	(?)
	idea on	(on)
	idea off	(off)
	star	(*)
	red star	(*r)
	green star	(*g)
	blue star	(*b)
	yellow star	(*y)

Because you can never have enough neat-looking graphics.

Macros

- {note}, {information}, {warning} and {tip} macros let you add coloured notes to your page.
- The {section} and {column} macros give you an alternative to wiki-markup tables
- The {jiraportlet} macro allows you to embed a JIRA 3 portlet into a Confluence page
- The {excerpt} macro has an optional `hidden=true` parameter to hide the contents of the excerpt within the page
- The {excerpt-include} macro has an optional `nopanel=true` paramater to display the excerpt without any decoration
- The {jiraissues} macro has an optional `count=true` parameter to display only the number of issues found, not the details of those issues. It also has an oprtional `cache=off` parameter to ensure that a list of Jira issues will refresh on each request.
- The {blog-posts} macro consistently displays blog-posts in reverse chronological order
- The {search} macro can be limited to particular types of content
- The notation guide has been reorganised to be more friendly to users who don't know what a macro is

Improvements

Referrer Management



For public sites, referrer management has been vastly improved in Confluence 1.3, with a new management screen, the ability to block unwanted referrers right from where they are displayed, and the ability to turn off referrer tracking with a single click.

Setup Wizard

The Confluence setup wizard has been rewritten from the ground up. We've made it task-oriented, and stripped it back to just the stuff you need to get Confluence running. The result is a much smoother, much faster installation.

User Interface

We've made a lot of improvements to Confluence's user interface. You'll notice some improvements on the Dashboard, making it easier to see exactly what's changed recently. You'll notice some *huge* changes to the space summary/space administration section of the site.

Backups

- You can now **exclude attached files from your backups**. Of course, this means you have to back up your attachment directory separately, but if you already have a good backup regime for your filesystem (and can thus restore attachments separately), it means your Confluence data backups will take a lot less space
- We now **include important system configuration in your backups**, so that when you restore a site from backup, it will work the same way as it did when you backed it up.

Indexing

We have improved the way we index content within Confluence, which means your searches are even more likely to find the right result.

If the primary language of your Confluence site is not English, you should change the "Indexing Language" preference in Confluence's General Configuration.

Site Performance

We've identified a number of places that were slowing down the performance of Confluence sites, from the dashboard to the spaces list page, to the search indexer, to the storing of referrers. Confluence should now perform faster than ever.

Also...

- Pages in an export are now in alphabetical, not creation order.
- The remote API can be accessed by anonymous users (this must be turned on in General Configuration)
- When previewing a page, you can continue editing without having to go back to the edit page
- You can link to anything in Confluence if you know its ID in the database (currently this is how you must link directly to mail) using the following link format: [$1234]
- You can link to anything relative to the root of the Confluence installation (useful for pointing to parts of the site that can't otherwise be linked) using three leading slashes: [///pages/editpage.action?pageId=1234]
- A whole lot more that we've forgotten...

Notable Bug Fixes

We resolved a *lot* of issues between Confluence 1.2.3 and Confluence 1.3. The best way to see what we've fixed is to ask JIRA, the world's best issue-tracker: Issues Resolved for 1.3

Confluence Presentation

This presentation was relevant to Confluence 1.3, so it was rather outdated.
We have removed it from this documentation space. You can find a copy of it here.

Issues Resolved for 1.3

Errors were reported by the JIRA trusted connection.

- APP_UNKNOWN; Unknown Application: {0}; ["confluence:4557196"]

JIRA Issues (200 issues)

Type	Key	Summary	Status	Resolution
	CONF-2241	1.3 Release Tasks	Closed	Fixed
	CONF-1953	Library upgrade: TextMining	Closed	Fixed
	CONF-1948	Library upgrade: HTTP Client	Resolved	Fixed
	CONF-1952	Library upgrade: WebWork	Closed	Fixed
	CONF-2032	Library upgrade: EHCache	Resolved	Fixed
	CONF-2022	Improve unit testing on attachments being exported	Closed	Fixed
	CONF-2048	Library upgrade: Hibernate 2.1.6	Closed	Fixed
	CONF-1949	Library upgrade: Spring	Resolved	Fixed
	CONF-1954	Library upgrade: PDFBox	Closed	Fixed

	CONF-1878	Plugin-ify Confluence macros	Resolved	Fixed
	CONF-1877	Incorporate atlassian-plugins module into Confluence	Closed	Fixed
	CONF-2185	Task Macro should be installed in default build of Conf.	Closed	Fixed
	CONF-1955	Library upgrade: Lucene	Resolved	Fixed
	CONF-1951	Library upgrade: Sitemesh upgrade	Resolved	Fixed
	CONF-1950	Library upgrade: Seraph	Closed	Fixed
	CONF-1964	Make the capitalisation of TrackBack consistent	Resolved	Fixed
	CONF-1958	EDITSPACE permission is confusingly named now that it's really "create page"	Closed	Fixed
	CONF-2138	Confluence has performance issue that the frontend Apache Proxy did not get response sometimes.	Closed	Handled by Support
	CONF-2292	revise 'Setting up Confluence'	Closed	Fixed
	CONF-2286	back out fix on ConfluenceAuthenticator which forced to lower case	Closed	Fixed
	CONF-2244	Get DamageControl UATs running again	Closed	Fixed
	CONF-2267	Document Blog-entry macro	Closed	Fixed
	CONF-2061	Create demo content for first-time users	Closed	Fixed
	CONF-2248	Remove paths from admin	Closed	Fixed
	CONF-2280	Test FatCow on 1.3 final	Closed	Fixed
	CONF-2243	Merge 1_2_STABLE into HEAD	Closed	Fixed
	CONF-2242	Remove 1.2.4 from JIRA. Check issues fixed in HEAD.	Closed	Fixed
	CONF-2069	Improve final setup screen	Resolved	Fixed
	CONF-2149	Add not-yet-configured warnings to confluence admin console	Closed	Fixed
	CONF-2013	Search Interface for Mail	Resolved	Fixed
	CONF-2096	Link to single message	Closed	Fixed
	CONF-2053	Add delete mail functionality	Closed	Fixed
	CONF-2119	Add Mail Operations menu	Closed	Fixed
	CONF-2052	Add delete mail permission	Closed	Fixed
	CONF-2055	Notify of unviewed mail matches in search	Closed	Fixed

	CONF-2051	Index messageid column in database	Closed	Fixed
	CONF-2054	Add Mail icon to spaces list on dashboard	Closed	Fixed
	CONF-2057	Display attachments in view mail page	Closed	Fixed
	CONF-2012	View Single Mail	Closed	Fixed
	CONF-2018	Test mail accounts	Closed	Fixed
	CONF-1974	POP mailbox polling	Closed	Fixed
	CONF-2058	Database checking added to setup	Resolved	Fixed
	CONF-2068	Data setp in setup - demo, no data, import	Closed	Fixed
	CONF-2014	Mail Browser	Closed	Fixed
	CONF-2024	Remove mails from "recent changes" lists	Closed	Fixed
	CONF-1972	Mail domain objects/manager/dao	Resolved	Fixed
	CONF-2015	Extract Attachments	Closed	Fixed
	CONF-1976	Mail indexing	Closed	Fixed
	CONF-2288	Upgrade wiki.theserverside.com and remove referrers	Closed	Fixed
	CONF-2278	Link API docs from everywhere	Closed	Fixed
	CONF-2255	Check unit tests are running 100%	Closed	Fixed
	CONF-2249	Base URL admin	Closed	Fixed
	CONF-2247	Remove dummy.gif	Closed	Fixed
	CONF-2226	Add JIRA portlet macro	Closed	Fixed
	CONF-2016	Mail Account Management	Resolved	Fixed
	CONF-2269	Document gallery macro	Closed	Fixed
	CONF-2067	Remove paths configuration from setup	Closed	Fixed
	CONF-1894	Confluence needs a note macro	Resolved	Fixed
	CONF-2212	Add preference to disable remote API	Closed	Fixed
	CONF-1610	Count {include:Page} as a link to "Page"	Closed	Fixed
	CONF-2110	Link to any content object by ID	Resolved	Fixed

	CONF-2029	Lucene indexing queue	Resolved	Fixed
	CONF-2059	Threading and Related Mail	Resolved	Fixed
	CONF-1375	How about more complex permission?	Closed	Fixed
	CONF-1856	Decorator Themes	Resolved	Fixed
	CONF-1870	Access key for edit	Resolved	Fixed
	CONF-1586	Allow referrers to be turned off	Closed	Fixed
	CONF-1779	Make exporting a permission	Closed	Fixed
	CONF-1136	Radical idea, archive external email in Confluence	Resolved	Fixed
	CONF-851	Deleted Pages - Need a 'recycle bin'	Closed	Fixed
	CONF-1871	Include default demo space as a new install option	Closed	Fixed
	CONF-1654	Ability to add extra colour settings to a colour scheme	Resolved	Fixed
	CONF-2256	Layout macros to enable people to create complex page layouts	Resolved	Fixed
	CONF-2224	Improve JIRA macros	Closed	Fixed
	CONF-2220	Self-documenting macros	Closed	Fixed
	CONF-2195	Create system link capability	Closed	Fixed
	CONF-1992	Retrieve page by space key and page title	Resolved	Fixed
	CONF-1965	Confluence Mail Archive	Resolved	Fixed
	CONF-2221	Limit remote API search by space/date/content types	Closed	Fixed
	CONF-1631	Page Edition / Preview should be in one	Resolved	Fixed
	CONF-2271	Add "nopanel" parameter to excerpt-include macro	Closed	Fixed
	CONF-2028	Space summary recently updated panel should show new comments, blogs, etc	Closed	Duplicate
	CONF-2139	Make full thread view not a popup	Closed	Fixed
	CONF-2270	Add "hidden" parameter to excerpt macro	Resolved	Fixed
	CONF-2211	Improve referrer links performance	Resolved	Fixed
	CONF-2210	Fix Setup UI for select db connection type	Closed	Fixed
	CONF-1999	Move "Start watching this space" next to the RSS icon - and make it just an icon (use the tooltip for descrioption)	Closed	Fixed

	CONF-2201	Remove space link in spacelist.vm has 'Add Page' title text on the img	Closed	Fixed
	CONF-2049	Setup Wizard Improvements	Closed	Fixed
	CONF-1639	Anonymous SOAP and XML-RPC access	Closed	Fixed
	CONF-1793	Handle things better if we get an error loading confluence.cfg.xml	Resolved	Fixed
	CONF-2193	Mail archive graphic and description should be placed on initial Content pane in Space Summary	Closed	Fixed
	CONF-2194	Index macro to display excerpts	Closed	Fixed
	CONF-2209	Ensure that setup UI is consistent	Resolved	Fixed
	CONF-1261	Simplify the setup wizard	Closed	Fixed
	CONF-2114	Sort home-page drop down in edit space alphabetically	Closed	Fixed
	CONF-1237	Hide unsupported databases in setup pick-list.	Closed	Fixed
	CONF-2125	Search performance improvements	Resolved	Fixed
	CONF-2172	Remove "pool size" setting from database setup	Closed	Fixed
	CONF-1435	Allow attachments to be backed up separately	Resolved	Fixed
	CONF-2050	Improve Mail for DR4	Resolved	Fixed
	CONF-2109	Display relative times in "recent updates" lists	Closed	Fixed
	CONF-2004	Change "Remove Blog Post" icon to trashcan icon used for "Remove Page"	Resolved	Fixed
	CONF-2019	jiraissue macro now takes a cache='on' or 'off' argument	Closed	Fixed
	CONF-1605	Long Blog titles dont wrap very well	Closed	Fixed
	CONF-1946	Make help icon point to manual on c.a.c	Resolved	Fixed
	CONF-1865	Provide navigation options for screen after creating a new template	Resolved	Fixed
	CONF-1330	Make blog-posting a separate permission	Resolved	Fixed
	CONF-1891	Permissions Screen suggestion	Resolved	Won't Fix
	CONF-1674	Add JIRA-style user-picker component	Closed	Fixed
	CONF-1764	Improve UI for permissions editing	Closed	Fixed
	CONF-1928	Only list a referring page once	Closed	Fixed
	CONF-1864	Don't escape shortcut links if no parameters	Closed	Fixed

	CONF-1562	Seperate Space Summary and Space Administration	Closed	Fixed
	CONF-1176	When moving an page, it should not select the first project	Closed	Fixed
	CONF-1700	Calendar L&F not integrated with the rest of Confluence	Resolved	Fixed
	CONF-1740	When text file import fails, report name of file that died.	Closed	Fixed
	CONF-1858	Noisy dashboard 1	Closed	Fixed
	CONF-1551	Editing a blog post is a little difficult to find. Not intuitive.	Closed	Fixed
	CONF-2245	Allow macros to choose which page documentation occurs on	Closed	Fixed
	CONF-2223	Improve mail excerpting slightly	Closed	Fixed
	CONF-2151	Indicate where a thread continues back or forward in mail view	Closed	Fixed
	CONF-2148	'browse templates' in admin. screen breaks with the admin decorator	Closed	Fixed
	CONF-2000	Make all of "Create a new blog post" url-ified	Closed	Fixed
	CONF-1221	When previewing a page, display edit box below the preview	Closed	Fixed
	CONF-1052	Remove "Path" section of admin config	Resolved	Fixed
	CONF-2160	Speed up spaces list page	Closed	Fixed
	CONF-1720	Blog calendar has no Next/Previous month links	Closed	Fixed
	CONF-1513	Help for user macros in interface	Closed	Fixed
	CONF-1862	Add the (?) emoticon	Closed	Fixed
	CONF-1578	Remove/Delete Space should be on "Spaces" page	Closed	Fixed
	CONF-2197	Reorganise Notation Guide	Closed	Fixed
	CONF-2182	Improve performance of getSpace(key)	Closed	Fixed
	CONF-2094	excerpt on mail search result contains return path	Resolved	Fixed
	CONF-2044	Hook email address hiding preference into mail display	Resolved	Fixed
	CONF-2046	Exclude attachments from backup	Closed	Fixed
	CONF-1661	NPE using TaskList Macro	Resolved	Fixed
	CONF-2756	The backup that doesnt!	Resolved	Fixed
	CONF-2186	&s break links in Confluence	Closed	Invalid

	CONF-1687	{blog-posts} wrong sort order when using "time" parameter	Resolved	Fixed
	CONF-1787	embedded tab characters break {{code}} filter	Closed	Fixed
	CONF-1947	cannot remove myself from list	Closed	Incomplete
	CONF-1919	Tasklist Macro renders as null in preview mode	Closed	Duplicate
	CONF-2020	Confluence won't allow Groups with upper case letter names	Closed	Fixed
	CONF-1966	export 32 pixel gif missing on Space Summary - Content	Closed	Fixed
	CONF-1789	{content-by-user} Macro shows items in Restricted Spaces	Closed	Fixed
	CONF-1828	When emoticon used as link text, outgoing link section is garbled	Closed	Fixed
	CONF-1887	Nullpointer on Global Admin View Templates link	Closed	Fixed
	CONF-1838	After creating user, go to user details page	Closed	Fixed
	CONF-2311	Atlassian stops getting mail if it can't parse one message	Closed	Fixed
	CONF-2165	Single-page XML exports can be imported without a space	Closed	Fixed
	CONF-2306	xerces-2.1.1.jar is corrupt	Resolved	Fixed
	CONF-1287	"400 Bad Request" response when viewing page with quotes in title on Orion	Closed	Fixed
	CONF-2124	PDF Space Export includes deleted page	Closed	Fixed
	CONF-2074	Can create duplicate user by appending space to an existing username	Closed	Fixed
	CONF-1933	NullPointerException when removing attachment	Closed	Fixed
	CONF-1886	Browse templates breaks space colour scheme	Closed	Fixed
	CONF-1555	Edit My Profile alignment is off	Closed	Fixed
	CONF-1827	Javascript errors on unknown page links	Closed	Fixed
	CONF-2171	Incoming Links vanished from page-operations	Closed	Fixed
	CONF-1409	Restore doesn't restore space colours	Closed	Fixed
	CONF-1198	search is not working with non-ascii characters	Closed	Fixed
	CONF-948	Internal links don't work if a page is included using the #includePage macro	Closed	Fixed
	CONF-925	Create admin account step in Setup dies with an ungraceful DuplicateEntityException	Resolved	Fixed
	CONF-1758	Page store should still succeed even if indexing fails	Resolved	Duplicate

	CONF-1834	Space summary does not show new comments	Closed	Fixed
	CONF-1826	cannot add new spaces or update existing ones	Resolved	Fixed
	CONF-1814	Still need confluence-admin group to access /admin pages	Closed	Fixed
	CONF-1867	improper encoding in confluence.cfg.xml	Closed	Fixed
	CONF-1854	The new space permissions page doesnt use my colors	Resolved	Fixed
	CONF-1853	Gallery macro error when previewing create page	Resolved	Fixed
	CONF-1895	Please upgrade Tomcat used for building standalone distribution	Closed	Fixed
	CONF-1888	New global template reports 'Undefined Space' in breadcrumb	Closed	Fixed
	CONF-1425	On upgrade 1.1 -> 1.1.1 Confluence freezes on some time.	Resolved	Cannot Reproduce
	CONF-1424	Setting 'User email visibility' doesn't restored from backup	Closed	Fixed
	CONF-1481	NPE calling Search function on conf.atlas.com	Resolved	Fixed
	CONF-1537	ResourceNotFoundException when viewing Site Decorator for a Space	Closed	Fixed
	CONF-1557	Wrong URL causes IllegalStateException	Resolved	Fixed
	CONF-1552	Wonky alignment in edit blog post if post is very short	Closed	Fixed
	CONF-1583	Searching particular PPT fails	Closed	Fixed
	CONF-1618	Pages are exported in creation, not alphabetic order	Resolved	Fixed
	CONF-1672	jira-issues macro fails for URLs containing brackets	Closed	Fixed
	CONF-2235	Could not initialize proxy - the owning Session was closed	Closed	Resolved Locally
	CONF-2236	Unable to build search query: null	Closed	Fixed
	CONF-2216	friendly meg size is not being set in velocity var. (admin - backup and restore)	Closed	Fixed
	CONF-2206	Repair setup process - demo content && installed successful links to demo content homepage	Closed	Fixed
	CONF-2192	Password field for creating a mail account is plain textfield	Closed	Fixed
	CONF-2199	Errors upgrading from 1.2 -> 1.3 under MySQL: net.sf.hibernate.PropertyAccessException: exception setting property value with CGLIB	Closed	Fixed
	CONF-2176	Email footers have errors	Closed	Fixed
	CONF-2170	Upgrade from DR3 to DR4 broke shortcut links on confluence.atlassian.com	Closed	Fixed
	CONF-2169	Need to patch the "duplicate page" bug	Resolved	Fixed

	CONF-2150	Next and previous links in mail broken when context path is /	Closed	Fixed
	CONF-2147	Installation of Demo content does not bring up a demo content welcome screen on custom installation	Resolved	Fixed
	CONF-2159	Fix merge comment in setupdbchoice.vm	Closed	Fixed
	CONF-2158	Restore setup step does not validate existence of file	Closed	Fixed
	CONF-2157	Character encoding issues in archiving Mail messages	Closed	Fixed
	CONF-2161	Export page HTML has both upper & lower case space id	Closed	Fixed
	CONF-2141	Graceful error message upon reaching attachment filesize limit.	Closed	Fixed
	CONF-2363	Rebuild Index run forever	Closed	Fixed
	CONF-1996	Junit Macro not rendering	Closed	Fixed

What's New in 1.3

Many of the changes between Confluence 1.2 and 1.3 are really of interest only to Confluence administrators. Here we detail the changes that will be important to all Confluence users.

The Trash

Previously, when you deleted a page or blog post, it was unrecoverable. In 1.3, deleting a page moves it to the Trash. Space administrators can retrieve pages from the trash, or purge them forever. If a page has been deleted (or you have deleted a page by mistake), contact a space administrator to get it back.

Note Macros

There are new macros for inserting coloured notes into a page:

{note}:

This is a note
A note tells you about something that may be important to you.

{tip}:

This is a tip
A tip tells you something you might not have thought of yourself.

{info}:

This is informative
Info was an excuse to have a blue note.

{warning}:

Beware!
Warnings can be dangerous if overused, because people start ignoring them.

For more information, check out the notation guide.

Mail Archiving

Confluence now includes the capability to store email. This allows you to store mailing-list archives, or records of conversations amongst your team inside Confluence. Mail archiving must be set up by a space administrator.

Have a look at [an individual mail], or the view of an entire mail thread (We are using confluence.atlassian.com to keep archives of the confluence-user mailing list).

Searched for **excerpt AND macro** in **Confluence Documentation** excluding **Mail.**
Also found **15 matches** in **Mail.** Include these matches.

All mail is indexed, but by default we exclude it from search results because the volume of email can often overwhelm the content of the Confluence site. When you do a search, you might see this reminder that there could be an answer to your question in the email archives:

To link to an email from a Confluence page, you must find the numeric ID of that mail from the end of its URL, and put that in your link like so: [$12572], which Confluence will draw like this: [$12572]



Improved JIRA Macros

The {jiraissues} macro has two new optional parameters:

- `count=true` makes the macro only return the number of issues matched by your filter, not the whole list of issues
- `cache=off` makes the macro retrieve the filter results from JIRA every time the page is loaded, ensuring the results are accurate. (Be careful if the filter returns a lot of results, though, you don't want to overload your servers)

There is a new {jiraportlet} macro that allows you to retrieve any portlet from a server running JIRA 3, and display it in a Confluence page, like so:

Project: **Confluence** (CONF)	
Project Lead: Per Fragemann [Atlassian] **Reports:** Summary | Issues | Road Map | Change Log | Popular Issues | Calendar | Labels | Versions | Components | Agile	
Open Issues: (By Priority)	
Filter Issues: • All • Resolved recently • Outstanding • Added recently • Unscheduled • Updated recently • Most important	

For more information, check the notation guide



Easy Column Layout

You can use the {section} and {column} macros to organise your page into columns. This is especially useful when you combine it with the {jiraportlet} macro: you can arrange a Confluence page just like a JIRA dashboard!

Here's a simple two-column layout:

Lorem ipsum dolor sit amet, consectetuer adipiscing elit. Aliquam fermentum vestibulum est. Cras rhoncus. Pellentesque habitant morbi tristique senectus et netus et malesuada fames ac turpis egestas. Sed quis tortor. Donec non ipsum. Mauris condimentum, odio nec porta tristique, ante neque malesuada massa, in dignissim eros velit at tellus. Donec et risus in ligula eleifend consectetuer. Donec volutpat eleifend augue. Integer gravida sodales leo. Nunc vehicula neque ac erat. Vivamus non nisl. Fusce ac magna. Suspendisse euismod libero eget mauris.

Lorem ipsum dolor sit amet, consectetuer adipiscing elit. Aliquam fermentum vestibulum est. Cras rhoncus. Pellentesque habitant morbi tristique senectus et netus et malesuada fames ac turpis egestas. Sed quis tortor. Donec non ipsum. Mauris condimentum, odio nec porta tristique, ante neque malesuada massa, in dignissim eros velit at tellus. Donec et risus in ligula eleifend consectetuer. Donec volutpat eleifend augue. Integer gravida sodales leo. Nunc vehicula neque ac erat. Vivamus non nisl. Fusce ac magna. Suspendisse euismod libero eget mauris.

For more information, check the notation guide



Other Things

- Take a look at how the space summary page has been reorganised
- If you hit alt-E on any wiki page or blog post (ctrl-E if you're using a Mac), you'll be taken to the edit page.
- If you start a link with three slashes, you can link to something relative to the root of the Confluence installation. This is useful for

creating links to pages that are part of Confluence, such as the dashboard ([///]), or the space list ([///spaces/listspaces.action]).
- The {excerpt} macro can take a `hidden=true` parameter to hide the contents of the excerpt within the page
- The {excerpt-include} macro can take a `nopanel=true` parameter to display the excerpt without any decoration or tables

Release Notes 1.3.1

Confluence 1.3.1 is a maintenance release that fixes some bugs that users may have encountered using Confluence 1.3.

1.3.1 is a free upgrade for all existing Confluence customers.

Who should upgrade?

Confluence 1.3.1 fixes a number of bugs that were found in Confluence 1.3. However, none of the bugs that were fixed were considered critical or likely to cause data-loss. Administrators should only upgrade Confluence if they are affected by (or feel they would be affected by) one of the issues resolved by this release.

Upgrade Procedure

Upgrading Confluence should be pretty easy. If you are upgrading from Confluence 1.3, you can find instructions here. *We strongly recommend that you backup your `confluence.home` directory and database before upgrading!*

If you are upgrading from Confluence 1.2.3 or earlier, be sure to read the upgrade instructions in the Confluence 1.3 release notes.

Changes in 1.3.1

See also: Issues Resolved for 1.3.1

- Preferences for hiding or masking email addresses are no longer ignored on some screens – CONF-2352
- The plugin manager will no longer fail with an error when trying to return a plugin to its default state after activating or deactivating it – CONF-2396
- Disabling hot referrers in the referrer manager screen now works as expected – CONF-2397
- The save option has been restored to the comment preview screen – CONF-2321
- Deleting a blog post or a page no longer leaves possible orphaned comments in recent changes lists – CONF-2323
- Also, some edge cases in the setup wizard were fixed, the demonstration content was tidied up, and a few minor UI issues were resolved.

Issues Resolved for 1.3.1

Errors were reported by the JIRA trusted connection.

- APP_UNKNOWN; Unknown Application: {0}; ["confluence:4557196"]

JIRA Issues (24 issues)				
Type	**Key**	**Summary**	**Status**	**Resolution**
	CONF-2390	Demo space needs to be prepared (again) for point and dev. releases	Closed	Fixed
	CONF-2277	Put API docs online	Closed	Fixed
	CONF-2319	Add cache=true or false details for jiraissues macro to notation guide	Closed	Fixed
	CONF-2380	Sort user group names alphabetically in 'Edit Groups' screen	Closed	Fixed
	CONF-1987	remove cvs author tags from source build	Closed	Fixed
	CONF-2302	Viewing email is disallowed if POP boxes aren't configured	Closed	Cannot Reproduce
	CONF-2391	Remove mail should remove associated attachments	Closed	Fixed
	CONF-2396	Error activating macro	Closed	Fixed
	CONF-2366	Wording: "Your password successfully changed"	Closed	Fixed
	CONF-2386	External user management flag does not disble user management links	Closed	Fixed
	CONF-2362	Preview shouldn't jump you down to the edit area	Closed	Fixed
	CONF-2400	investigate reported problem when restoring data during the setup	Closed	Fixed
	CONF-2361	Email hiding preferences ignored	Closed	Fixed
	CONF-2310	Erraticity in state of applicationConfig	Closed	Fixed
	CONF-2320	Restore backup option in setup is represented incorrectly in setup step menu	Closed	Fixed
	CONF-2323	Deleting a blog post doesn't delete its comment	Closed	Fixed

	CONF-2321	Comment preview no longer lets you submit comment	Closed	Fixed
	CONF-2317	note/warning/info/tip macro icons not showing up in notation help	Closed	Fixed
	CONF-5460	CLONE -Wording: "Your password successfully changed"	Closed	Fixed
	CONF-2314	demo-site exportDescriptor.properties contains @BUILD_DATE@	Closed	Fixed
	CONF-2304	Imported email is logged into catalina.out	Closed	Fixed
	CONF-2389	demonstration space has two broken links	Closed	Fixed
	CONF-2397	Hot referrers switch not working	Closed	Fixed
	CONF-2334	Long "word" in post can make summary too wide in dashboard	Closed	Fixed

Release Notes 1.3.2

Confluence 1.3.2 is a maintenance release which includes 30 bug fixes and improvements that users may have requested using Confluence 1.3.

1.3.2 is a free upgrade for all existing Confluence customers.

Who should upgrade?

We recommend that all 1.3.x users upgrade to Confluence 1.3.2 . It includes, among other things, the fix of a memory leak which had been occuring in our error monitoring (see below or CONF-2540); if you have been encountering performance issues within Confluence, this could be of benefit.

Upgrade Procedure

Upgrading Confluence should be pretty easy. If you are upgrading from Confluence 1.3 or 1.3.1, you can find instructions here. *We strongly recommend that you backup your `confluence.home` directory and database before upgrading!*

If you are upgrading from Confluence 1.2.3 or earlier, be sure to read the upgrade instructions in the Confluence 1.3 release notes.

Changes in 1.3.2

See also: Issues Resolved for 1.3.2

We've closed 30 issues, so we will not name them all. Among the most promiment in the list are:

- A memory leak in our collection of logging events as a threadlocal, see CONF-2540.
- Consistent UI handling for editing and previewing blogs and comments, see CONF-2479 and CONF-2470
- Improved IO handling when building exports, see CONF-2510
- Better cleaning up attachments when pages are removed, see CONF-2567
  etc ...

Issues Resolved for 1.3.2

Errors were reported by the JIRA trusted connection.

- APP_UNKNOWN; Unknown Application: {0}; ["confluence:4557196"]

JIRA Issues (25 issues)

Type	Key	Summary	Status	Resolution
	CONF-2308	Verify oracle upgrade from 1.2 -> 1.3	Resolved	Fixed
	CONF-2355	Bandana upgrade error 1.2 -> 1.3	Closed	Cannot Reproduce
	CONF-2536	Markers around variable content of each page.	Closed	Fixed
	CONF-2479	Comment preview should mirror Page preview	Closed	Fixed
	CONF-2566	Import from disk should have a browse filesystem button	Closed	Invalid
	CONF-2490	Error while trying to draw the last-n pages!	Closed	Cannot Reproduce
	CONF-2567	Removing a page should remove its attachments	Closed	Fixed
	CONF-2414	XML-RPC methods are not allowed to return void	Closed	Fixed
	CONF-2449	Unsupported Database setup option does not let you pick a hibernate dialect	Closed	Fixed

	CONF-2458	MS SQL Server connection string wrong	Resolved	Fixed
	CONF-2426	There is no documentation for userlister	Closed	Fixed
	CONF-2442	attachments names with spaces get messed up	Resolved	Duplicate
	CONF-2413	display the user friendly alias for the system link in page summary	Closed	Fixed
	CONF-2324	Limit length of traceback links on blog pages	Closed	Fixed
	CONF-2507	rss macro creates bad links when feed doesn't contain links	Closed	Fixed
	CONF-2495	editmyprofile.action should blow up if you're not logged in	Closed	Fixed
	CONF-2420	disable change password link in profile	Closed	Fixed
	CONF-2484	html renders in recently updated	Closed	Fixed
	CONF-1771	Shortcut links containing an '@' do not work	Closed	Fixed
	CONF-2488	Import/Restore fails: truncation error	Resolved	Fixed
	CONF-2470	Blog title with illegal character looses blog entry	Closed	Fixed
	CONF-2441	Notation guide error	Closed	Fixed
	CONF-2418	Creating a blog post with bad characters in the title loses page content	Closed	Fixed
	CONF-2552	Check for existence of attachments before trying to delete attachments	Closed	Fixed
	CONF-2540	Ensure limit on ThreadLocalAppender for Log4J	Closed	Fixed

Release Notes 1.3.4

Confluence 1.3.4 is a maintenance release which includes a few important bug fixes for anyone running Confluence. These release-notes include information about the (briefly available) 1.3.3 version of Confluence, which 1.3.4 has superseded.

1.3.4 is a free upgrade for all existing Confluence customers.

Who should upgrade?

We recommend that all 1.3.x users upgrade to Confluence 1.3.4. For customers running 1.3.2 and earlier, it includes a fix to the security issue described in Confluence Security Advisory 2005-02-09, and it thus an important upgrade for anyone who is still running an un-patched system.

For customers running 1.3.3, the upgrade is also recommended as this version fixes CONF-2740, a regression in 1.3.3 which could cause referrer and trackback data to disappear.

Upgrade Procedure

Upgrading Confluence should be pretty easy. If you are upgrading from Confluence 1.3 or 1.3.1, you can find instructions here. *We strongly recommend that you backup your `confluence.home` directory and database before upgrading!*

If you are upgrading from Confluence 1.2.3 or earlier, be sure to read the upgrade instructions in the Confluence 1.3 release notes.

Changes in 1.3.4

See also: Issues Resolved for 1.3.4

Since only five bug-fixes were made for this release, here is a quick run-down.

- The patch for the security issue described in Confluence Security Advisory 2005-02-09 has been applied to 1.3.4
- A regression in 1.3.3 that caused trackback and referrer data to be lost has been fixed – CONF-2731
- An intermittent ClassCastException when viewing blog-posts has been fixed – CONF-1671
- The image cache for Apache FOP is now cleared between exports, preventing a case where the wrong version of an attachment may end up in a PDF export – CONF-2647
- A spurious "Connection already closed" log message in the JIRA user management bridge no longer occurs CONF-2656

Issues Resolved for 1.3.4

Issues resolved for 1.3.3

Errors were reported by the JIRA trusted connection.

- APP_UNKNOWN; Unknown Application: {0}; ["confluence:4557196"]

JIRA Issues (6 issues)				
Type	**Key**	**Summary**	**Status**	**Resolution**
	CONF-2438	Do not send daily update email if no updates have taken place	Closed	Cannot Reproduce
	CONF-2554	Jira integration database connection issue	Resolved	Duplicate
	CONF-2624	Exception occurred inside setter of com.atlassian.confluence.links.OutgoingLink.destinationPageTitle during import	Resolved	Fixed
	CONF-2550	Missing resourcebundle, fails long running task	Closed	Fixed
	CONF-2647	PDF Export does not pick up image updates.	Closed	Fixed
	CONF-2656	FATAL [user.providers.jira.JiraJdbcProfileProvider] Unable to close connection: Connection is closed.	Closed	Fixed

Issues resolved for 1.3.4

Errors were reported by the JIRA trusted connection.

- APP_UNKNOWN; Unknown Application: {0}; ["confluence:4557196"]

JIRA Issues (3 issues)				
Type	**Key**	**Summary**	**Status**	**Resolution**
	CONF-1671	Possible ClassCastException in Blog.getDatePath	Resolved	Fixed
	CONF-2740	Referrers not showing	Closed	Fixed
	CONF-2731	Trackbacks display error $generalUtil.wordwrap($tbl.title, 32)	Closed	Fixed

Release Notes 1.3.5

Confluence 1.3.5 is a maintenance release which fixes a number of bugs found in earlier 1.3.x Confluence releases. Some areas improved in this release are space export/import, search and LDAP user management. 1.3.5 is also the first Confluence version to be successfully tested against Resin 3.0: see below for the special steps you need to take to get it running.

1.3.5 is a free upgrade for all existing Confluence customers.

Who should upgrade?

We recommend that all 1.3.x users upgrade to Confluence 1.3.5. For customers running 1.3.2 and earlier, it includes a fix to the security issue described in Confluence Security Advisory 2005-02-09, and it thus an important upgrade for anyone who is still running an un-patched system.

For customers running 1.3.4, the upgrade is also recommended as this version fixes CONF-2750, a regression in 1.3.4 that made it impossible to page through search results.

Upgrade Procedure

Upgrading Confluence should be pretty easy. If you are upgrading from Confluence 1.3 – 1.3.4, you can find instructions here. *We strongly recommend that you backup your `confluence.home` directory and database before upgrading!*

If you are upgrading from Confluence 1.2.3 or earlier, be sure to read the upgrade instructions in the Confluence 1.3 release notes.

Changes in 1.3.5

Resin 3.0 Compatibility

Confluence 1.3.5 is the first Confluence version to be successfully tested against Resin 3.0.12. Customers wishing to run Confluence under Resin 3 should read the special instructions at Known Issues for Resin 3.x

Issues Resolved

- Paging through search results no longer results in an error. CONF-2750
- Space export and import now performs much more reliably. CONF-2678
- Setup wizard no longer complains about an incomplete setup if you restore a backup from disk. CONF-2637
- Special characters in search no longer cause parser to explode. CONF-2527, CONF-2532, CONF-2728, CONF-2735...
- Attachment filenames containing spaces are no longer truncated when downloading using Firefox. CONF-2739

- For the full list of fixes, see Issues Resolved for 1.3.5

Issues Resolved for 1.3.5

Errors were reported by the JIRA trusted connection.

- APP_UNKNOWN; Unknown Application: {0}; ["confluence:4557196"]

JIRA Issues (33 issues)				
Type	**Key**	**Summary**	**Status**	**Resolution**
	CONF-1000	Get Confluence working on Resin 3.x	Closed	Fixed
	CONF-2781	Backups and Restore MUST WORK	Closed	Fixed
	CONF-2318	Space Summary - Admin main page should have icons and explanations as per Content	Resolved	Won't Fix
	CONF-2788	Preview doesnt show page title	Resolved	Fixed
	CONF-2527	Searching for ***... BANG!	Closed	Fixed
	CONF-2739	File names are truncated when accessing attachments	Resolved	Fixed
	CONF-2750	Error viewing Next >> seach results	Closed	Fixed
	CONF-2728	Leading Wildcard "*searchterm" leads to exception in search	Closed	Fixed
	CONF-2532	"System Error" when searching for ~username	Closed	Fixed
	CONF-2678	Demo Space import fails	Resolved	Duplicate
	CONF-1511	Confluence forces login on every page view (Resin 3.0 incompatibility)	Resolved	Fixed
	CONF-2415	Pages with long titles are inaccessable	Resolved	Fixed
	CONF-2322	NPE in jiraissues macro	Resolved	Fixed
	CONF-2086	Problem deploying on Jboss under SuSe or FreeBSD	Resolved	Fixed
	CONF-2820	Error when paging through search results	Resolved	Duplicate
	CONF-2837	Broken mail configuration makes admin console inaccessible	Resolved	Fixed
	CONF-2779	RemoteSpaceSummary.hashCode has NPE if key is not set	Resolved	Fixed
	CONF-2771	MappingException in ReferralTaskQueue	Resolved	Fixed
	CONF-2762	Content link dies if target is a comment.	Resolved	Fixed
	CONF-2561	LazyInitializationException deleting a user	Resolved	Fixed
	CONF-2627	Can't thumbnail a file called attachments.png	Resolved	Fixed
	CONF-2634	Confluence import creates two velocity directories instead of one (breaking decorators)	Resolved	Fixed
	CONF-2637	Tried to finish setup but had not run through the whole wizard?	Resolved	Fixed
	CONF-2655	Special characters break search	Closed	Fixed
	CONF-2387	quick search breaks when the [character is used	Closed	Fixed
	CONF-1509	All '/display/' links redirect to login page for Resin 3.0.8	Closed	Fixed
	CONF-2819	Linked pages in "Move Page" screen contains bogus spaces ($content.space.name)	Closed	Fixed
	CONF-2768	User management with LDAP ignores users registered only in Confluence	Closed	Fixed
	CONF-2769	returning a null if referring content is not a space and walking into a NPE	Closed	Fixed
	CONF-2764	Database Dialect always shows up as N/A in system dump	Closed	Fixed
	CONF-2735	Search error on confluence.atlassian.com	Closed	Fixed
	CONF-2751	Exported space fails import	Resolved	Fixed
	CONF-2737	LDAP and non LDAP user login at Confluence	Closed	Fixed

Release Notes 1.3.6

Confluence 1.3.6 is a special maintenance release for customers who are using Confluence with an Oracle database. It contains a number of fixes, originally developed for Confluence 1.4, to Confluence's behaviour against Oracle databases.

1.3.6 is available as a free upgrade for all existing Confluence customers from the archive download page.

Who should upgrade?

We recommend that customers skip Confluence 1.3.6 and upgrade to Confluence 1.4. Confluence 1.4 includes a great many new features and bug-fixes that are not included in 1.3.6. We do, however, recognise that 1.4 is a significant upgrade and some customers wish to continue with Confluence 1.3 for the time being.

Confluence 1.3.6 is recommended only for Confluence customers who:

- are not yet ready to upgrade to Confluence 1.4, and
- are running Confluence with an Oracle database

Upgrade Procedure

Upgrading Confluence should be pretty easy. If you are upgrading from Confluence 1.3 – 1.3.4, you can find instructions here. *We strongly recommend that you backup your `confluence.home` directory and database before upgrading!*

Changes in 1.3.6

Confluence 1.3.6 addresses a number of situations in which Confluence could lock up, consume additional resources, or fail to perform backup or restore operations against an Oracle database.

Release Notes 1.4

Atlassian Software is proud to present Confluence 1.4 (otherwise known as Hunter). Existing customers who wish to upgrade, or new users who wish to try out Confluence for 30 days can download either the standalone or WAR distributions from the Atlassian website: http://www.atlassian.com/software/confluence

Hunter is the fourth major update to Confluence, and the biggest yet in terms of new features and improvements both visible to the user or hidden "under the hood". Between 1.3.5 and 1.4, we resolved a massive 480 issues.

A big thanks to everyone who reported bugs, and offered suggestions over the last few months, especially everyone who helped by trying out our Development Releases.

See also: Issues Resolved for 1.4

Contents

- Upgrading
- Migration
- New Features
- Notable Bug Fixes
- Outstanding Bugs

Upgrading From a Previous Version of Confluence

Upgrading Confluence should be pretty easy: you can find [instructions here]. *We strongly recommend that you backup your confluence.home directory and database before upgrading!*

Upgrading from 1.3.5

Because of the significant UI changes between 1.3 and 1.4 (see below), we highly recommend disabling any themes or UI customisations before upgrading Confluence.

After upgrading to 1.4, administrators will need to rebuild the site's search-index to ensure all the new search features are enabled. Do this from the *Content Indexing* section of the global administration menu.



Upgrade Migration
The database migration tasks that must be run to upgrade Confluence from 1.3 to 1.4 may take several minutes to run. During this time, Confluence will be unavailable.

Confluence License Expiration
A commercial Confluence license entitles you to a year of free upgrades. As such, Confluence 1.4 is the first release of Confluence to which some customers may not be able to upgrade.

If you were issued your Confluence license before May 24th 2004, you will *not* be able to upgrade to Confluence 1.4. Confluence 1.4 will refuse to run with any license issued before this date, and you will be required to downgrade to a previous Confluence version.

Upgrading from 1.2.3 or Earlier

Users upgrading from an earlier version of Confluence should check the release-notes of the other major Confluence releases:

- Release Notes 1.3
- Release Notes 1.2
- Release Notes 1.1

Migration

The best place to find out about the new features available in Confluence 1.4 is by checking out our What's New in Confluence 1.4 pages. However, there are a few features of specific interest to Confluence administrators that will be mentioned here:

v2Renderer

The engine by which Confluence converts wiki-text to HTML has been rewritten from scratch for Confluence 1.4. This was necessary, as the complexity of Confluence's markup was overwhelming the capabilities of the previous engine and leading to an increasing number of bugs, some of which had the potential to bring down the entire server.

v2Renderer is in every way more powerful (and better, more predictable) than the previous engine, and has undergone a lot of testing to make sure that it renders wiki pages the same way as its predecessor. However, pages that made use of bugs or undocumented features in the original engine may no longer render the same in 1.4 as they did in 1.3. If you come across something that is a bug or a missing feature in the new renderer, please file an issue in JIRA.

One side-effect of the new renderer is that custom Java macros written for Confluence 1.3 and earlier may not be compatible with Confluence 1.4 (although most should continue to function). For more information, macro authors should read this document

New User Interface

In response to a *lot* of feedback from customers and users of Confluence, the User Interface has undergone a major overhaul between Confluence 1.3 and 1.4. While the changes are mostly self-explanatory, and we believe the new interface is significantly easier for *everyone* to use and understand than its predecessor, it may be a good idea to make the following resource available to your users as part of the transition: 1.4 Interface - Where Did Everything Go?

Because of the substantial changes to the interface, themes and UI customisations made for Confluence 1.3 and earlier are not compatible with Confluence 1.4, and should be disabled before you upgrade.

Blogs have become News

Also in response to a lot of feedback, blogs in Confluence have now been renamed to *news*. We feel this makes it a lot easier to explain them to new and non-technical Confluence users. None of the functionality of blogs has been changed, just the name. 

New Features

Confluence 1.4 has a lot of cool new stuff. Over the next week we'll be releasing a new section of the Confluence website that will examine the new features in detail. Stay tuned.

Notable Bug Fixes

We resolved a *lot* of issues between Confluence 1.3.5 and Confluence 1.4. The best way to see what we've fixed is to ask JIRA, the world's best issue-tracker: Issues Resolved for 1.4

Outstanding Bugs

Some bugs were introduced during the Confluence 1.4 development cycle that we could not fix in time for the final release. Of note are:

- Index rebuilding may fail on multi-processor systems (for a workaround, see CONF-3168)
- You can not install an XWork plugin by uploading through the web interface, it must be copied into `WEB-INF/lib` (CONF-3184)
- Uploading a malformed plugin through the web interface may make other plugins cease to function (CONF-3183)
- Under certain circumstances, the mbox mail import may fail against Oracle databases (CONF-3284)
- The Insert Image and Insert Link popups jump the cursor to the top of the text input area on browsers other than Internet Explorer (CONF-3232)

Once again, if you find any bugs in Confluence, or have any feature suggestions, you can report them online in JIRA.

Issues Resolved for 1.4

Error formatting macro: jiraissues: java.lang.RuntimeException: A value with ID '11191' does not exist for the field 'project'.

Release Notes 1.4.1

Confluence 1.4.1 is a maintenance release that resolves some issues users may have encountered using Confluence 1.4. Issues include an occasional failure to display Confluence pages, Oracle and MySQL database issues, and a slow memory-leak.

1.4.1 is a free upgrade for all customers who purchased their Confluence license after June 2nd, 2004.

Who should upgrade?

Confluence 1.4.1 is a recommended upgrade for all users of Confluence 1.4, as it addresses a memory-leak which, while slow, would progressively degrade Confluence's performance and ultimately cause the server to exhaust available memory.

Upgrade Procedure

Upgrading Confluence should be pretty easy. If you are upgrading from Confluence 1.4, you can find instructions here. *We strongly recommend that you backup your `confluence.home` directory and database before upgrading!*

If you are upgrading from Confluence 1.3.6 or earlier, be sure to read the upgrade instructions in the Confluence 1.4 release notes.

Changes in 1.4.1

See also: Issues Resolved for 1.4.1

- A slow memory leak has been fixed. (CONF-3347)
- An issue where Confluence pages would intermittently only display their header when loading has been resolved. (CONF-3295)
- Confluence will now deploy reliably on MySQL 4.1 with UTF-8 encoding. (CONF-3306)
- Several issues regarding Confluence's use of the Oracle database have been resolved.
- A recurring bug preventing PNG images from being exported to PDF on some application servers has been fixed. (CONF-731)
- Inconsistencies with custom colour-schemes have been resolved (CONF-3314)

Issues Resolved for 1.4.1

Errors were reported by the JIRA trusted connection.

- APP_UNKNOWN; Unknown Application: {0}; ["confluence:4557196"]

JIRA Issues (16 issues)				
Type	Key	Summary	Status	Resolution
	CONF-3306	Update os-propertyset key size to ensure that it works with mysql	Resolved	Fixed
	CONF-2720	Tasklist should allow for unnamed list	Closed	Fixed
	CONF-3400	Pages with edit-permissions are indexed as if they have view permissions	Resolved	Fixed
	CONF-3295	Page mysteriously fails to load, sometimes	Resolved	Fixed
	CONF-2672	Error while generating PDF	Closed	Fixed
	CONF-3284	Mail import failing on oracle.	Resolved	Fixed
	CONF-3255	Links in the {children} macro are mis-rendered in exported PDFs	Resolved	Cannot Reproduce
	CONF-3304	Non-profit footer too long	Resolved	Fixed
	CONF-3303	NPE in link macro?	Resolved	Fixed
	CONF-3314	Custom color scheme not applied to spacelist	Closed	Fixed
	CONF-3315	RemoveOldCommentVersionFromParentComment upgrade task fails on MySQL (and other case sensitive databases)	Closed	Fixed

	CONF-731	PNG images (as well as some of the emoticons which are PNG) can not be exported as PDF!!	Resolved	Fixed
	CONF-1478	Search of numbers yiels no result	Closed	Fixed
	CONF-783	Icons and images don't get exported in PDFs	Resolved	Fixed
	CONF-3347	Memory leak caused by registering Webwork actions with Spring	Resolved	Fixed
	CONF-3128	Links to Attachment disappeared in exported HTML page	Resolved	Fixed

Release Notes 1.4.2

Confluence 1.4.2 is a maintenance release that resolves some issues users may have encountered using previous Confluence 1.4 releases. 1.4.2 introduces Websphere and DB2 compatibility, and fixes issues related to content indexing and the mail queue.

1.4.2 is a free upgrade for all customers who purchased their Confluence license after June 30th, 2004.

Who should upgrade?

Confluence 1.4.2 is a bugfix release. Customers should consult the list of issues resolved for this release to decide whether it is worth their while upgrading. Since this release includes a number of important performance and reliability fixes for content indexing, anyone who relies on Confluence's search functionality should consider upgrading.

Upgrade Procedure

Upgrading Confluence should be pretty easy. If you are upgrading from Confluence 1.4 or 1.4.1, you can find instructions here. *We strongly recommend that you backup your `confluence.home` directory and database before upgrading!*

If you are upgrading from Confluence 1.3.6 or earlier, be sure to read the upgrade instructions in the Confluence 1.4 release notes.

Database Connection Pool Changes

Dependent on the size of your Confluence installation, the new re-indexing task may use up to 10 database connections simultaneously. As such, you may need to increase the size of your database connection pool in order to allow Confluence to continue to operate during and after rebuilding the search index.

If Confluence has been set up to use direct JDBC connections, then you will find the setting for Confluence's connection pool size in `confluence.cfg.xml` in your confluence home directory. Confluence should update this value automatically when you upgrade, so after starting up Confluence 1.4.2 for the first time, you should check to make sure it has been updated to 15 connections, as seen below:

```
1.<property name="hibernate.c3p0.max_size">15</property>
```

If Confluence has been configured to use a datasource supplied by the application server, you should ensure the application server is configured to supply Confluence with sufficient connections to run.

Changes in 1.4.2

See also: Issues Resolved for 1.4.2

- Confluence has now been tested on Websphere 5.1.1.3 (see Known Issues for Websphere)
- Confluence has now been tested against DB2 8.1 (Linux and Windows)
- The mechanism by which Confluence rebuilds its full-text search index has been rewritten to be significantly less memory-intensive (but may be 10-15% slower). CONF-3340
- An issue that could cause comments or attachments to be removed from the search index when a page was edited has been resolved. CONF-3489
- An issue that could cause index rebuilding to fail (especially on multi-processor systems) has been resolved. CONF-3168
- Issues related to the size and reliability of the mail queue were resolved. CONF-3334 CONF-3474
- An exception when creating global templates was fixed. CONF-3386
- Some Javascript errors that were being flagged on Internet Explorer no longer occur. CONF-3422, CONF-3482
- Exporting a space to HTML via XML-RPC now succeeds. CONF-3401
- Creating a blog post via XML-RPC now succeeds. CONF-3412
- ALL CAPS WORDS are no longer linked when CamelCase linking is activated. CONF-3361
- The display of user attachments does not break when the user has not yet set a profile. CONF-3420

Issues Resolved for 1.4.2

Errors were reported by the JIRA trusted connection.

- APP_UNKNOWN; Unknown Application: {0}; ["confluence:4557196"]

JIRA Issues (35 issues)

Type	Key	Summary	Status	Resolution
	CONF-3001	Website improvements	Resolved	Fixed
	CONF-3474	THe Mail Queues Error Queue gets Filled with Mail Exceptions	Closed	Fixed
	CONF-3294	Make Navigation Consistent	Resolved	Fixed
	CONF-3204	Proof-read, clear and deploy	Resolved	Fixed
	CONF-3202	Make front page spiffy	Resolved	Fixed
	CONF-3341	Support DB2 (Windows/Linux)	Resolved	Fixed
	CONF-3705	Don't show comments text area in print view	Resolved	Fixed
	CONF-3354	Turn off logging when exporting to PDF	Closed	Fixed
	CONF-3340	Reindexing shouldn't load all objects into memory at once	Resolved	Fixed
	CONF-3413	Incomplete i18n properties for core.dateutils.DateUtils	Resolved	Fixed
	CONF-3473	Cannot un-check "Macro has a body" flag	Closed	Fixed
	CONF-3361	CamelCase behaves much differently in 1.4.1 than 1.3.5	Resolved	Fixed
	CONF-3348	Blogpost preview shows Dashboard > Undefined Space > $helper.getText("action.name")	Closed	Fixed
	CONF-3489	Indexing a page removes the indexes for the comments on attachments belonging to the page	Resolved	Fixed
	CONF-1467	PDF Export of page with large image breaks	Closed	Fixed
	CONF-3405	Servlet plugins not covered by standard confluence servlet filters	Resolved	Fixed
	CONF-1234	Websphere 5.1 installation problems	Closed	Fixed
	CONF-2568	Unable to create blog post - bad sql grammar exception (DB2)	Resolved	Fixed
	CONF-3377	Upgrade issue from 1.3 to 1.4.1 with DB2	Resolved	Fixed
	CONF-3391	Invalid formatting in space.vmd of URL links to blog comments in News Items page if non-ascii characters present in blog title	Resolved	Fixed
	CONF-3422	JavaScript error when creating a page from link in undefined pages list	Resolved	Fixed
	CONF-3401	Exporting a space to HTML via XML-RPC results in a ClassCastException	Resolved	Fixed
	CONF-3412	storeBlogEntry() fails with NullPointerException if a publishDate is not set	Resolved	Fixed
	CONF-3410	Mail stops sending, MailExceptions in stdout.log, all mail moves to Error Queue	Closed	Duplicate
	CONF-3334	Default email MIME type is not valid	Closed	Fixed

	CONF-3443	Upgrade bug on Sybase: Column not NULL by default	Closed	Fixed
	CONF-3482	Javascript error in IE when clicking on non-existant page link	Resolved	Fixed
	CONF-3680	Problem with turkish characters. New line starts while it is not supposed to be with turkish characters.	Closed	Fixed
	CONF-1738	Problem sending email notification can cause email flooding	Resolved	Cannot Reproduce
	CONF-3375	Livesearch macro rendering (breaks panels for one)	Closed	Fixed
	CONF-1478	Search of numbers yiels no result	Closed	Fixed
	CONF-3379	Typo on no permissions page	Resolved	Fixed
	CONF-3381	$body always defined for user macros	Resolved	Fixed
	CONF-3386	Creating global template throws null pointer exception in getSpaceName	Resolved	Fixed
	CONF-3420	Attachments tab sometimes shows a macro instead of (0)	Resolved	Fixed

Release Notes 1.4.3

Confluence 1.4.3 is a maintenance release that resolves some issues users may have encountered using previous Confluence 1.4 releases. 1.4.3 introduces a new SOAP service provider, and fixes issues related to incoming link tracking, notifications management, serving resources from dynamically loaded plugins, and much more.

1.4.3 is a free upgrade for all customers who purchased their Confluence license after August 17th, 2004.

Who should upgrade?

Confluence 1.4.3 is a bugfix release. While Atlassian recommends customers always run the most recently available stable Confluence release, customers should consult the list of issues resolved to decide whether it is worth their while upgrading.

Upgrade Procedure

Upgrading Confluence should be pretty easy. If you are upgrading from Confluence 1.4 to 1.4.2, you can find instructions here. *We strongly recommend that you backup your `confluence.home` directory and database before upgrading!*

If you are upgrading from Confluence 1.3.6 or earlier, be sure to read the upgrade instructions in the Confluence 1.4 release notes.

Changes in 1.4.3

See also: Issues Resolved for 1.4.3

Remote API Changes

Confluence now ships with an second SOAP provider running Apache Axis, alongside the existing Glue SOAP provider. This is the first step towards migrating entirely to Axis as our sole SOAP provider over the next few major Confluence releases. If you use Confluence's SOAP API, you should read this document for information about the migration process and timeframe.

> ⚠ All SOAP clients that worked with Confluence 1.4.2 will continue to work with Confluence 1.4.3 with no change.

Other changes made to the remote API:

- A condition under which SOAP faults could be masked by HTML error pages on some application servers was resolved. (CONF-3043)
- Deleting a page via the remote API now places that page in the trash. (CONF-3403)
- Adding space level permissions to a user through the remote API now works as expected. (CONF-3596)

Other Changes and Fixes

- A bug that would cause a page's most recent editor to change when the page was previewed has been fixed (CONF-3424)
- A bug that could cause a page's incoming links not to be displayed has been fixed (CONF-3509)
- A bug that could cause "Watch this space" not to function has been fixed (CONF-3510)

- A bug that could cause Confluence to run out of available file descriptors when exporting a space has been fixed (CONF-3517)
- A bug that could cause an email notification to be sent regarding a page that was restricted by page-level permissions has been fixed (CONF-3573)
- The attachment upload form in the "Insert Image" is more robust against error (CONF-3677, CONF-3676)
- Searching for a page in the Parent Page Picker Popup now functions as expected (CONF-3364)
- Deleting a user now correctly deletes their email notifications (CONF-3619)
- Restoring a page from the trash now re-adds that page's comments to the search index (CONF-3564)
- Attachment links are now not rendered if the user can not view the attachment due to page level permissions (CONF-3553)
- The list of recently edited pages in a user's profile is now more accurate (CONF-2430)
- Plugins uploaded through the web interface can now correctly serve downloadable resources (CONF-3387)
- The {recently-updated} macro and recent blogs pages now work under JDK 1.5 (CONF-3528, CONF-3601)
- The {section} macro's "border" property now works correctly (CONF-3736, CONF-3433)
- The {noformat} macro now has a "nopanel" parameter that will suppress drawing a background or border (CONF-3656)
- Several issues regarding the {tasklist} macro were fixed (CONF-3622, CONF-3633, CONF-3632)
- The {dynamictasklist} macro no longer breaks PDF exports (CONF-3513)
- The informational macros no longer center their text when viewed in Internet Explorer (CONF-3537)
- Links are now parsed in the title of a {panel} macro (CONF-3560)
- The string ?!?! is no longer interpreted by the Wiki/HTML renderer as a broken image

See also: Issues Resolved for 1.4.3

Confluence SOAP Provider Migration

Starting with the release of Confluence 1.4.3, and stretching through the next three major releases, Confluence will be transitioning from using the Glue library to provide a SOAP remote API, to using Apache Axis. Unfortunately, while the SOAP services will stay the same, the WSDL that these libraries generate to interact with the same services *will change*, so SOAP applications that interact with Confluence will need to migrate with us.

The migration should be relatively painless. Since the underlying objects represented by the WSDL are still the same, the process should involve regenerating your SOAP stubs, and a few cosmetic code changes.

The XML-RPC API is unaffected by this change.

Current Status

1.4.3 and 1.5/2.0	Deploy Axis SOAP service alongside Glue	Completed
(unscheduled)	Allow configuration of default SOAP provider, the default at installation being Axis	Incomplete
(unscheduled)	Remove Glue SOAP provider	Incomplete

Progress on these issues can also be tracked via CONF-3141

The Plan

As of version 1.4.3, Confluence ships with three SOAP endpoints:

Endpoint	WSDL URL	Provider
`/rpc/soap/confluencesoapservice-v1`	`/rpc/soap/confluenceservice-v1.wsdl`	Glue
`/rpc/soap-glue/confluencesoapservice-v1`	`/rpc/soap-glue/confluenceservice-v1.wsdl`	Glue
`/rpc/soap-axis/confluencesoapservice-v1`	`/rpc/soap-axis/confluenceservice-v1?wsdl`	Axis

The Axis and Glue providers produce slightly different WSDL URLs


Third-party SOAP [RPC Plugins] deployed in Confluence will be similarly deployed in three locations

Over the next three major Confluence releases we will:

1. Make the provider of `/rpc/soap/confluenceservice-v1` configurable, then
2. Change the default provider to Axis (but leave it configurable), then finally
3. Remove the Glue provider entirely

What This Means for SOAP Clients

Confluence major releases occur every three or four months. Thus, authors of SOAP clients should keep in mind this timeline, starting with

the release of Confluence 2.0 (November 2005)

0-3 months	• All existing SOAP clients written against the Glue provider will continue to function normally. • New clients should be written to the Axis provider in `/rpc/soap-axis` • Existing clients should change their SOAP endpoints to point to `/rpc/soap-glue` (or, better still, move directly to Axis)
3-6 months	• Confluence will require configuration to work with any client of the Glue provider that has not changed its endpoint • New clients should be written to the Axis provider in `/rpc/soap-axis` • Existing clients will need to migrate to the Axis provider
6+ months	• Confluence will no longer work with clients written against the Glue provider • Both `/rpc/soap` and `/rpc/soap-axis` endpoints will continue to be served by the Axis provider

Issues Resolved for 1.4.3

Errors were reported by the JIRA trusted connection.

- APP_UNKNOWN; Unknown Application: {0}; ["confluence:4557196"]

JIRA Issues (50 issues)

Type	Key	Summary	Status	Resolution
	CONF-3658	Create axis soap service alongside glue service	Resolved	Fixed
	CONF-3670	Restrict the number of top level pages shown in the {children} macro.	Resolved	Fixed
	CONF-801	Support different sortings of page children	Resolved	Fixed
	CONF-3656	Add option to suppress panel in noformat macro	Resolved	Fixed
	CONF-3654	Edit panel and buttons disappear when previewing a global template	Closed	Fixed
	CONF-3424	Canceling edit after preview changes last editor	Resolved	Fixed
	CONF-3601	Ognl exception while getting property blogPosts	Resolved	Duplicate
	CONF-3505	ClassCastException in Recently Updated Macro	Closed	Fixed
	CONF-1843	Confluence doesn't reconnect to database after MySQL is restarted	Closed	Fixed
	CONF-3736	SectionMacro:border=true not responding.	Resolved	Fixed
	CONF-3619	Deleting user does not delete there notifications.	Resolved	Fixed
	CONF-3577	Need to handle multiple "watches" / "notifications" gracefully.	Resolved	Fixed
	CONF-3632	Duplicate tasklist macro items cause SQL duplicate key exception on restore	Closed	Fixed
	CONF-3564	Comments are not re-indexed when a page is restored from the trash	Resolved	Fixed
	CONF-3575	The actual order of pages in a PDF Space export does not follow the same order of the index	Closed	Fixed
	CONF-3633	Tasklist macro allows duplicate items if items contain trailing spaces	Closed	Fixed

	CONF-3711	Page Picker search results are always empty	Closed	Fixed
	CONF-3507	Creating a page containing the RSS macro may cause the server to hang against DB2	Resolved	Fixed
	CONF-3364	When adding a page and designating its parent, searching for a page does not work.	Resolved	Fixed
	CONF-1749	lotus notes emails screwy	Closed	Fixed
	CONF-2287	RPC: removing a non-existent space results in a 'No permission to remove spaces' XmlRpcSxception	Resolved	Fixed
	CONF-2430	User Profile's 'Recently Edited' items list is wrong.	Resolved	Fixed
	CONF-1324	Username can't have uppercase characters	Resolved	Fixed
	CONF-3387	Plugin resources not served when plugin installed in ${confluence.home}/plugins	Resolved	Fixed
	CONF-3433	border=true attribute in {section} tag turns on or off borders of tables in section but does not put border around section	Resolved	Fixed
	CONF-3428	Using {recently-updated} macro under JDK 1.5.0	Resolved	Duplicate
	CONF-3403	Delete page via SOAP api doesnt put it in the Trash	Resolved	Fixed
	CONF-3227	Cannot type "?!?!" characters without 404 image renderings	Resolved	Fixed
	CONF-3620	Bad action - after send forgot password	Closed	Fixed
	CONF-3676	insert image popup shows upload form even though the user does not have permissions.	Closed	Fixed
	CONF-3677	500 Internal Server Error wehn submitting empty upload form in insert image popup	Resolved	Fixed
	CONF-3510	Watch this space link doesn't toggle, adds multiple subscriptions	Resolved	Fixed
	CONF-3509	Incoming links are broken	Resolved	Fixed
	CONF-3517	Unclosed file handles.	Resolved	Fixed
	CONF-3537	Info macros have centered text on IE6.	Resolved	Fixed
	CONF-3530	Information Macros center content text when rendered in IE	Resolved	Fixed
	CONF-3572	confluence:docs maven target does not create javadoc for the com.atlassian.renderer packages.	Resolved	Fixed
	CONF-3573	Security issue with 'watching'	Resolved	Fixed
	CONF-3608	Remove all uses of EqualsBuilder	Resolved	Fixed
	CONF-3503	If an attachment is created with a null (rather than empty) comment, setting a new one causes an NPE	Resolved	Fixed
	CONF-3492	Granting a user/group export space permission actually grants "admin space" permission	Closed	Duplicate
	CONF-1478	Search of numbers yiels no result	Closed	Fixed

	CONF-3612	Incoming links went south?	Resolved	Fixed
	CONF-3641	Administration link: closing a tag typo	Closed	Fixed
	CONF-3506	Daily backup date format and prefix being set after bandana has been saved	Resolved	Fixed
	CONF-3553	Link to attachment on restricted page displays "401 Unauthorized" error	Resolved	Fixed
	CONF-3513	Error while exporting space in PDF (Possibly due to dynamictyasklist macro)	Resolved	Fixed
	CONF-3538	{panel} macro does not respond to properties unless a title is specified.	Resolved	Won't Fix
	CONF-3560	Links are not parsed in title of panel macro	Resolved	Fixed
	CONF-3596	Adding space permissions for a user via XML-RPC instead adds permissions for a gorup	Resolved	Fixed

Release Notes 1.4.4

Confluence 1.4.4 is a maintenance release that resolves some issues users may have encountered using previous Confluence 1.4 releases. 1.4.4 fixes a number of issues, including many related to the reliability of uploaded plugins

1.4.4 is a free upgrade for all customers who purchased their Confluence license after September 23rd, 2004.

Who should upgrade?

Confluence 1.4.4 is a bugfix release. While Atlassian recommends customers always run the most recently available stable Confluence release, customers should consult the list of issues resolved to decide whether it is worth their while upgrading.

Upgrade Procedure

Upgrading Confluence should be pretty easy. If you are upgrading from Confluence versions between 1.4 and 1.4.3, you can find instructions here. *We strongly recommend that you backup your `confluence.home` directory and database before upgrading!*

If you are upgrading from Confluence 1.3.6 or earlier, be sure to read the upgrade instructions in the Confluence 1.4 release notes.

Changes in 1.4.4

Uploaded Plugin Fixes

Prior to Confluence 1.4.4, plugins containing Spring components, XWork actions or velocity templates could not be uploaded through the web interface, and needed to be installed manually in `WEB-INF/lib`. Confluence 1.4.4 fixes these issues.

Database Indexes

Confluence 1.4.4 now correctly creates database indexes on a clean installation. If you are upgrading an existing Confluence instance, you should consult Performance Tuning for information on indexes that can improve Confluence's performance significantly.

Other Changes and Fixes

See: Issues Resolved for 1.4.4

Issues Resolved for 1.4.4

Errors were reported by the JIRA trusted connection.

- APP_UNKNOWN; Unknown Application: {0}; ["confluence:4557196"]

JIRA Issues (47 issues)				
Type	**Key**	**Summary**	**Status**	**Resolution**
	CONF-3980	Spelling error "hiearchy: on moving/renaming pages between spaces	Closed	Fixed
	CONF-3220	Document the dangers of turning external user management off	Resolved	Fixed

	CONF-4793	Link to news item	Closed	Answered
	CONF-3658	Create axis soap service alongside glue service	Resolved	Fixed
	CONF-3989	AttachmentViewEvent	Resolved	Fixed
	CONF-3411	login page should display alternative content if user is already logged in	Resolved	Duplicate
	CONF-3973	Logo doesnt link to user's "home" - instead hard-links to dashboard	Resolved	Fixed
	CONF-2132	Include recent changes as a macro	Closed	Fixed
	CONF-2999	Fix release process	Resolved	Fixed
	CONF-3895	When editing an existing page, clicking the Choose Page image to select a Parent Page causes a Javascript error in both IE and Mozilla	Resolved	Fixed
	CONF-4015	Component plugins are not removed from the spring context when disabled.	Resolved	Fixed
	CONF-3905	Colour of Breadcrumb text cannot be configured	Closed	Fixed
	CONF-3988	Missing #comments and #addcomment anchors when viewing page	Resolved	Fixed
	CONF-3664	Deactivate user link does not display when viewing a user's profile	Closed	Fixed
	CONF-3838	Export Space to XML may produce "Too many open files" error	Closed	Fixed
	CONF-4107	Page titles containing ? may not be linked to on some application servers	Closed	Fixed
	CONF-3969	Welcome greeting color not using colour scheme colours	Resolved	Fixed
	CONF-4000	View system info shows "1 , 1 hour, 55 minutes, 53 seconds" as the system uptime.	Resolved	Fixed
	CONF-3971	Layout off on Themes page	Resolved	Fixed
	CONF-3800	Start watching this space	Closed	Fixed
	CONF-3843	includePage macro does not work when called from an inline decorator	Closed	Fixed
	CONF-3567	Umlauts and page-titles - few work, many don't.	Closed	Fixed
	CONF-3855	DuplicateNotificationsCleanupUpgradeTask breaks on MySQL 3	Resolved	Fixed
	CONF-3884	Could not execute action	Closed	Answered
	CONF-3383	[file:// links break with \	Resolved	Won't Fix
	CONF-3878	Uploading macros via web interface is broken	Resolved	Fixed
	CONF-3805	MAILTO Links on PDF export are broken	Closed	Duplicate
	CONF-3758	%span% markup should be removed from notation guide.	Closed	Fixed
	CONF-3893	Email containing no body generates exception.		Fixed

			Closed	
	CONF-3610	Exported space (pdf or html) always has blank 'Available Pages' section	Closed	Fixed
	CONF-3288	JavaScript error in IE when re-editing a blog post	Resolved	Fixed
	CONF-3325	Attachment and Anchor links appear on undefined pages list	Closed	Fixed
	CONF-3059	User Management Delegated to JIRA	Resolved	Fixed
	CONF-3781	HTTPS Links on PDF export are broken	Resolved	Fixed
	CONF-3803	News and Comments	Closed	Fixed
	CONF-3790	Broken links to pages with non-ASCII titles	Resolved	Fixed
	CONF-3797	Resoruces not loading for plug-ins	Resolved	Fixed
	CONF-3849	Email and News updated	Closed	Fixed
	CONF-4024	SQL run by DuplicateNotificationsCleanupUpgradeTask gives syntax error on MySQL	Resolved	Duplicate
	CONF-3997	Job DEFAULT.indexQueueFlushJob threw an unhandled Exception	Closed	Fixed
	CONF-3970	wrong icon on Plain Website theme	Closed	Fixed
	CONF-3860	Broken "attached" link in Confluence demo space	Closed	Fixed
	CONF-3858	Unnecessary db hits in ViewPageAction	Closed	Fixed
	CONF-4782	Missing Apostrophe in Page Permission Error Message	Closed	Duplicate
	CONF-2780	Missing database indexes and slow performance	Resolved	Fixed
	CONF-3184	XWork actions in uploaded plugins do not work	Resolved	Fixed
	CONF-3886	XMLRPC API storePage ignores parentId=0	Resolved	Fixed

Release Notes 2.0

Atlassian Software is proud to present Confluence 2.0 (otherwise known as Yarra). Yarra is the result of five months of solid work by the Confluence team, and we're really glad to be able, finally, to share it with the world. Existing customers who wish to upgrade, or new users who wish to try out Confluence for 30 days, can download Confluence from the Atlassian website: http://www.atlassian.com/software/confluence

Yarra is the fifth major update to Confluence. Among the improvements in Confluence 2.0 are an easy-to-use WYSIWYG editor for writing pages, labels for categorising them, and a powerful RSS builder for keeping track of what's new.

Confluence 2.0 is a free upgrade for any customer who purchased their Confluence license after November 16th, 2004. If the maintenance period of your license has expired, or is about to expire, why not contact our friendly sales staff and get it renewed? It's the only way to keep up with all the great new features we're adding.

A big thanks to everyone who reported bugs and offered suggestions over the last few months, especially everyone who helped by trying out our Development Releases. Also, congratulations to the Socceroos for getting Australia into the World Cup for the first time in 32 years. It *almost* makes up for losing the Ashes.

See also: Issues Resolved for 2.0

Contents

- Upgrading
- Migration
- New Features
- Notable Bug Fixes
- Outstanding Bugs

Upgrading From a Previous Version of Confluence

Upgrading Confluence should be pretty easy. *We strongly recommend that you backup your confluence.home directory and database before upgrading!*

Upgrades from 1.4.2 and earlier
If you are upgrading from Confluence 1.4.2 or earlier, you may experience problems with some space-related functionality (see CONF-4765), such as adding a space as a favourite. To work around this, please restart the Confluence instance after the upgrade. This will be fixed in Confluence 2.0.2.

Important Migration Notes

Macro/Plugin Compatibility
Necessary changes were made to the Confluence rendering subsystem during the development of Confluence 2.0 that may render some third-party plugins (especially macros) inoperable. If you upgrade Confluence and find that macros or plugins are not operating correctly, try removing *all* files from `[confluence-home]/plugins` and restarting Confluence.

HSQL 1.8 Upgrade
If you are using the embedded HSQL database, it is possible that Confluence will not be able to automatically upgrade your data. If this happens, Confluence 2.0 will refuse to start, and you will be directed to the following Confluence page which contains instructions on how to upgrade the database manually: Upgrading From HSQL 1.7.1 to 1.8

SSO update
If you are using some third-party Seraph authenticator with Confluence, or have written your own, you should read CONF-4581 before upgrading. Confluence now uses the Seraph defined login.link.url property to define its login link urls, which may cause issues with authenticators that relied on Confluence's previous, incorrect behaviour.

Weblogic Performance
Confluence 2.0 may perform very badly under Weblogic. There is a workaround for this problem described in CONF-4634, and a full fix IS included in 2.0.1.

JDK 1.5
If you are running Confluence 2.0 on the JDK 1.5, you will need to download some additional dependencies as described in CONF-4643. A full fix is included in a 2.0.1.

MySQL 5
If you are running Confluence 2.0 on the MySQL 5 database, you may encounter some problems. One fix is described here .

Upgrading from 1.4.4

After upgrading to 2.0, administrators will need to rebuild the site's search-index to ensure all the new search features are enabled. Do this from the *Content Indexing* section of the global administration menu.

Upgrading from 1.3.5 or Earlier

Users upgrading from an earlier version of Confluence should check the release-notes of the other major Confluence releases:

- Release Notes 1.4
- Release Notes 1.3
- Release Notes 1.2
- Release Notes 1.1

New Features

The four major new features in Confluence 2.0 are:

- Rich Text Editing
- Labels for content
- New Dashboard Features for managing sites with large numbers of spaces.
- A dynamic RSS Builder

..but there's a lot more on top of that.

Rich Text Editing

> **Browser Compatibility**
> The Confluence Rich Text editor is currently only compatible with Internet Explorer 6 on Windows, plus Mozilla and Firefox across platforms. Javascript must be enabled in the browser for the editor to function. Support for Safari under Mac OS X is currently not available. To track Safari compatibility, please follow this JIRA issue: CONF-3864

It almost goes without saying that the most highly requested feature in Confluence has been the ability to create pages without having to learn wiki markup. We're glad we can finally offer a powerful "what you see is (pretty much) what you get" rich text editor built into Confluence, making it easier for anyone to contribute to the site.



The WYSIWYG editor is enabled when you install or upgrade to Confluence 2.0. Global Administrators can disable the editor if they want to stick with pure wiki markup, and can also choose which editor users should be presented with by default. (The setting is under "General Configuration" in the administrative console). Users can also choose which editor they prefer simply by clicking on the "Make this my default editor" link that appears on the edit screen.

For the "feature mad" amongst us, here are some neat things you can do with the WYSIWYG editor:

- Full screen view - really useful for editing large pages. Click in the menu bar.
- Quickly switch between WYSIWYG and Wiki markup without a page refresh
- Change the size of your editing window to suit your browser. Your size preference is remembered across sessions. To change it, drag the handle in the bottom right hand corner of the editor.
- Undo and redo!

Labels

Another highly requested feature was the ability to categorise content within Confluence beyond the rigid heirarchy allowed by spaces and parent-child relationships between pages. To this end we have introduced labels: simple one-word 'tags' that can be added to any page or blog-post the user has permission to edit. Labels can be used to categorise content, bookmark it, flag it for attention, or anything else you can think of.

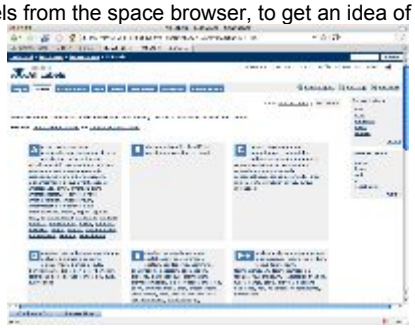
> **A Tag By Any Other Name**
> Picking a name for labels wasn't easy. Google's GMail service calls them labels, while other collaborative categorisation systems such as del.icio.us and Flickr call them tags. We decided that 'label' was a more natural description. For more information about our choice, see Labels vs Keywordson our new Developer Blog and for more on the overall philosophy behind labels and tags, check out the Folksonomy page on Wikipedia.



Labels can be added to any page from the edit screen, or through a dynamic interface right when you're viewing a page.



Once a page is labeled, then clicking on the label's name allows you to browse other pages with the same label, or view related labels that commonly occur on the same pages.
You can also view the space's most popular labels from the space browser, to get an idea of the most popular topics within the space.



Personal Labels
If you prepend `my:` to a label (for example, `my:todo` or `my:favourite`, then the label is a personal label - only visible to you. Personal labels allow you to tag content for your own purposes: for example to keep track of pages you feel need your attention, or that contain information you refer to frequently. You can browse your personal labels from your user profile. Any user can add their personal labels to any page, even when they don't have editing permission.



Favourites

Favourites are a special personal label: `my:favourite` or `my:favorite`. Whenever you see the icon, it means you can label this content as being your favourite, and whenever you see the icon, it means that the content is currently in your list of favourites. You can view your favourites from the Labels tab of your user profile, or keep track of them on your dashboard.
Label-Aware Macros
Many existing macros have been improved to allow you to filter content based on labels: including the {recent-pages}, {recently-updated} and {blog-posts} macros. We've also added macros that provide more information about labels and labelled content:

- {related-labels} gives a list of labels that might be related to a page
- {listlabels} lists all the labels in a space
- {contentbylabel} lists content that has a particular label
- {recently-used-labels} lists labels that have been recently added or applied
- {navmap} draws a nice-looking table of links to pages with a particular label

Dashboard Features

The Confluence dashboard has been improved to make it easier for you to keep track of only those spaces you are interested in: a big improvement for Confluence sites with large numbers of spaces. The list of spaces is now divided into four tabs (although all four may not be visible):

- **All** shows you all spaces
- **My** shows you all spaces you have marked as your favourites (you can mark a space as your favourite from the *All* tab by clicking on the icon)
- **Team** shows you all the available 'teams', and the spaces that have been assigned to them
- **New** shows you any space that has been added in the last week

The recently updated content list on the dashboard will reflect the spaces in your chosen tab. So if you're looking at the **My** tab, the dashboard will only be showing you the recent updates in your favourite spaces.

Teams!dashboard-team-tab.png|align=center!
Teams are a simple, and very wiki-like way to group spaces together. Space administrators can add "team labels" to a space, which are then

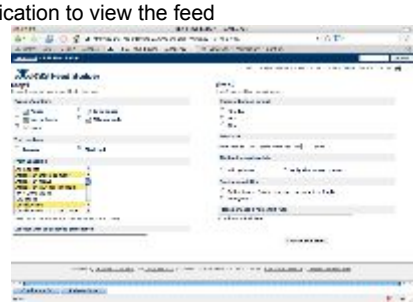
used to group those spaces under the team tab on the dashboard. So if your wiki has 100 spaces, but only five of them are of any interest to your sales team, just add a "sales" team label to those five spaces. That will group those spaces together on the dashboard under the 'team' tab, and your sales team need never look at the other 95 spaces.

RSS Builder

Confluence has always provided a brace of useful RSS feeds, but the problem is that for every feed we provided, users asked for half a dozen more. The obvious answer is to let users build RSS feeds based on their own chosen criteria. You can access the RSS builder from the Confluence dashboard

Once in the builder, you can choose

- Which spaces to include in the feed
- Which types of content should be tracked
- Which labels, if any, you are interested in
- How many items to include in the feed
- Whether you want a single RSS entry per page or one for each time the page is edited
- Whether you want an RSS 2.0 or Atom 0.3 feed
- Whether Confluence should require authentication to view the feed



Once you have decided what you want, Confluence will give you a URL to paste into your RSS reader. These URLs can be shared with other Confluence users, although they will only ever be allowed to see content that they have permission to view. If you have asked to authenticate, Confluence will require HTTP Basic Authentication, which is supported by most RSS readers.

We've also taken the opportunity to improve the presentation of our RSS feeds - including a lot more information in each feed so you can follow your Confluence site entirely from your newsreader.

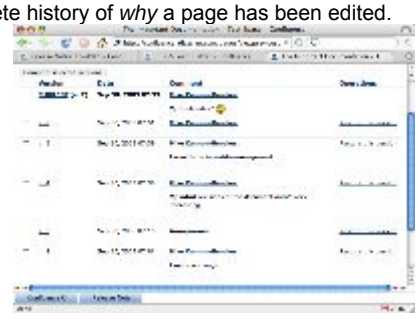


The Atom Working Group has blessed the final Atom 1.0 standard, and are recommending that all applications remove support for Atom 0.3. However, as of the release of Confluence 2.0, there is no stable Atom 1.0 library from which we could generate feeds. In some future release of Confluence, Atom 0.3 support will be removed in favour of Atom 1.0.

Other New Features

Record "Change Comments" When Editing a Page

There is now a field on the edit screen for recording a "change comment" when you edit a page. These comments are stored in the page history, and can be used to keep a more complete history of *why* a page has been edited.



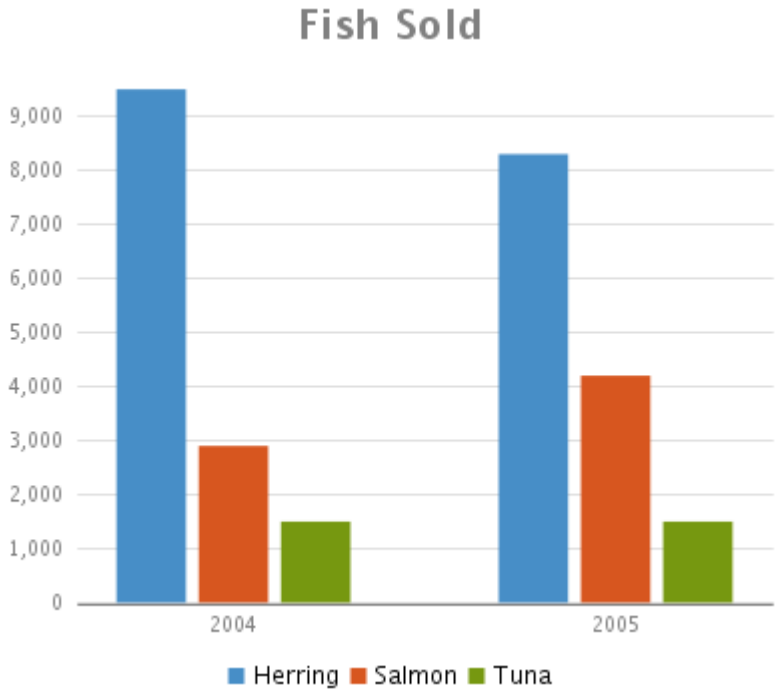
Embed Flash and Movies
You can now embed Flash content or movies (Quicktime or Windows Media) into a page as easily as you can an image: just attach the Flash or movie file to the page, then include it as you would include an image (!filename.mov!).
Export Pages as Word Documents
You can now export pages straight into Word from the Info tab. This is extremely useful for emailing around content to non-Confluence users, printing a document or just creating a backup in Word.



Copy Pages
Also on the Info tab is a "Copy" link that allows you to clone a page in a single click - including making copies of any attachments.
Improved Search Interface
Results returned from Confluence's search engine now have:

- Improved contextual results, showing the most important text around where your query was matched in the page
- Contextual results for any attachment: see where a search was matched even **inside** PDF, Word, PowerPoint or Excel documents!
- Search results for attachments give you more (and clearer) information about what the attachment is, and where it's from!search-result-fragments.png|align=center!
  Chart Plugin

The Chart Macro is now shipped with Confluence, allowing you to dynamically generate neat looking charts like this:



Wiki Source

```
{chart:title=Fish Sold|type=bar|width=400|height=350|legend=true}
|| Fish Type || 2004 || 2005 ||
|| Herring | 9,500 | 8,300 |
|| Salmon | 2,900 | 4,200 |
|| Tuna | 1,500 | 1,500 |
{chart}
```

Improved Gallery Macro

The gallery macro has been spruced up, and now has a slideshow view:

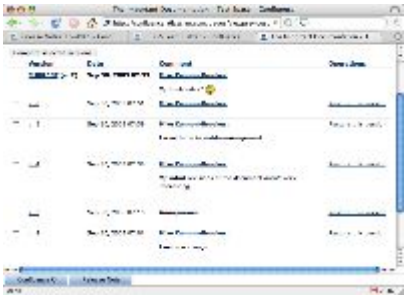
Confluence 2.0 Screenshots





My personal labels! Only for me!
Popular labels for a given space.
Alphabetically organised labels for a given space.





Some release notes exported to Word (even on a Mac!) - great for printing too.
The new Dashboard space tabs, here showing the "jira" team tab and one favourite space.
A very basic (quite boring - sorry, it's late) example of change summaries.





Screenshot of the new 2.0 rich text editor
Improved search showing fragments from attachments, file sizes and types.
The new RSS builder lets you construct RSS feeds of exactly the content you want.



View page label interface showing autocomplete.

Additions to the Remote API

Additions to the Confluence Remote API include:

- Comment manipulation
- Label manipulation
- Attachment uploading and editing
- Improved user- and permissions management

> ⚠ Confluence now uses version 2.0 of the Apache XML-RPC library. Java XML-RPC clients using earlier versions of the Apache XML-RPC libraries (i.e. 1.3 or earlier) may experience problems with responses containing non-ASCII data.

Also...

- You can download all the attachments on a page in a single zip-file
- Import and restore now have progress indicators
- Backup and restore use significantly less memory
- The embedded database has been upgraded to HSQL 1.8, which should be significantly more reliable

- Collapsed breadcrumbs now expand with a single mouse click

Notable Bug Fixes

We resolved a *lot* of issues between Confluence 1.4.4 and Confluence 2.0. The best way to see what we've fixed is to ask JIRA, the world's best issue-tracker: Issues Resolved for 2.0

Outstanding Bugs

Some bugs were introduced during the Confluence 2.0 development cycle that we could not fix in time for the final release. Of note are:

- Some pernicious Javascript errors when changing styles in the rich text editor under Internet Explorer
- The rich text editor may not perfectly handle complex pages with structural macros
- Attachments containing high-bit characters in their filenames may not be correctly retrieved by Confluence

Once again, if you find any bugs in Confluence, or have any feature suggestions, you can report them online in JIRA.

The Confluence 2.0 Team

Development
Tom Davies
Jeremy Higgs
David Loeng
Charles Miller
Daniel Ostermeier
Jens Schumacher

Documentation
Vidya Madabushi

Oversight & ~~Mis~~ ~~management~~
Mike Cannon-Brookes
Scott Farquhar

Nerf Target-Practice
Nick Faiz

Well that's all folks - if you're still reading - thank you for getting this far!

> To keep up with all the latest developments in the next 2.1 release (codenamed: Bogan) - subscribe to our developer blog.

Issues Resolved for 2.0

Error formatting macro: jiraissues: java.lang.RuntimeException: A value with ID '11191' does not exist for the field 'project'.

Release Notes 2.0.1

Confluence 2.0.1 is a maintenance release that resolves some issues users may have encountered using Confluence 2.0. In particular, this includes issues relating to failures to upgrade and javascript problems.

2.0.1 is a free upgrade for all customers who purchased their Confluence license after November 28th, 2004.

Who should upgrade?

Confluence 2.0.1 is a recommended upgrade for all users who have not yet upgraded to 2.x and for those users of 2.0 that are encountering issues fixed in this release.

Upgrade Procedure

Upgrading Confluence should be pretty easy. If you are upgrading from Confluence 1.4, you can find instructions here. *We strongly recommend that you backup your `confluence.home` directory and database before upgrading!*

If you are upgrading from Confluence 1.4.4 or earlier of Confluence, please check the release-notes of the other major Confluence releases:

- Release Notes 2.0
- Release Notes 1.4
- Release Notes 1.3
- Release Notes 1.2
- Release Notes 1.1

Changes in 2.0.1

See also: Issues resolved for 2.0.1

- Javascript failing with a class not found. (CONF-4643)
- Confluence 2.0 hangs on weblogic. (CONF-4634)
- Disabling WYSIWYG editor causes AJAX error in Preview mode. (CONF-4745)
- Ancestors table hangs on to foreign key relationships if it can't be deleted. (CONF-4700)

Issues resolved for 2.0.1

Errors were reported by the JIRA trusted connection.

- APP_UNKNOWN; Unknown Application: {0}; ["confluence:4557196"]

JIRA Issues (42 issues)				
Type	**Key**	**Summary**	**Status**	**Resolution**
	CONF-4612	Add RSS Autodiscovery to the results pages of the feed builder	Resolved	Fixed
	CONF-4292	Sort page tree pages alphabetically	Resolved	Fixed
	CONF-4678	Provide users with license screen to enter a new license if their existing one has expired	Closed	Fixed
	CONF-4748	Mail importer should strip out illegal characters in attachment filenames	Closed	Fixed
	CONF-4598	replace PermissionCheckDispatcher with PermissionHelper	Resolved	Duplicate
	CONF-4642	Login page shows 'You do not have permission to access null'. null??	Resolved	Fixed
	CONF-4653	{blog-posts: x} returns Error while trying to draw the last-n pages!	Resolved	Fixed
	CONF-4646	global-reports macro doesn't render properly	Resolved	Fixed
	CONF-4659	No body macros are duplicated when switching between WIKI and WYSIWYG editing.	Closed	Cannot Reproduce
	CONF-4680	Cycling between WIKI and WYSIWYG shows out of date data.	Resolved	Fixed
	CONF-4665	Space list drop down listbox in Link selector has empty string options	Closed	Fixed
	CONF-4689	storeBlogEntry method via XML-RPC throws invalid argument exception	Closed	Fixed
	CONF-4699	Remove WYSIWYG checkbox from the user profile.	Resolved	Fixed
	CONF-4654	Exception while getting property calendarForThisMonth	Resolved	Fixed
	CONF-4749	Line breaks missing for the display of code macros in exported pdf pages	Closed	Fixed
	CONF-4533	Macros with forms interfere with "Save" and "Cancel" buttons in WYSIWYG editor	Resolved	Fixed
	CONF-4578	blog-posts macro only shows news items from current month	Resolved	Fixed
	CONF-4639	Image popup window wrong size?	Resolved	Fixed
	CONF-4737	Newlines being stored as version comments	Resolved	Fixed
	CONF-4698	Confluence claims that a restore has been completed even though the restore is still running in the background.	Resolved	Fixed
	CONF-4700	Ancestors table hangs on to foreign key relationships if it can't be deleted	Resolved	Fixed

	CONF-4768	Exception thrown while accessing "News"	Closed	Duplicate
	CONF-4745	Disabling WYSIWYG editor causes AJAX error in Preview mode	Resolved	Fixed
	CONF-4643	Basic functions not working on JDK 1.5	Resolved	Fixed
	CONF-4631	Ampersand not handled correctly when switch between renderers	Resolved	Fixed
	CONF-4641	Strange screen if you delete a comment twice	Resolved	Fixed
	CONF-4634	Confluence 2.x hangs on weblogic.	Resolved	Fixed
	CONF-4649	Update the chart macro to the latest version (1.2)	Resolved	Fixed
	CONF-4667	{blog-posts:x} macro displays only news items that were created or edited after upgrading to v2.0	Resolved	Fixed
	CONF-4666	viewrecentblogposts.action no longer displays calendar for navigation.	Closed	Duplicate
	CONF-4668	Switching between Rich Text and Wiki Markup tabs sucks links that follow tables into the table itself	Resolved	Fixed
	CONF-4661	blog-posts macro time parameter doesn't work as expected	Resolved	Fixed
	CONF-4561	RSS feed for pages marked "favourite"	Resolved	Fixed
	CONF-4487	Dynamic search and dynamic tasks macros don't work well placed together	Resolved	Fixed
	CONF-4720	Duplicate Notification upgrade task is failing.	Resolved	Fixed
	CONF-4650	WYSIWYG bug with inner phrases	Resolved	Fixed
	CONF-4719	1.3.5 backup can't be restored to 2.0 at config time, but can be restored later.	Resolved	Fixed
	CONF-4388	Empty headings should contain a half space, so they can be selected	Resolved	Fixed
	CONF-4714	Sync SOAP / XML-RPC API for addAttachment.	Resolved	Fixed
	CONF-4702	Anonymous users are offered link to edit 'Space Labels' in space where permissions for this space are restricted to 'view' only. Clicking the 'edit' link throws exception	Resolved	Fixed
	CONF-4682	NullPointerException when saving a page with content pasted from MS Word	Resolved	Fixed
	CONF-4526	Do not remove blank lines after headings when switching between WYSIWYG and RichText.	Resolved	Fixed

Release Notes 2.0.2

Confluence 2.0.2 is a maintenance release that resolves a security issue, along with various issues users may have encountered using Confluence 2.0.1.

A security flaw as described by the Confluence Security Advisory 2005-12-05 has been identified to exist in Confluence 1.4.x and 2.0.x. This has been fixed in 2.0.2. We recommend to all customers that they either upgrade to 2.0.2 or follow the instructions provided on the Confluence Security Advisory 2005-12-05 to patch there installation.



The release of Firefox 1.5 resulted in new bugs and issues with the Rich Text editor. Therefore every version of Confluence up to 2.0.2 isn't very compatible with this browser in terms of Rich Text editing.

Confluence 2.0.3 will feature an updated version of the editor which will solve most of the problems with Firefox 1.5 (CONF-4809).

2.0.2 is a free upgrade for all customers who purchased their Confluence license after December 5th, 2004.

Who should upgrade?

Confluence 2.0.2 is a recommended upgrade for all users as it contains a security patch for the Confluence Security Advisory 2005-12-05. If you are unable to upgrade to 2.0.2, then please see Confluence Security Advisory 2005-12-05 for details on how to patch your installation.

Upgrade Procedure

Upgrading Confluence should be pretty easy. If you are upgrading from Confluence 2.0, you can find instructions here. *We strongly recommend that you backup your `confluence.home` directory and database before upgrading!*

If you are upgrading from Confluence 1.4.4 or earlier of Confluence, please check the release-notes of the other major Confluence releases:

- Release Notes 2.0
- Release Notes 1.4
- Release Notes 1.3

Changes in 2.0.2

See also: Issues resolved for 2.0.2

- Search results page needs to XML encode the query string provided by the user (CONF-4825).
- Recently updated does not list any items (CONF-4770).
- CamelCase linking interferes with rendering of attachment link (CONF-3447).
- Umlaute in links are causing encoding problems with the Rich Text editor (CONF-4775).
- Disabled accounts still receive (blank) daily reports (CONF-4802).

Issues resolved for 2.0.2

Errors were reported by the JIRA trusted connection.

- APP_UNKNOWN; Unknown Application: {0}; ["confluence:4557196"]

JIRA Issues (21 issues)				
Type	**Key**	**Summary**	**Status**	**Resolution**
	CONF-4672	NullPointerException when adding space as favourite	Resolved	Answered
	CONF-4810	Make a note in the 2.0.2 release notes about problems with WYSIWYG editor on Firefox 1.5 for the Mac	Closed	Fixed
	CONF-4151	Allow sorting of attachments	Resolved	Fixed
	CONF-3272	Allow attachments to be ordered by date, size and name in the attachments macro	Resolved	Fixed
	CONF-4786	Need to support UNICODE characters in MSSQL server.	Closed	Fixed
	CONF-2790	Images in historical versions of pages are not rendered	Resolved	Fixed
	CONF-4770	recently updated does not list any items	Closed	Resolved Locally
	CONF-3447	CamelCase linking interferes with rendering of attachment link	Resolved	Fixed
	CONF-3923	CamelCase linking breaking image rendering if image titled "ThisIsAnImage.jpg"	Resolved	Fixed
	CONF-4771	editor does not work in firefox 1.5	Closed	Handled by Support

	CONF-4775	Umlaute in links are causing encoding problems with the Rich Text editor	Resolved	Fixed
	CONF-4799	Extra "<" in labels-include.vm causes problems with page loading	Resolved	Fixed
	CONF-4804	Contentbylabel macro separates labels by spaces rather than commas	Resolved	Fixed
	CONF-4825	Search results page needs to XML encode the query string provided by the user.	Resolved	Fixed
	CONF-4766	LazyInitializationException in ResourceUpgradeTask	Closed	Fixed
	CONF-4750	Contentbylabel macro does not work if there is no space separating label names	Resolved	Fixed
	CONF-4469	Favorites Pages shows only 5 page!!	Resolved	Fixed
	CONF-4807	Cannot sort attachments by name	Resolved	Fixed
	CONF-4802	Disabled accounts still receive (blank) daily reports	Closed	Fixed
	CONF-4776	Upgrade tasks do not flush the cache after updating content	Resolved	Fixed
	CONF-4765	Space Descriptions may not have their Space ID set, causing problems when the space is referenced	Resolved	Fixed

Release Notes 2.0.3

Confluence 2.0.3 is a maintenance release. It contains a upgrade of the WYSIWYG editor and numerous bug fixes.

2.0.3 is a free upgrade for all customers who purchased their Confluence license after December 12th, 2004.

Who should upgrade?

Confluence 2.0.3 is a recommended upgrade for all users who are having problems with the WYSIWYG editor. This release contains an upgrade to the WYSIWYG module that contains numerous bug fixes.

Upgrade Procedure

Upgrading Confluence should be pretty easy. If you are upgrading from Confluence 2.0, you can find instructions here. *We strongly recommend that you backup your `confluence.home` directory and database before upgrading!*

If you are upgrading from Confluence 1.4.4 or earlier of Confluence, please check the release-notes of the other major Confluence releases:

- Release Notes 2.0
- Release Notes 1.4
- Release Notes 1.3

Changes in 2.0.3

See also: Issues resolved for 2.0.3

- Upgrade to TinyMCE 2.0.1 (CONF-4808)
- Backup restore progress monitor does not refresh during restore, making it look like the backup is hanging (CONF-4895)
- Information leak when accessing url directly (CONF-4794)
- Personal labels showing up on the {contentbylabel} macro (CONF-4894)

Issues resolved for 2.0.3

Errors were reported by the JIRA trusted connection.

- APP_UNKNOWN; Unknown Application: {0}; ["confluence:4557196"]

JIRA Issues (23 issues)				
Type	**Key**	**Summary**	**Status**	**Resolution**
	CONF-4880	Wrong SQL-Statement in JiraJdbcProfileProvider?	Resolved	Fixed

	CONF-4875	Modify Standalone Tomcat configuration to stop NotSerializableException	Resolved	Fixed
	CONF-4819	Increase size of search box	Resolved	Fixed
	CONF-4893	Hide "Add Content" Links in Left Navigation Theme if user does not have permissions to add any kind of content.	Resolved	Fixed
	CONF-4846	Update jfreechart to 1.0.0	Resolved	Fixed
	CONF-4837	RSS feedbuilder link hidded when Tableless Theme selected	Resolved	Fixed
	CONF-4835	Warn if page exists even when user lacks view/edit page permission	Resolved	Fixed
	CONF-4510	Selecting text colour in WYSIWYG causes previous text to change colour	Resolved	Fixed
	CONF-4894	Personal labels showing up on the {contentbylabel} macro.	Resolved	Fixed
	CONF-4849	Image Resource Problems - working in 1.4.4 and broken in 2.0.2	Resolved	Fixed
	CONF-4895	Backup restore progress monitor does not refresh during restore, making it look like the backup is hanging.	Resolved	Fixed
	CONF-4794	Information leak when accessing url directly	Resolved	Fixed
	CONF-4822	{navmap} doesn't respect server base url	Resolved	Fixed
	CONF-4809	Rich Text editor doesn't display the WYSIWYG content the first time the edit page is loaded (Firefox 1.5 MAC)	Resolved	Fixed
	CONF-4808	Upgrade to TinyMCE 2.0.1	Resolved	Fixed
	CONF-4824	Switching quickly between Rich Text and WikiMarkup / Preview causes content to disappear	Resolved	Fixed
	CONF-4843	Object and Embedd tag parameters for embedded objects in a page are missing quotes.	Resolved	Fixed
	CONF-4752	Parent page lost when creating page with template	Resolved	Fixed
	CONF-4891	Remove duplicate results from the {contentbylabel} macro.	Resolved	Fixed
	CONF-4889	"Citation" markup is not converted between Rich Text and Markup modes	Resolved	Fixed
	CONF-3604	java.io.NotSerializableException: org.apache.log4j.Logger	Resolved	Fixed
	CONF-4820	Cannot apply strikethrough to links in rich text editor	Resolved	Fixed
	CONF-4685	Fullscreen button freezes IE and reports error in Firefox	Resolved	Fixed

Release Notes 2.1

Atlassian is happy to offer our customers an early Christmas present: Confluence 2.1 (otherwise known as Bogan). Existing customers who wish to upgrade, or new users who wish to try out Confluence for 30 days, can download Confluence from the Atlassian website: http://www.atlassian.com/software/confluence

Bogan is the Sixth major update to Confluence, offering vastly improved LDAP integration through our new atlassian-user library, as well as introducing the much-requested autosave feature to protect you from losing your precious edits.

Upgrading from 2.0

Upgrading Confluence should be pretty easy: you can find instructions here. *We strongly recommend that you backup your `confluence.home` directory and database before upgrading!*

Plugin Compatibility
Any plugin written for Confluence 2.0 and earlier that interfaces with the User system will need to be updated for Confluence 2.1. Plugin developers can find more information on the Atlassian developer blog

LDAP Integration Configuration Changes
Customers who have already integrated Confluence with LDAP through the OSUser LDAP providers will need to make changes to their `osuser.xml` file before upgrading. You can find full details in step 3 of LDAP Authentication with OSUser.

OSUSer or Seraph Customisations
Customers who have performed their own customisations on OSUser or Seraph within Confluence **must** be sure to test their changes with Confluence 2.1 before upgrading any production system. While we have tried to maintain backwards compatibility, the integration of atlassian-user may adversely affect existing customisations.

Upgrading from 1.4 and earlier

Users upgrading directly from 1.4 or earlier should also read the 2.0 Release Notes for caveats regarding the 1.4 -> 2.0 upgrade.

Contents

1. New Features
2. Improvements
3. Notable Bug-fixes

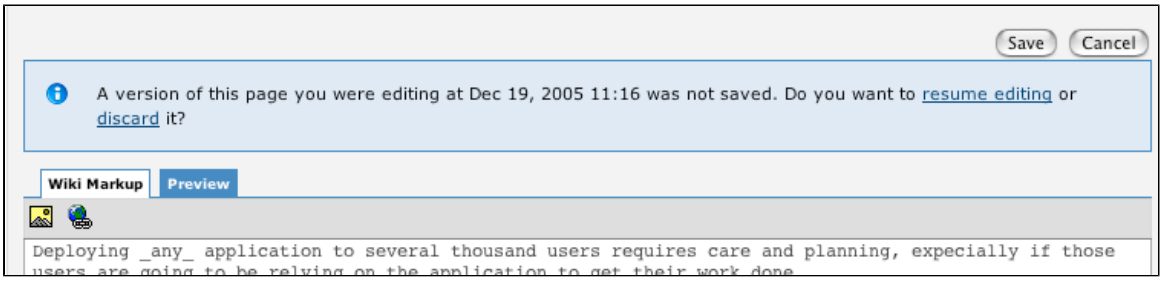
See also: Issues Resolved for 2.1

New Features

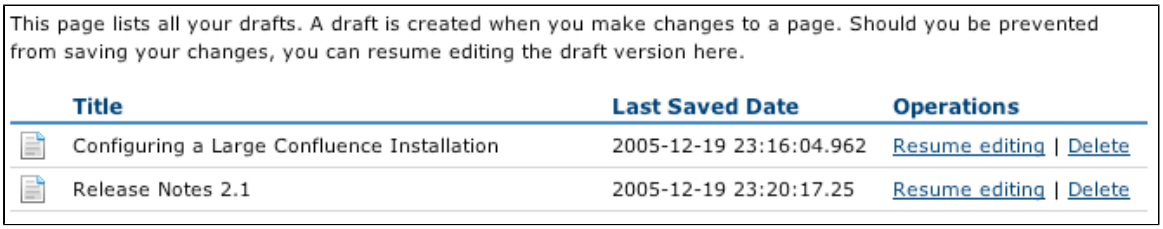
Autosave

If you've used wikis for long enough, you know the pain of losing a long, involved editing session to a browser crash, session timeout, or just not thinking and clicking on a link to another page. As of Confluence 2.1, this is no longer a problem: the most you'll lose is half a minute of your work.

Every thirty seconds (this interval is configurable by the global administrator), the browser will save your unfinished edit to the server. This autosave will hang around until you either hit the "Save" or "Cancel" button yourself, even if the server is restarted. So if you lose your changes, all you have to do is navigate back to the edit page, and will be given the chance to begin editing again where you left off.

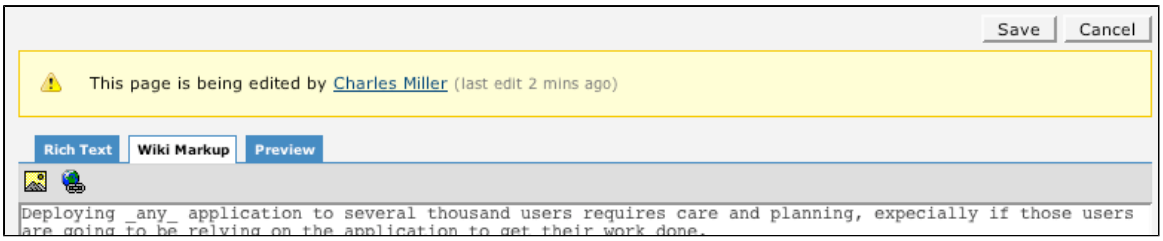


You can view your outstanding autosaved documents from your profile under the "Drafts" tab.



Concurrent Edit Warnings

A useful side-effect of autosave is that now Confluence has an accurate way of measuring who is editing which page. If two people start editing the same page at the same time, Confluence will display a warning message telling you who else is editing the page. You can then negotiate between yourselves who gets to save first.



Even better, Confluence will let you know how long it was since the other editor made any changes to the page, so if somebody has just left their browser open on the edit page for a couple of hours, you'll know you can sneak your changes in while they're asleep.

User Management

Confluence 2.1 sees the landing of our new `atlassian-user` user management library. Most Confluence customers will not see any significant change from this move, but for anyone looking to integrate Confluence with an external user-base (especially LDAP), atlassian-user is a big step forward.

Confluence can now integrate fully with an LDAP directory server, without the previously annoying practice of having to mirror the users and groups locally on Confluence: Confluence LDAP Documentation Index

Improvements

Performance

A lot of work was done improving the performance of the Confluence dashboard and edit pages, especially for customers who may have thousands of spaces and hundreds of thousands of pages in their Confluence installation. Both of these pages should now respond significantly faster.

Other

- You can now manually set the MIME type under which embedded objects should be served – CONF-4906
- The system info and error pages include more information about your database configuration – CONF-4957
- ~~Email attachments no longer show up on the recently updated list~~ (this seems to have regressed in the final release) – CONF-4684
- Added an icon mapping for sub-tasks in the jiraissues macro – CONF-4921
- Allow attached Windows Bitmap files to be displayed as embedded resources – CONF-4922

Notable Bug-fixes

Confluence 2.1 includes all bug-fixes that were made up to Confluence 2.0.3, and also includes fixes for the following issues:

- Preview now works correctly when rich text editor is disabled – CONF-4935
- Info tab no longer shows incoming links from pages in the trash – CONF-4815
- Trying to set a page as its own parent now presents a validation error instead of a system error – CONF-4941
- Confluence no longer removes every second character from the filenames of email attachments – CONF-4938
- JiraJdbcProfileProvider problems reliably accessing profile information from JIRA resolved – CONF-4933
- Unresolved images now replaced with placeholders in WYSIWYG editor – CONF-4929
- Removing the {excerpt} macro from a page now deletes the excerpt – CONF-4918

The Confluence 2.1 Team

Development
Tom Davies
Jeremy Higgs
David Loeng
Charles Miller
Daniel Ostermeier
Jens Schumacher

Atlassian-User
Nick Faiz

Documentation
Vidya Madabushi

Oversight & ~~Mis management~~
Mike Cannon-Brookes
Scott Farquhar

Issues Resolved for 2.1

Errors were reported by the JIRA trusted connection.

- APP_UNKNOWN; Unknown Application: {0}; ["confluence:4557196"]

JIRA Issues (43 issues)				
Type	**Key**	**Summary**	**Status**	**Resolution**
	CONF-4335	Upgrade task - osuser2hibernate	Closed	Duplicate
	CONF-4227	Replace group select menus with a GroupPicker	Closed	Duplicate
	CONF-4692	Grouppicker should replace select menus of groups	Closed	Fixed
	CONF-4711	Formalize set up procedure r.e. existing users and existing groups in a delegation	Closed	Fixed
	CONF-4859	Broken page table structure in Edit mode after upgrade from 1.4.4 to 2.0.1	Closed	Fixed
	CONF-4693	GroupPicker in edit space perms	Closed	Fixed
	CONF-4696	GroupPicker in global admin permissions	Closed	Fixed
	CONF-4694	GroupPicker in page permissions form	Closed	Fixed
	CONF-1193	Warn when page is concurrently edited by multiple users	Closed	Fixed
	CONF-973	Autosave of editing box	Closed	Fixed
	CONF-4217	Add GMail style auto-save when editing new or existing pages	Resolved	Fixed
	CONF-4233	Create utility to move Entites from OSUser tables into Atlassian User tables.	Closed	Fixed
	CONF-4232	Create new rdbms structure for a Hiberate implementation of Atlassian User	Closed	Fixed
	CONF-4499	SmartListManager performance degrades as the number of groups increases	Closed	Fixed
	CONF-4234	Global Export should only export users in Confluence's database.	Closed	Fixed
	CONF-1806	Merge Feature for Pages that Were Edited at the Same Time	Resolved	Fixed
	CONF-4906	Allow syntax to override the file MIME type in embedded objects	Resolved	Fixed
	CONF-4236	Activate/deactive user now manipulates the user's relation to USE_CONFLUENCE permission	Closed	Fixed
	CONF-4965	Improve dashboard performance	Resolved	Fixed
	CONF-4735	All spaces a user can see should be retrieved with one query	Resolved	Fixed
	CONF-4966	Improve edit page performance	Resolved	Fixed
	CONF-2844	Make attachment MIME type displayed and editable	Closed	Fixed
	CONF-4229	UserPicker should use the new query system and the PagingIterator	Closed	Fixed
	CONF-4688	PagerPaginationSupport must retrieve more results for build more pages	Closed	Fixed
	CONF-4957	Include Database driver information on the "System Info" and error pages	Resolved	Fixed
	CONF-4933	JiraJdbcProfileProvider incorrectly accesses user profile information.	Resolved	Fixed
	CONF-4941	IllegalArgumentException when setting page as its own parent page	Resolved	Fixed
	CONF-4929	Unresolved Image Links are replaced in WYSIWYG mode	Resolved	Fixed
	CONF-4947	Unsaved changes in the rich text editor should disregard whitespace	Resolved	Fixed
	CONF-3629	Format tags in to word	Resolved	Fixed

	CONF-4504	Unable to embed resource	Resolved	Fixed
	CONF-4230	PermissionDelegates are stored via keys based on class names	Closed	Fixed
	CONF-4935	Preview is not working when richeditor is disabled.	Resolved	Fixed
	CONF-4922	Windows bitmap files not recognised by the embedded resource renderer	Resolved	Fixed
	CONF-4918	excerpt in page, still shows in children even after the except macro was removed	Resolved	Fixed
	CONF-4815	Info tab of a page displays incoming links from deleted pages	Resolved	Fixed
	CONF-4763	add a user picker to global perm.s page	Closed	Fixed
	CONF-4967	JIRA + LDAP OSUser integration not supported	Closed	Fixed
	CONF-4760	Edit groups for user brings up a null pointer	Closed	Fixed
	CONF-3396	AJAX doesn't work with Safari v2 on mac	Closed	Fixed
	CONF-4938	Names of mail attachments are wrong	Closed	Fixed
	CONF-4932	Image Linking Broken	Closed	Duplicate
	CONF-4921	Jiraissues does not have an icon mapping for subtasks	Resolved	Fixed

Release Notes 2.1.1

Confluence 2.1.1 is a maintenance release that resolves some issues users may have encountered using Confluence 2.1, including issues related to user management, the editing UI and email notifications.

2.1.1 is a free upgrade for all customers who purchased their Confluence license after December 24th, 2004.

Who should upgrade?

Due to the severity of the issues that it resolves, Confluence 2.1.1 is a recommended upgrade for all Confluence customers.

Upgrade Procedure

Upgrading Confluence should be pretty easy. If you are upgrading from Confluence 2.1, you can find instructions here. *We strongly recommend that you backup your `confluence.home` directory and database before upgrading!*

If you are upgrading from Confluence 2.0.3 or earlier, please check the release-notes of the other major Confluence releases:

- Release Notes 2.1
- Release Notes 2.0
- Release Notes 1.4
- Release Notes 1.3
- Release Notes 1.2
- Release Notes 1.1

Changes in 2.1.1

See also: Issues resolved for 2.1.1

- Users can not change their passwords CONF-5005
- Cursor jumps to the beginning of the rich text editor periodically CONF-4993
- Change comments not being displayed correctly CONF-4979
- Email notifications do not correctly report the user who made the change CONF-4973
- Pages can not be created or saved in some custom atlassian-user or OSUser configurations CONF-4978

Issues resolved for 2.1.1

Errors were reported by the JIRA trusted connection.

- APP_UNKNOWN; Unknown Application: {0}; ["confluence:4557196"]

JIRA Issues (9 issues)				
Type	**Key**	**Summary**	**Status**	**Resolution**

	CONF-4984	Children macro gives bad error message when page not found	Resolved	Fixed
	CONF-4988	JiraJdbcPropertySet is read only	Closed	Fixed
	CONF-4978	Useraccounts are read-only when using LDAP user management	Resolved	Fixed
	CONF-4993	The cursor jump to beginning of first line when you stop typing a few seconds in the wysiwyg editor	Resolved	Fixed
	CONF-5005	Users can not change there passwords.	Resolved	Fixed
	CONF-4979	Comments are being rendered litterally as $page.renderedVersionComment	Resolved	Fixed
	CONF-4974	Concurrent modification exception in SimpleDisplayServlet	Resolved	Fixed
	CONF-4973	Some notifications are not reporting the user	Resolved	Fixed
	CONF-4980	Autosave in rich text resets cursor	Closed	Fixed

Release Notes 2.1.2

Confluence 2.1.2 is a maintenance release that resolves some issues users may have encountered using Confluence 2.1.1 or earlier, including issues related to page editing, Javascript errors in Firefox and Safari and PDF exports.

2.1.2 is a free upgrade for all customers who purchased their Confluence license after January 12th, 2005.

Who should upgrade?

Confluence 2.1.2 fixes a number of bugs found in Confluence 2.1.1 and earlier. Customers are recommended to upgrade if they are experiencing any of the issues fixed in this release.

Upgrade Procedure

> ⚠ If you are using MySQL, please apply the patch detailed in this bug report http://jira.atlassian.com/browse/CONF-5153. That is, please download 2.1.2, extract it, and copy the above mentioned patch into the unpacked 2.1.2 distribution *and then* go about your upgrade. We will be releasing a version with this patch bundled soon.

Upgrading Confluence should be pretty easy. If you are upgrading from Confluence 2.1 or 2.1.1, you can find instructions here. *We strongly recommend that you backup your `confluence.home` directory and database before upgrading!*

If you are upgrading from Confluence 2.0.3 or earlier, please check the release-notes of the other major Confluence releases:

- Release Notes 2.1
- Release Notes 2.0
- Release Notes 1.4
- Release Notes 1.3
- Release Notes 1.2
- Release Notes 1.1

Changes in 2.1.2

See also: Issues resolved for 2.1.2

Improvements

- JiraPortlet macro now has an optional "baseurl" parameter for when Confluence accesses JIRA from a different URL to regular users – CONF-4897
- Children listed at the bottom of pages are now sorted alphabetically – CONF-4878
- Confluence warns space administrators when they permit anonymous access to a space, but global anonymous access is disabled – CONF-4898

Bugs Fixed

- NS_ERROR_NOT_AVAILABLE popup no longer appears when editing certain pages in Firefox – CONF-5038
- "Error converting parameters" popup no longer appears when editing certain pages in Safari – CONF-4976
- Edit page no longer returns NoSuchElementException error under some circumstances – CONF-5007
- "Last week" and "Last month" searches no longer fail when time period spans the new year – CONF-5056
- News items marked as favourites are now listed on the dashboard – CONF-4998

- Mail attachments no longer show up in recent changes on dashboard – CONF-4684
- PDF Export no longer fails with error parsing background-color attributes – CONF-5026
- Anonymous users do not cause a crash when exporting a space – CONF-5129
- Disabling a servlet plugin no longer causes a NullPointerException – CONF-5021
- Notification emails no longer prepend the site's context path twice to certain URLs – CONF-5024
- It is now possible to move a page between spaces and change its parent in the same operation – CONF-5103
- Dashboard favourite selection now works in Opera 8.5 – CONF-5012
- Username links no longer cause a NullPointerException in certain configurations – CONF-5028

Issues resolved for 2.1.2

Errors were reported by the JIRA trusted connection.

- APP_UNKNOWN; Unknown Application: {0}; ["confluence:4557196"]

JIRA Issues (37 issues)				
Type	**Key**	**Summary**	**Status**	**Resolution**
	CONF-5090	Change landing URL for 'get license' link during setup	Resolved	Fixed
	CONF-5017	Add the file.encoding system property to the system info page.	Closed	Fixed
	CONF-5020	List of child pages should be sorted alphabetically	Resolved	Duplicate
	CONF-4898	Warn users when granting anonymous users space access without granting then user-confluence access	Resolved	Fixed
	CONF-4897	Using JIRA base-URL for JIRAPORTLET and JIRAISSUES links	Resolved	Fixed
	CONF-4878	the children listing on the bottom of the page should sort the child pages alphabetically	Resolved	Fixed
	CONF-4971	export space returns zip file with Space Details only	Resolved	Fixed
	CONF-5059	Link to 'discard' draft does not discard draft unless clicked twice.	Resolved	Fixed
	CONF-5046	Line separating drafts under 'User Profile' breaking up when draft is without a title	Closed	Fixed
	CONF-5028	Linked user name Wiki tag crashes Confluence	Closed	Fixed
	CONF-5024	Notification emails are adding an additional confluence directory (contextPath) in certain URL's	Closed	Fixed
	CONF-5087	Confusing alert message when you have a draft on page create	Resolved	Fixed
	CONF-5082	Changes to the length of "Recently Updated" list on dashboard.action page is lost immediately after visiting another page	Closed	Fixed
	CONF-4976	Drafts JavaScript error when editing page on Safari	Resolved	Fixed
	CONF-5038	wysiwyg-javascript exception NS_ERROR_NOT_AVAILABLE	Resolved	Fixed
	CONF-4684	I don't want email attachments (e.g. digital sigs) to show up in the 'recently updated' list	Resolved	Fixed
	CONF-5089	System info page is showing $action.getDatabaseDriverName()	Resolved	Fixed
	CONF-4998	News items don't appear in favorite pages section	Resolved	Fixed
	CONF-4996	NullPointerException when searching for users via the Manage User Administration screen	Resolved	Fixed
	CONF-5026	Crash generating PDF from the online documentation	Resolved	Fixed

	CONF-4829	Pop-up page for inserting links doesn't size properly	Closed	Fixed
	CONF-5126	Warning icon in the concurrent editing message does not respect the context path	Resolved	Fixed
	CONF-5129	IllegalArgumentException when exporting a space as the anonymous user	Resolved	Fixed
	CONF-5119	Email format in Profile and notifications pages should be consistent	Resolved	Fixed
	CONF-5118	Drafts tab is not visible in the user profile when external user management is enabled	Resolved	Fixed
	CONF-5154	User picker in Edit Space Permissions page is only accessible by super users	Closed	Fixed
	CONF-5141	Upgrade stalls on AdditionalIndexes1UpgradeTask	Closed	Fixed
	CONF-5056	Search by "Last Week" and "Last Month" fail in the new year.	Resolved	Fixed
	CONF-5052	Explicit background colours break PDF export	Resolved	Fixed
	CONF-5103	When moving a page, you can't change space and parent in same step	Resolved	Fixed
	CONF-6194	Editing space layouts edits the main layout instead	Closed	Duplicate
	CONF-5008	NPE getChangesSinceLastEdit Error getting changes since last edit:	Resolved	Fixed
	CONF-5100	Anonymous comment warning image does not respect the context path	Resolved	Fixed
	CONF-5012	Dashboard favourite space selection dies in Opera 8.5 (and perhaps other versions too)	Resolved	Fixed
	CONF-5011	Javascript error on AJAX heartbeat	Resolved	Duplicate
	CONF-5007	Edit Page returns System Error	Resolved	Fixed
	CONF-5021	NullPointerException when disabling a servlet plugin	Resolved	Fixed

Release Notes 2.1.3

Confluence 2.1.3 is a maintenance release that resolves some issues users may have encountered using Confluence 2.1.2 or earlier, including issues related to Rich Text editing, LDAP user management, restoring to MySQL, and Unicode on MS SQL Server

2.1.3 is a free upgrade for all customers who purchased their Confluence license after January 23rd, 2005.

Who should upgrade?

Confluence 2.1.3 fixes a number of bugs found in Confluence 2.1.2 and earlier. Customers are recommended to upgrade if they are experiencing any of the issues fixed in this release.

Confluence is a **recommended upgrade** for all customers using Confluence's new atlassian-user LDAP user management with dynamic group mappings. For more information see USER-95

Upgrade Procedure

Upgrading Confluence should be pretty easy. If you are upgrading from Confluence 2.1 to 2.1.2, you can find instructions here. *We strongly recommend that you backup your `confluence.home` directory and database before upgrading!*

If you are upgrading from Confluence 2.0.3 or earlier, please check the release-notes of the other major Confluence releases:

- Release Notes 2.1
- Release Notes 2.0
- Release Notes 1.4

- Release Notes 1.3
- Release Notes 1.2
- Release Notes 1.1

Changes in 2.1.3

See also: Issues resolved for 2.1.3

Improvements

- You can now suppress camel-case linking with {nl:CamelCaseWord} – CONF-3700
- Confluence RSS macro is now compatible with Google News RSS – CONF-4892
- The {contentbylabel} macro supports a wider range of options – CONF-5137

Bugs Fixed

- Many fixes related to Rich Text editing, and switching between Rich Text and Wiki Markup editing.
- Further PDF export fixes – CONF-4930
- The upload option in the attachments macro works – CONF-5127
- Fixed a bug where Word exports would open in the wrong application – CONF-5163
- Handling of situation where site exceeds its licensed user limit and then removes users is improved – CONF-5208
- Restoring a site to MySQL now properly removes and re-creates database tables – CONF-5153
- Confluence now uses correct NVARCHAR type for multibyte text in MS SQL Server – CONF-5204

Issues resolved for 2.1.3

Errors were reported by the JIRA trusted connection.

- APP_UNKNOWN; Unknown Application: {0}; ["confluence:4557196"]

JIRA Issues (38 issues)				
Type	**Key**	**Summary**	**Status**	**Resolution**
	CONF-5173	Add License ID to View License Page	Resolved	Fixed
	CONF-5204	Merge SQLServerIntlDialect from confluence_2_0_stable branch to HEAD and confluence_2_1_stable	Resolved	Fixed
	CONF-3783	Strike through does not work for links	Resolved	Fixed
	CONF-5380	Allow Editing mode to be a preference	Resolved	Answered
	CONF-3700	Suppression of CamelCase link generation for certain words	Resolved	Fixed
	CONF-5212	Display database connection URL/datasource information on the system info page	Resolved	Fixed
	CONF-5208	You have exceeded the maximum number of users for your license error	Closed	Fixed
	CONF-5190	Catch all errors that may be thrown from a macro	Resolved	Fixed
	CONF-4892	Google News RSS requires a user agent for access	Resolved	Fixed
	CONF-5137	Add more options to contentbylabel macro	Closed	Fixed
	CONF-5211	Handle datasource driver exceptions on the system info page gracefully	Resolved	Fixed
	CONF-5169	SearchEntitiesAction displays $user.name, $user.fullName, $user.email instead of the proper values when searching for a group	Resolved	Fixed
	CONF-5167	Links with an apostrophe fails to render	Resolved	Fixed
	CONF-4260	JTidy removes empty divs, spans	Resolved	Fixed
	CONF-4960	Mixing ordered and unordered lists confuses WYSIWYG at times	Resolved	Fixed

	CONF-5077	JTidy error message when saving edited page	Resolved	Fixed
	CONF-5153	Error restoring on MySQL	Resolved	Fixed
	CONF-4251	User management has design flaw !!	Resolved	Fixed
	CONF-5342	User Picker throws xwork ConfigurationException	Resolved	Fixed
	CONF-5165	Ampersands and backslashes multiply...	Resolved	Fixed
	CONF-5163	Export a page as a word document opens exportword.action in IE instead of pagename.doc	Closed	Fixed
	CONF-5175	jira external users managment: Null pointer exception after try to browse all users or all group in confluence administrative task	Closed	Fixed
	CONF-5182	Swicthing between Rich Text and Wiki Markup loses content between {html} macro tags	Resolved	Fixed
	CONF-4874	Image popup shows ${lang_conf_ok} and ${lang_conf_cancel} after you have attached an image.	Resolved	Cannot Reproduce
	CONF-4930	Error trying to export confluence 2.0 documentation to PDF	Resolved	Duplicate
	CONF-4865	Rich Text editor can mangle horizontal rule tag ('----')	Resolved	Fixed
	CONF-4896	Table containing heading is broken by WYSIWYG editor.	Resolved	Fixed
	CONF-4915	WYSIWYG editor not accepting certain edits on Save	Resolved	Fixed
	CONF-4847	Using header notation and the children macro on the same line brings up a JavaScript alert	Resolved	Fixed
	CONF-4624	Losing html macro contents in Rich Text editor	Closed	Duplicate
	CONF-4681	Daily notification email displays wrong date.	Resolved	Fixed
	CONF-4493	{code} element produces syntax errors when trying to preview	Closed	Cannot Reproduce
	CONF-5127	Option 'upload' does not work for the attachment macro	Closed	Fixed
	CONF-5201	HTML produced when newline typed in table cell doesn't produce correct markup	Resolved	Fixed
	CONF-5049	Bad behaviour of the Wysiwyg editor with paragraphs	Resolved	Fixed
	CONF-5042	Link contents not changed when a link is edited directly	Resolved	Fixed
	CONF-3351	Sign-up takes 20 minutes to load	Resolved	Fixed
	CONF-5233	Cross-site scripting vulnerability in the full name user profile field	Closed	Fixed

Release Notes 2.1.4

Confluence 2.1.4 is a maintenance release that resolves some issues users may have encountered using Confluence 2.1.3 or earlier. Confluence 2.1.4 fixes more than 55 issues, including bugs related to space import/export, the WYSIWYG editor, page breadcrumbs, and many more

2.1.4 is a free upgrade for all customers who purchased their Confluence license after February 16th, 2005.

Who should upgrade?

Confluence 2.1.4 fixes a number of bugs found in Confluence 2.1.3 and earlier. Customers are recommended to upgrade if they are experiencing any of the issues fixed in this release.

Confluence is a **recommended upgrade** for all customers using Confluence's new atlassian-user LDAP user management, as it fixes important issues related to user login validation.

Upgrade Procedure

Upgrading Confluence should be pretty easy. If you are upgrading from an earlier Confluence 2.1 release, you can find instructions here. *We strongly recommend that you backup your `confluence.home` directory and database before upgrading!*

> If you are running a **Sybase** or **Microsoft SQL Server** database, please replace confluence/WEB-INF/classes/com/atlassian/confluence/upgrade/UpgradeUtils.class with this version before upgrading.
>
> You should also install this if your upgrade fails with the message: **Session is currently disconnected**

If you are upgrading from Confluence 2.0.3 or earlier, please check the release-notes of the other major Confluence releases:

- Release Notes 2.1
- Release Notes 2.0
- Release Notes 1.4
- Release Notes 1.3
- Release Notes 1.2
- Release Notes 1.1

Changes in 2.1.4

More than 55 issues were resolved between 2.1.3 and 2.1.4. For the complete list, see: Issues resolved for 2.1.4

Note about markup for embedded content

A bug was fixed in 2.1.4 that previously allowed spaces as attribute separators in embedded content wiki markup (CONF-5406). This means `!image.jpg|alt=Great pic!` now works correctly, but `!image.jpg|border=1 alt=test!` doesn't. The correct and documented way to write the latter continues to be `!image.jpg|border=1,alt=test!`.

Improvements

- {livesearch} macro results can now be limited to a particular space – CONF-3432
- A setting was added to General Configuration to control the maximum number of attachments that can be uploaded at once – CONF-5447
- A setting was added to General Configuration to enable or disable GZip content encoding – CONF-5257
- "Insert Link" and "Insert Image" now work correctly on Safari in OS X 10.4.4 and later – CONF-5435
- A draft is automatically saved when moving from 'Edit' to 'Preview' – CONF-5366
- A `getPermissionsForUser` method was added to the remote API to allow administrators to retrieve the permissions of particular users – CONF-5439

Notable Bugs Fixed

- Many fixes related to the maintenance of the ancestors table while moving or renaming pages, which may have resulted in the misplacing of page breadcrumbs or inherited page-level permissions – CONF-5104, CONF-5243, CONF-5244
- Many fixes related to Rich Text editing, including the handling of tables and whitespace – CONF-5299, CONF-5231, CONF-5492, CONF-5294, CONF-5324, CONF-5247, CONF-5362, CONF-5265
- Space import no longer overwrites part of global configuration – CONF-5370
- Confluence no longer tries to create indexes with names bigger than DB2 can handle – CONF-5316
- "Insert Link" and "Insert Image" no longer erase contents of edit fied in Safari on OS X 10.4.4 and later – CONF-5329
- Links from the {labels} macro are now space-relative – CONF-5346
- emoticons no longer mistaken for embedded image markup – CONF-3369
- "Profile:" no longer mistaken as the start of a file: URL

Issues resolved for 2.1.4

Errors were reported by the JIRA trusted connection.

- APP_UNKNOWN; Unknown Application: {0}; ["confluence:4557196"]

JIRA Issues (68 issues)				
Type	**Key**	**Summary**	**Status**	**Resolution**
	CONF-5263	Modify encoding test to UPPER and LOWER results in the database	Resolved	Fixed

	CONF-5403	Slow edit page load time	Closed	Fixed
	CONF-5401	Slow dashboard due to getPermittedEntities()	Closed	Fixed
	CONF-4546	Modifying livesearch plug-in to search by space, pages, documents.	Closed	Fixed
	CONF-5257	Add gzip response encoding to general configuration	Resolved	Fixed
	CONF-5290	Display enabled plugins in the error page	Resolved	Fixed
	CONF-5366	Ensure a draft is saved when flipping between Edit and Preview	Resolved	Fixed
	CONF-4353	Open source the standard Confluence plugins and macros	Resolved	Fixed
	CONF-5382	Remove hand-coded 'style' attributes from generated HTML	Resolved	Fixed
	CONF-5435	Support 'insert link' and 'insert image' on Safari	Closed	Fixed
	CONF-5439	New getPermissions method for remote API to return permissions for a given user to a space	Closed	Fixed
	CONF-3282	Conflue needs hasPermission RPC	Closed	Fixed
	CONF-3432	Allow option in livesearch to limit search to a specific space	Resolved	Fixed
	CONF-5447	Create UI for maxAttachmentsInUI setting	Closed	Fixed
	CONF-5303	Using JIRA base-URL for JIRAISSUES links	Resolved	Fixed
	CONF-5770	Allow enabling/disabling the WYSIWYG-Editor per User	Resolved	Fixed
	CONF-5244	Page breadcrumbs get out of sync when moving pages between spaces	Resolved	Fixed
	CONF-5266	Type the word *Profile:* it gets messed up	Resolved	Fixed
	CONF-5265	Weird modifications done by the Rich Text editor	Resolved	Fixed
	CONF-5262	NullPointerException when include macro references an invalid space	Resolved	Fixed
	CONF-5241	Getting ClassCast Exception when using seraph-paths.xml to secure additonal directories in confluence webapp	Closed	Fixed
	CONF-5326	breadcrumb for when you view the "drafts" tab from your profile is wrong.	Resolved	Fixed
	CONF-5304	Space Index always shows even restricted pages	Resolved	Fixed
	CONF-5349	In the plain website theme, users with edit permission should see the standard Confluence interface.	Closed	Fixed
	CONF-5247	WYSIWYG editor replaces "src" with "xsrc" in code snippets	Resolved	Fixed
	CONF-4687	System error clicking on view change	Closed	Fixed
	CONF-5200	Anchor Link deleted when using WYSIWYG editor.	Resolved	Fixed
	CONF-5406	Embedded content properties split on spaces as well as commas		Fixed

			Resolved	
	CONF-5399	Alt tags on images are broken	Resolved	Fixed
	CONF-5329	"Insert Link" feature clears whole edit area in Safari	Resolved	Fixed
	CONF-5434	Remote calls to getPermissions() fail if user is not superuser	Closed	Fixed
	CONF-5387	When a Site Homepage is set to a Space in General Config, user profiles set to Dashboard don't go to the dashboard but to the Site Homepage	Closed	Fixed
	CONF-5492	table mangled by rich text editor	Resolved	Fixed
	CONF-5503	exporting a space from a test server and importing to a live system replaced the site welcome message on the live system	Resolved	Duplicate
	CONF-5370	Space Export/Import transfer global bandana and conf/*	Resolved	Fixed
	CONF-5448	page version numbers off by 1 in "page edited" HTML email	Resolved	Duplicate
	CONF-5521	Failed to upgrade from 2.1.3 to 2.1.4	Closed	Duplicate
	CONF-5316	On upgrade from 2.0 to 2.1.3 failed to create two indexes since names longer than 18 chars.	Resolved	Fixed
	CONF-5464	Group Picker only shows first 49 groups with no option for paging	Resolved	Fixed
	CONF-5243	Ancestors table gets out of sync when moving pages to another space	Resolved	Fixed
	CONF-5340	Change URL has wrong version number in update email	Resolved	Fixed
	CONF-5252	Group search shows 10 groups, but no page controls	Resolved	Fixed
	CONF-5659	attachment link is incorrect...	Resolved	Fixed
	CONF-5498	Breadcrumbs showing wrong path after page with children moved	Closed	Fixed
	CONF-5681	Links to attachments break when page re-edited	Resolved	Duplicate
	CONF-5313	Watch Mail View Changes link incorrect	Resolved	Fixed
	CONF-5355	Cannot store unicode characters in Site Welcome Message	Closed	Fixed
	CONF-5107	NullPointerException on edit page	Closed	Fixed
	CONF-5346	Labels macro links aren't space relative	Resolved	Fixed
	CONF-5331	rss: Error parsing RSS feed after editing General Configuration	Resolved	Fixed
	CONF-5293	Embedded images do not contain full URL in Word exports	Resolved	Fixed
	CONF-5327	{spaces} tag renders an error in Preview during editation	Resolved	Fixed
	CONF-5324	loosing anchor links	Resolved	Fixed
	CONF-5321	WSYIWYG editor is introducing spaces		Fixed

			Resolved	
	CONF-5352	Class Cast Exception on Blog versioning	Closed	Fixed
	CONF-5235	Blog/News headings style setting forces color to black	Resolved	Fixed
	CONF-5249	Bug in HandleProfileAttachmentsAction.isPermitted() ?	Closed	Fixed
	CONF-5487	NullPointerException listing users of LDAP group	Closed	Fixed
	CONF-5410	Text-only notification email of new blog post does contain unresolved velocity variable	Resolved	Fixed
	CONF-5362	Anonymous user Edit & immediate Cancel produces NullPointerException	Resolved	Fixed
	CONF-5299	WYSIWYG space-adding errors	Resolved	Fixed
	CONF-3369	(!) emoticon sometimes interpretated as image link	Resolved	Fixed
	CONF-5104	Page level permissions set on a page are not updated on page move	Closed	Fixed
	CONF-5294	Existing anchor links are converted to page links by Rich Text editor	Resolved	Fixed
	CONF-5210	Adding lines in bullet lists breaks things	Resolved	Fixed
	CONF-5250	rich text editor corrupts tables with bullets and newlines	Resolved	Fixed
	CONF-5474	Unable to edit page with html markup	Resolved	Fixed
	CONF-6162	Crashed editing a page. (initially nothing special happening from my point of view. repeatable every time I edit this page now)	Resolved	Fixed

Release Notes 2.1.5

Confluence 2.1.5 is a maintenance release that resolves some issues users may have encountered using Confluence 2.1.4 or earlier. Confluence 2.1.5 fixes around 50 issues, including a number of fixes for LDAP support, the WYSIWYG editor, and other areas. 2.1.5 is a free upgrade for all customers who purchased their Confluence license after March 16th, 2005.

Confluence 2.1.5 has a problem which disables the rich text editor link dialog. Please replace the file WEB-INF/classes/com/atlassian/confluence/user/actions/PagePickerAction.properties with this file: PagePickerAction.properties

This is fixed in **2.1.5a**.

Known Issues in 2.1.5a

Anonymous users may get a Javascript error under some circumstances. See CONF-5765 for details and a patch.

When using LDAP user management, changing passwords for local users won't work. See CONF-5775 for details and a patch.

Who should upgrade?

Confluence 2.1.5 fixes a number of bugs found in Confluence 2.1.4 and earlier. Customers are recommended to upgrade if they are experiencing any of the issues fixed in this release.

Confluence is a **recommended upgrade** for all customers using Confluence's new atlassian-user LDAP user management, as it fixes important issues related to LDAP integration.

Upgrade Procedure

Upgrading Confluence should be pretty easy. If you are upgrading from an earlier Confluence 2.1 release, you can find instructions here. *We strongly recommend that you backup your `confluence.home` directory and database before upgrading!*

If you are upgrading from Confluence 2.0.3 or earlier, please check the release-notes of the other major Confluence releases:

- Release Notes 2.1
- Release Notes 2.0
- Release Notes 1.4
- Release Notes 1.3
- Release Notes 1.2
- Release Notes 1.1

Changes in 2.1.5

More than 48 issues were resolved between 2.1.4 and 2.1.5. For the complete list, see: Issues resolved for 2.1.5

Improvements

- LDAP queries are cached to improve performance - CONF-5288
- Administrators can insert custom HTML on every Confluence page, see Administration, Custom HTML - CONF-5350
- Tables created in the rich text editor don't have non-breaking spaces added to each cell - CONF-5044
- News improvements - Info tab available including list of version, News can be exported as PDF - CONF-1989, CONF-5582, CONF-5594

Notable Bugs Fixed

- Anchor, user profile and attachment links now round-trip properly in the rich text editor - CONF-5678
- Paging of LDAP users works properly - CONF-5438
- Servlet plugins can be disabled - CONF-5598
- Forgotten username emails work again - CONF-5530
- User profiles fixes with LDAP - CONF-5225, CONF-5549

Issues resolved for 2.1.5

Errors were reported by the JIRA trusted connection.

- APP_UNKNOWN; Unknown Application: {0}; ["confluence:4557196"]

JIRA Issues (57 issues)				
Type	**Key**	**Summary**	**Status**	**Resolution**
	CONF-5549	The facility to operate on the LDAP profile (based on AtlassianUser) not only in the ReadOnly mode	Closed	Fixed
	CONF-5350	Ability to add stuff to every page on Confluence (e.g., omniture tracking)	Resolved	Fixed
	CONF-1989	No blog entry versioning?	Resolved	Fixed
	CONF-5594	Export of news as PDF	Closed	Fixed
	CONF-5044	Do we need in table cells?	Resolved	Fixed
	CONF-5621	Include latest release of chart macro	Closed	Fixed
	CONF-5582	Display "Info" tab for News/Blog posts	Resolved	Fixed
	CONF-5513	Unserializable objects in the session: bucket.search.lucene.SearchWordsLister	Closed	Fixed
	CONF-4791	Error generating PDF when the title contains a '&' .	Closed	Fixed
	CONF-5533	Access denied on URL returned by remote API exportSpace()	Resolved	Fixed
	CONF-5488	Site Welcome Message (unicode) changed to question marks after restarting server	Closed	Fixed
	CONF-5511	Upgrade from 2.0.3 to 2.1.4 fails on ReduceIndexNameLengthUpgradeTask for index sp_permusername_idx	Resolved	Fixed
	CONF-5517	Confluence 2.1.4 fails to compile due to missing maven dependencies		Fixed

			Resolved	
	CONF-5530	Forgotten username doesn't work	Resolved	Fixed
	CONF-5497	NPE ErrorQueuedTaskQueue	Closed	Fixed
	CONF-5584	Page restrictions UI doesn't update correctly	Resolved	Fixed
	CONF-5591	Breadcrumb expansion does not work on left-nav theme	Resolved	Fixed
	CONF-5601	Merge performance is incredibly bad on long pages	Resolved	Fixed
	CONF-5598	Servlet plugins do not disable or uninstall	Resolved	Fixed
	CONF-5225	User profiles not editable with LDAP user management	Closed	Fixed
	CONF-5609	ClassCastException on Membership Check	Closed	Fixed
	CONF-5605	Users with no groups still receive daily update emails	Closed	Fixed
	CONF-5651	Gallery macro thumbnails height and width attributes are sometimes -1	Resolved	Fixed
	CONF-5561	DWREngine failure in IE and	Closed	Fixed
	CONF-5654	Page titles with a period '.' breaks	Resolved	Fixed
	CONF-5547	The error by adding new group when using LDAP (based on AtlassianUser)	Closed	Fixed
	CONF-5664	Can't create multiple space labels on space admin	Resolved	Fixed
	CONF-5662	make ALT+s submit the comment form	Resolved	Fixed
	CONF-5658	Typist's error: missing single quote	Closed	Fixed
	CONF-5630	Advanced page is not displaying Space Labels	Resolved	Fixed
	CONF-5673	Demonstration space has broken link on Thumbnail page	Resolved	Fixed
	CONF-5687	Jira Issues macro seems to add 4 issues when using the 'count' argument	Resolved	Fixed
	CONF-5525	Error formatting macro: navmap: java.lang.ClassCastException	Closed	Fixed
	CONF-5693	Unknown Group error on setting page level permission	Closed	Fixed
	CONF-5695	NullpointerException thrown from MergedPager	Closed	Fixed
	CONF-5438	Pagination of users is buggy beyond the 10th paged result	Resolved	Fixed
	CONF-5699	Missing image button in rich text editor when adding a new comment	Closed	Fixed
	CONF-5562	Name of file downloads doesn't work for non-ASCII characters	Resolved	Fixed
	CONF-5675	Fetching Mail leads into exception		Fixed

			Resolved	
	CONF-5578	Mail Import FAIL	Closed	Fixed
	CONF-5105	When using LDAP, In "Manage Users" fullname is incorrectly displayed sometimes..	Closed	Fixed
	CONF-6151	Link Properties in Rich Text Mode Not Working	Closed	Handled by Support
	CONF-5437	Text +like this+ should use the sematic <ins> rather than <u>.	Closed	Won't Fix
	CONF-6571	CLONE -Page titles with a period '.' breaks	Resolved	Cannot Reproduce
	CONF-8654	CLONE -Page titles with a period '.' breaks daily backup	Resolved	Handled by Support
	CONF-3645	Malformed links in Daily Change Email	Resolved	Fixed
	CONF-4614	When downloading Multi-byte named attachments, Its character is broken.	Resolved	Fixed
	CONF-5709	ClassCastException when trying to fetch members	Closed	Fixed
	CONF-5777	WYSIWYG Linking Pop-up Does Not Work - specifying link and clicking OK results in nothing	Closed	Fixed
	CONF-5730	Create external link in WYSIWYG does not work	Resolved	Fixed
	CONF-5746	System error when trying to set permissions for a space	Resolved	Fixed
	CONF-5288	LDAP user and group queries need to cache	Resolved	Fixed
	CONF-5500	Cancelling of the Copy page function returns users to a blank page	Resolved	Fixed
	CONF-5678	corrupted anchor link on round trip in create or copy page	Resolved	Fixed
	CONF-5679	Local user accounts cannot login if LDAP or Active Directory server is down	Closed	Fixed
	CONF-5545	Edit Profile tab missing for ldap users	Closed	Fixed
	CONF-5537	Breadcrumbs in wrong order	Closed	Fixed

Release Notes 2.2

Atlassian is proud to announce the release of Confluence 2.2, otherwise known as Shoalhaven Existing customers who wish to upgrade, or new users who wish to try out Confluence for 30 days, can download Confluence from the Atlassian website: http://www.atlassian.com/software/confluence

Shoalhaven is the seventh major update to Confluence (in two years!). It introduces 'personal spaces', support for localisation/internationalisation, CAPTCHA spam protection, a multitude of new extension points for plugin developers, a simpler LDAP configuration syntax, and more.

Upgrading from 2.1

Upgrading Confluence should be pretty easy: you can find instructions here. *We strongly recommend that you backup your `confluence.home` directory and database before upgrading!*

MySQL Driver Support
For users using MySQL with Confluence 2.2 (or higher), please ensure that you are using the latest (3.1.12) MySQL Java Connector. Earlier versions of the MySQL connector have a bug which may prevent Confluence from upgrading successfully. (more information below)

LDAP Authentication
If you are currently using LDAP authentication **through OSUser** (that is, you adjusted your `osuser.xml` file), please *do not upgrade to 2.2*. This feature is currently unavailable and an issue has been filed for it here. This issue has been resolved in Confluence 2.2.1.

User Management Bug
You must apply the patch attached on this bug report to be able to create users in this version. This issue has been resolved in Confluence 2.2.1.

Index Rebuilding
Upgrading to 2.2 will trigger a rebuild of Confluence's search index. This will cause certain features such as search, the dashboard and RSS feeds to behave unpredictably for up to a few minutes after upgrading.

LDAP Integration Configuration Changes
The configuration file syntax for atlassian-user has changed. Customers who configured Confluence 2.1 to use AtlassianUser LDAP integration must follow the instructions on this page.

Anonymous Permission Changes
The way that anonymous permissions are applied to users who have logged in has changed. Administrators should understand how they have changed below before upgrading.

Plugin Loading Changes
Prior to Confluence 2.2, it was possible for an outdated plugin, or a plugin with unsatisfied dependencies, to break the entire plugin subsystem. Confluence 2.2 is much more careful about loading plugins, and isolating them if they break. This may, however, cause Confluence to refuse to load an entire plugin, simply because one of its contained plugin modules will not load.

Upgrading from 2.0 and earlier

Users upgrading directly from 2.0 or earlier should also read the 2.1 Release Notes for caveats regarding the 2.0 -> 2.1 upgrade.

Contents

1. New Features
2. Improvements
3. Notable Bug-fixes
4. Important note for MySQL users

See also: Issues Resolved for 2.2

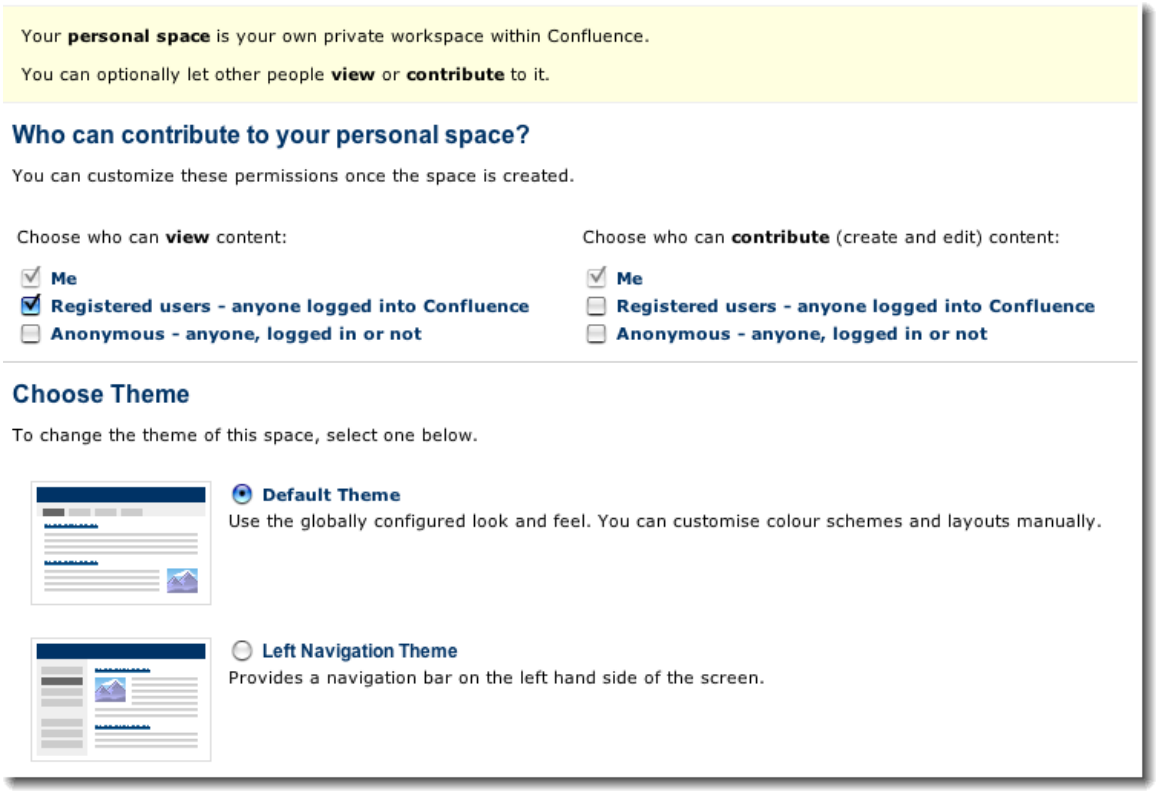
New Features

Personal Spaces

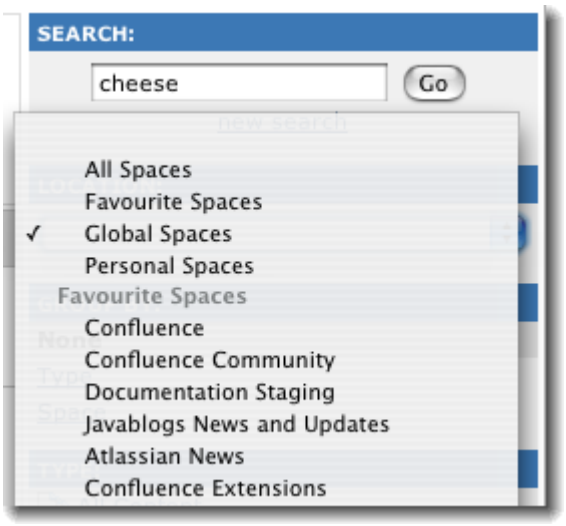
Two of the most frequent questions we get from Confluence customers have been: "How do I give my users their own wiki?" and "How do I give my users their own blog?" It seems everybody needs a little Personal Space.



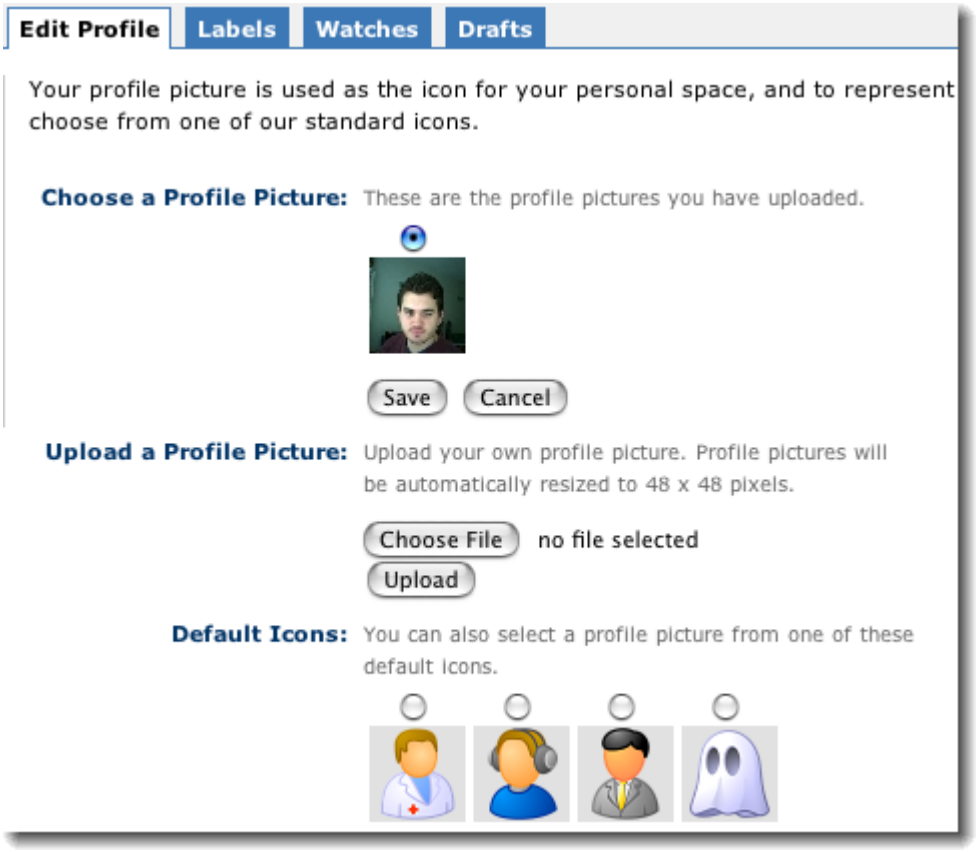
Personal spaces belong to particular users, and rather than being listed on the dashboard, are available from the user's profile. (Future versions of Confluence will feature a 'people browser' to make it easier to discover the interesting personal spaces on your server). They can contain pages and news items like any other space, be searched and browsed. They can be kept private, or opened up so the whole world can view and edit them, just like global spaces.



Confluence's search and RSS builder interfaces have been updated to make it easy to choose whether you're interested in personal spaces or not.



Oh, and if you create a personal space, don't forget to upload (or choose) your own profile picture, so people can see who you are.

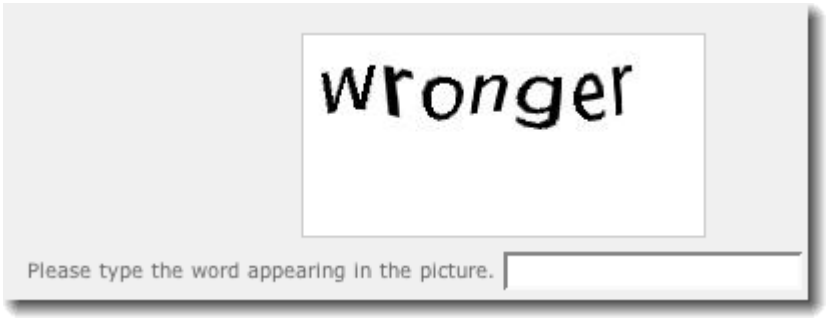


Localisation/Internationalisation

Confluence now supports drop-in language packs to change the language of the user interface. The global administrator can select a default language for the entire site, while individual users can set their preferred language in their preferences.

No language packs are currently available, but we are currently working with our global partners so we can begin to provide translations. If you are interested in translating Confluence into another language, you can find instructions on building a language pack here: Language Module.

CAPTCHA Support



Many of the more public Confluence wikis have been suffering at the hands of spammers. CAPTCHA support adds the familiar 'type in this word' question to signup, edit and comment forms, to defeat automated spamming bots. You can turn on CAPTCHA from the global administrative console, and also choose which users will, or will not be subject to the spam check.



Plugin Improvements

Each version of Confluence is more customiseable than the last. One of the most exciting things to come from our recent plugin competition was learning just how creative our plugin developers are, and discovering more ways we can help them add features to Confluence.

- Job Module and Trigger Module allow you to schedule periodic tasks to run within Confluence
- Web UI Module allow plugin developers to add new links, buttons and tabs to the Confluence interface, and make it easier for theme developers to stay up to date with changes to the Confluence UI.
- Code Formatting Module allow you to plug support for new languages into the {code} macro (or override the existing language support with something better).

We've also made a number of improvements to the plugin system, including:

- Making sure that a single 'rogue' plugin can't bring down the whole plugin subsystem.
- Enabling 'conditional get' for plugin resources, so browsers don't download them again and again.
- Fixing problems that may occur when a servlet or component plugin is reloaded or upgraded. (With many thanks to Dan Hardiker of Adaptavist.)

New atlassian-user Configuration Syntax

As promised, we have cleaned up the configuration file syntax for our atlassian-user user management library. This should make it much easier to configure Confluence to use external user repositories such as LDAP. You can find details of the new configuration file format here: Add LDAP Integration

Improvements

Permissions Changes

In Confluence 2.1 and earlier, permissions that were assigned to Anonymous users were not automatically assigned to logged-in users, leading to the confusing situation where you could view a page when not logged in, but not view it when you *were* logged in.

In Confluence 2.2 and later, permissions that are assigned to the Anonymous user are also assigned to all logged in users. The sole exception to this rule is the global Use Confluence permission, which must still be explicitly granted to any user who wishes to log in. (This exception is necessary due to licensing restrictions).

Search

With help from Kelvin Tan, we've overhauled much of the underlying infrastructure of Confluence's search engine. Search in Confluence should now be more efficient, and some searches that were problematic before (such as wildcard* searches) now work as you would expect

them to.

LDAP Performance

Confluence's performance against external LDAP user repositories should now be significantly faster.

Other

- The General Configuration screen is much better organised
- You can choose a space's initial permissions and theme when you create it
- The left-navigation theme has improved
- Themes may now be packaged with icons that will be displayed in the theme selector (see the screenshot of the "create personal space" screen above for an example)
- Confluence Standalone is now packaged with Tomcat 5.5
- Atom feeds are now generated to the Atom 1.0 standard
- The back-end storage of attachment files has been changed to resolve a number of problems that were caused when the underlying filesystem did not support the character-set of the attachment filename

Important note for MySQL users

For users using MySQL with Confluence 2.2 (or higher), please ensure that you are using the latest (3.1.12) MySQL Java Connector. Earlier versions of the MySQL connector have a bug which is triggered by improvements in Confluence 2.2. These earlier connector versions will result in an error being recorded in your logs on upgrade (and will result in unstable operation of Confluence)

```
1.ERROR [hibernate.tool.hbm2ddl.SchemaUpdate] execute could not complete schema update
```

You can download the latest MySQL connector from the MySQL Java Connector 3.1 download page. Please be sure that you remove any older versions of the connector from your application server.

The Confluence 2.2 Team

Development and Support
Tom Davies
Jeremy Higgs
David Loeng
Charles Miller
Daniel Ostermeier
Christopher Owen
Matt Ryall
Jens Schumacher

Oversight & ~~Mis~~ management
Mike Cannon-Brookes
Scott Farquhar

Issues Resolved for 2.2

Errors were reported by the JIRA trusted connection.

- APP_UNKNOWN; Unknown Application: {0}; ["confluence:4557196"]

JIRA Issues (84 issues)				
Type	**Key**	**Summary**	**Status**	**Resolution**
	CONF-1470	Add Tomcat HTTPS connector commented out in Standalone	Resolved	Fixed
	CONF-5843	Disable Glue servlet on JRockit	Resolved	Fixed
	CONF-5209	Review the 2.2 Release Impact document	Resolved	Fixed
	CONF-5351	Upgrade Confluence Standalone to use Tomcat 5.5	Resolved	Fixed
	CONF-3725	Plugin Package and Module Configurations	Closed	Fixed
	CONF-193	Confluence services	Closed	Fixed
	CONF-3724	Pluggable Menus	Resolved	Fixed

	CONF-1315	Internationalization	Resolved	Fixed
	CONF-577	Tie blogs in with user profiles	Closed	Fixed
	CONF-1080	Personal Space	Closed	Fixed
	CONF-5477	Add support for localisation of help tips sidebar	Closed	Fixed
	CONF-6002	Throw an event on user searches	Closed	Fixed
	CONF-4273	Plugins hooking into the admin UI	Resolved	Fixed
	CONF-5155	move bucket cache manager into its own package	Resolved	Won't Fix
	CONF-4861	Review the writing/export of files to disk using content IDs, instead of filenames	Closed	Fixed
	CONF-5172	Change object id generation to be cluster friendly.	Resolved	Fixed
	CONF-4399	Make Plugins configurable from the admin interface	Closed	Duplicate
	CONF-5613	Add new plugin events	Closed	Fixed
	CONF-5834	Clean up General Configuration	Resolved	Fixed
	CONF-5846	Add feed icon to pages with discoverable RSS	Closed	Fixed
	CONF-5847	Add authentication to discoverable RSS feeds when user is logged in	Resolved	Fixed
	CONF-5854	Move Atom feeds to Atom 1.0	Resolved	Fixed
	CONF-5780	Make /admin/ redirect to /admin/console.action	Closed	Fixed
	CONF-5739	Create personal spaces via remote API	Resolved	Fixed
	CONF-5741	Convert space to personal space via Remote API	Resolved	Fixed
	CONF-5738	Add ability to upload profile photos via the remote API	Resolved	Fixed
	CONF-5022	Remove EHCache specific implementations from Confluence.	Resolved	Fixed
	CONF-5903	Logged in users should inherit access rights of anonymous users.	Closed	Duplicate
	CONF-5917	Add pluggable formatters to the code macro	Closed	Fixed
	CONF-5905	Confluence 2.2 - feedback	Resolved	Fixed
	CONF-5977	Parent Page field in page Edit screen is clearly too small for comfortable editing	Resolved	Fixed
	CONF-4860	Use an attachment's content ID instead of filename when writing to disk	Resolved	Fixed
	CONF-3208	Allow attachments to be stored in the database	Resolved	Fixed

	CONF-5057	Allow for manually flushing the Confluence caches.	Resolved	Fixed
	CONF-5470	RSS feeds don't honour page permissions?	Closed	Cannot Reproduce
	CONF-2471	Unix backup fails to restore on windows when attachment filenames contain accented characters	Resolved	Fixed
	CONF-5528	User profile pictures are stored twice in the one upload	Resolved	Fixed
	CONF-5529	User profile attachments are not deleted when the user is removed	Resolved	Fixed
	CONF-4953	Giving Anonymous Access does not give logged in user permission	Resolved	Fixed
	CONF-5422	Anonymous access permissions are not inherited - logged in users don't even share anonymous access privelages	Closed	Duplicate
	CONF-5598	Servlet plugins do not disable or uninstall	Resolved	Fixed
	CONF-3326	Uploading a modified plugin with the same name does not update velocity template	Closed	Fixed
	CONF-4903	Plugin keys must be lowercase	Closed	Fixed
	CONF-4478	plugins cannot use capitalization in their names	Closed	Fixed
	CONF-5666	On edit space labels page, clicking on team labels doesn't work	Resolved	Fixed
	CONF-5724	Create page doesn't use space logo	Resolved	Fixed
	CONF-5868	XML import adds leading newline to CDATA fields	Resolved	Fixed
	CONF-2948	Attachments unreliable due to file-names	Resolved	Fixed
	CONF-5580	Searching by username and fullname fails	Resolved	Fixed
	CONF-5653	Space import doesn't set mail 'From' address	Resolved	Fixed
	CONF-5918	Standalone package has apps in the webapp directory	Resolved	Fixed
	CONF-5943	Livesearch results print out {searchresultwithexcerpt}	Resolved	Fixed
	CONF-5526	Favourites (spaces and pages) are lost on site restore	Closed	Fixed
	CONF-5989	Layout all screwed up on Preview	Resolved	Fixed
	CONF-6000	NPE whne requesting new password	Closed	Fixed
	CONF-5413	Boilerplate text in a new Space still refers to 'You can edit the content of this page using the Edit link on the right.'	Closed	Won't Fix
	CONF-6046	Switching between rich text editor and wiki editor breaks markup	Closed	Obsolete
	CONF-5951	Attachment file names can be set to an empty string	Resolved	Fixed
	CONF-5909	japanese character search input does not send correct query	Resolved	Fixed

	CONF-5754	Accessing an unauthorized download directly returns a "401 Unauthorized" page, rather than "Page Not Found"	Resolved	Fixed
	CONF-5987	Entering a task with a % in {dynamictasklist} will make the page unusable	Resolved	Obsolete
	CONF-4879	Navigation map macro throws NullPointerException when there are no labels	Closed	Duplicate
	CONF-4777	Login does not redirect to last page location	Resolved	Fixed
	CONF-3168	Lucene reindexing fails and goes to 100% CPU on multi-processor systems	Resolved	Answered
	CONF-5255	HibernateSystemException: "query did not return a unique result" when performing label/space operations	Closed	Fixed
	CONF-5912	Xwork actions fail due to being unable to load the class from the plugin	Resolved	Fixed
	CONF-4503	Problems with special characters in file names	Closed	Fixed
	CONF-2885	Use error page if Confluence is running but missing JDBC drivers on classpath	Closed	Fixed
	CONF-5710	When inserting image onto a page via WYSIWIG Editor -> "Insert/Edit Image" icon, the "ok" and "cancel" buttons are no longer visible when there are too many image attachments	Resolved	Fixed
	CONF-5995	Attachments show up as anonymous on dashboard	Resolved	Fixed
	CONF-5056	Search by "Last Week" and "Last Month" fail in the new year.	Resolved	Fixed
	CONF-5205	Admin license details link hangs	Resolved	Fixed
	CONF-5109	UserSessionExpiryListener incompatible with certain application servers	Closed	Fixed
	CONF-5116	Non-breaking space html entities showing up in search result page	Resolved	Fixed
	CONF-5361	Search results do not take into account exclusions (NOT) nor wild card characters (?)	Closed	Fixed
	CONF-5585	Moving a page should also move children pages	Resolved	Fixed
	CONF-7025	Add Page with no title gives no error and does not save	Resolved	Fixed
	CONF-4974	Concurrent modification exception in SimpleDisplayServlet	Resolved	Fixed
	CONF-3183	Broken Plug-ins are fatal	Resolved	Fixed
	CONF-5761	Manage groups reports only 100 members per group for LDAP groups	Resolved	Fixed
	CONF-5915	The Space dropdown in the Move Page box on Edit view is not alphabetical	Resolved	Fixed
	CONF-5214	Searching for multiple user using the user-filter is currently not possible	Resolved	Fixed
	CONF-5418	Plugin resource servlet does not provide headers necessary for caching	Resolved	Fixed
	CONF-5236	Manage Groups produces a StackOverFlowError with LDAPDynamicGroups	Closed	Fixed

Release Notes 2.2.1

Confluence 2.2.1 is a maintenance release that resolves some issues users may have encountered using Confluence 2.2, including issues related to LDAP user management, i18n, personal spaces and drafts.

2.2.1 is a free upgrade for all customers who purchased their Confluence license after May 18th, 2004.

Who should upgrade?

Confluence is a recommended upgrade for all customers using Confluence with external LDAP user management. All other customers should consult the list of issues resolved, and determine if this release is necessary for their own deployments.

Upgrade Procedure

Upgrading Confluence should be pretty easy. If you are upgrading from Confluence 2.2, you can find instructions here. *We strongly recommend that you backup your `confluence.home` directory and database before upgrading!*

Draft Table Upgrade
Upgrading Confluence to 2.2.1 will cause all autosaved 'drafts' to be deleted from the database. Be sure everyone has saved their work before upgrading!

Confluence 2.2.1 and 2.2.1a
The current version available for download is 2.2.1a. 2.2.1 was originally released with an old version of the atlassian-plugins library, which caused the plugin fixes listed below not to be included in the release. Confluence 2.2.1a resolves that problem.

PageChildrenMacro IndexOutOfBoundsException error
Customers who downloaded 2.2.1 or 2.2.1a before the official announcement may encounter an error when using the {children} macro to view the children of other pages in the same space. There is a patch for this issue attached to CONF-6197

If you are upgrading from Confluence 2.1.5 or earlier, please check the release-notes of the other major Confluence releases:

- Release Notes 2.2
- Release Notes 2.1
- Release Notes 2.0
- Release Notes 1.4
- Release Notes 1.3
- Release Notes 1.2
- Release Notes 1.1

Changes in 2.2.1

For a complete list, see: Issues resolved for 2.2.1

External User Management

- Creating a new user will no longer cause a primary key violation on some Confluence installations. CONF-6092
- Fixed a regression where authentication against an external LDAP server via the old OSUser integration layer did not function. CONF-6052
- Fixed a possible ClassCastException on viewing a user's profile. CONF-6021

Plugins

Note: Due to a scheduling error, these fixes were not included in Confluence 2.2.1, but are in the supplementary 2.2.1a release.

- Component plugins are now correctly unregistered when they are removed. CONF-4041
- Plugins may now contain and reference dependent jar files. Thanks to Dan Hardiker of Adaptavist for the patch. PLUG-8
- Installing an older version of an already-installed plugin no longer causes an error. PLUG-12
- Removing an uploaded plugin no longer causes it to be deactivated the next time it is installed. PLUG-13

General Stability

- Confluence can now store drafts for spaces with long (>20 character) keys. CONF-6010
- Attachments stored in MySQL databases will no longer be silently truncated to 64KB. CONF-6120
- Global language setting is now saved correctly. CONF-6027
- Missing attachment data no longer causes PDF export to crash. CONF-6063
- Misleading "attachment missing" errors are no longer logged after a space import. CONF-6026

Improvements

- A Javascript alert warns users if they have left the CAPTCHA field blank. CONF-5984
- A link to the logged-in user's personal space is included on every page (from their full name). CONF-6056
- The {children} macro can now show the children of pages in another space. CONF-5986

Issues resolved for 2.2.1

Errors were reported by the JIRA trusted connection.

- APP_UNKNOWN; Unknown Application: {0}; ["confluence:4557196"]

JIRA Issues (46 issues)

Type	Key	Summary	Status	Resolution
	CONF-5986	Show children of a page from another space	Closed	Fixed
	CONF-6056	Link personal spaces from a user's name on the top right hand side	Resolved	Fixed
	CONF-6096	Don't placeFocus() on edit pages	Resolved	Fixed
	CONF-5984	Javascript warning if captcha is not filled in	Closed	Fixed
	CONF-6018	NullPointerException in recently updated macro for annymous use when profiling is enabled.	Closed	Fixed
	CONF-6021	View User Profile throws ClassCastException	Resolved	Fixed
	CONF-6024	"No more results" error when running atlassian-user migration JSP	Closed	Fixed
	CONF-6022	Unique Index Violations	Closed	Duplicate
	CONF-6026	Attachment filename fixer on space import attempts to fix filenames of attachments from other spaces	Closed	Fixed
	CONF-6027	Global language setting not honoured	Closed	Fixed
	CONF-6039	NullpointerException in PermittedPagesScope	Closed	Fixed
	CONF-6040	ClassCastException when viewing page information	Closed	Fixed
	CONF-6013	Unable to specify the publishDate of a blog entry using the SOAP API	Resolved	Fixed
	CONF-6025	NullPointerException in HibernateProfileProvider.getPropertySet	Closed	Fixed
	CONF-6063	PDF export throws error when attachment data is missing	Resolved	Fixed
	CONF-6073	Attachments can be renamed to an already existing attachment name	Resolved	Fixed
	CONF-5407	Entering a value containing a '$' dollar sign for a page template variable throws IllegalArgumentException	Resolved	Fixed
	CONF-6081	'Add comment' link is startlingly close to page body text at times - may confuse some users	Resolved	Fixed
	CONF-6087	Templates throw error when variables contain certain strings	Resolved	Fixed
	CONF-6048	Repeating italic and bold markers (_, *) in the same line makes the view gets confused (display some text as bold, some as italic, and some marks)	Closed	Fixed
	CONF-6107	Attachments are not included in backup when stored in database	Resolved	Fixed
	CONF-6065	Duplicated versions in page version history	Resolved	Won't Fix
	CONF-6010	Drafts spacekey column length restricted to 20 causing BatchUpdateException	Resolved	Fixed

	CONF-6050	System error clicking next on manage users	Resolved	Fixed
	CONF-6110	CLONE -LDAP Authentication via OSUser is broken	Closed	Fixed
	CONF-6092	Cannot create new user due to primary key violation or unique constraint error	Closed	Fixed
	CONF-6176	Servlet Plugins don't unload	Resolved	Fixed
	CONF-4676	Can't disable "Script executing and show flash macro" plugin.	Resolved	Fixed
	CONF-6197	{children} macro fails with IndexOutOfBoundsException when viewing children of another page in the same space.	Resolved	Fixed
	CONF-6125	Spaces in Login Name cause the User preferences Link to break.	Closed	Fixed
	CONF-5283	dynamictasklist tasks cannot be in Japanese (and presumably can't contain any multi-byte characters)	Closed	Fixed
	CONF-6083	Active directory users cannot login to Confluence after a certain time	Resolved	Fixed
	CONF-6045	Can not create new users after importing Confluence 2.1.5 full-export into 2.2	Resolved	Duplicate
	CONF-6268	Edit page with code macro displays error instead of code text	Resolved	Fixed
	CONF-6293	'Undefined Pages' shows link tips	Resolved	Fixed
	CONF-6023	Error starting up Confluence 2.2	Closed	Fixed
	CONF-4041	Uploaded component plugins aren't unregistered when the plugin is uninstalled	Resolved	Fixed
	CONF-6052	LDAP Authentication via OSUser is broken	Closed	Fixed
	CONF-6019	Language pack plugin directory not included - no build.xml file	Resolved	Fixed
	CONF-5966	uploadspacelogo.action throws exception	Resolved	Fixed
	CONF-6120	Attachments truncated to 64k in MySQL database storage	Resolved	Fixed
	CONF-6124	Changing the main decorator within a space updates the global main decorator as well	Closed	Fixed
	CONF-6075	Cannot create personal space if username contains international characters	Closed	Fixed
	CONF-6123	Invalid user search term throws NPE	Resolved	Fixed
	CONF-6028	upgrade from 2.1 -> 2.2 fails	Closed	Fixed
	CONF-6061	Spaces or international characters in usernames causes a 404 page when trying to view the preferences page	Closed	Fixed

Release Notes 2.2.2

Confluence 2.2.2 is a maintenance release that resolves some issues users may have encountered using Confluence 2.2, including issues related to PDF generation, internationalisation, backup/restore and more.

2.2.2 is a free upgrade for all customers who purchased their Confluence license after May 31st, 2005.

Who should upgrade?

Confluence 2.2.2 resolves CONF-6237, a bug related to the backup and restore of personal spaces. As such it is a recommended upgrade for customers wishing to restore or import data containing personal spaces. Other users should consult the list of resolved issues, and determine if the upgrade is necessary for their installation.

Upgrade Procedure

Upgrading Confluence should be pretty easy. If you are upgrading from Confluence 2.2 or 2.2.1, you can find instructions here. *We strongly recommend that you backup your `confluence.home` directory and database before upgrading!*

If you are upgrading from Confluence 2.1.5 or earlier, please check the release-notes of the other major Confluence releases:

- Release Notes 2.2
- Release Notes 2.1
- Release Notes 2.0
- Release Notes 1.4
- Release Notes 1.3
- Release Notes 1.2
- Release Notes 1.1

Changes in 2.2.2

Over 25 issues were resolved between Confluence 2.2.1 and 2.2.2. For a complete list, see: Issues resolved for 2.2.2

Backup/Restore

- An issue was resolved that could cause personal spaces to be imported incorrectly from backups. Existing backup files are still valid, the error was only during the re-importing process – CONF-6237

PDF Export

- PDF exports can now be generated for pages containing non-latin characters. To enable these exports you will need to provide Confluence with the correct font – CONF-1457
- PDF exports created in Confluence are now properly searchable, both in Confluence (CONF-1360) and in Adobe Acrobat (CONF-4529)

Other

- (UI) Fixed a Javascript error when saving a page with CAPTCHA disabled – CONF-6221
- (User management) Resolved an error when deleting a non-empty group – CONF-6082
- (i18n) Resolved issues with hard-coded English text in various places – CONF-6202
- (i18n) Resolved issues with internationalisation keys showing up in the permissions guide – CONF-6053

Issues resolved for 2.2.2

Errors were reported by the JIRA trusted connection.

- APP_UNKNOWN; Unknown Application: {0}; ["confluence:4557196"]

JIRA Issues (25 issues)				
Type	**Key**	**Summary**	**Status**	**Resolution**
	CONF-5910	Allow support for Japanese license keys in 2.2.1.	Resolved	Fixed
	CONF-1457	PDF export of unicode pages	Closed	Fixed
	CONF-6119	Default !embed! rendering to <img> if we can't determine the mime type	Closed	Fixed
	CONF-1360	Generated PDF is not searchable	Closed	Fixed
	CONF-6260	Dashboard favourites should link to /label/my:favourite	Resolved	Fixed
	CONF-6064	Hide PDF export link when viewing historical page version	Closed	Fixed
	CONF-6202	More Missing Strings (Notation Guide)	Closed	Fixed
	CONF-6203	Missing Linked Pages	Closed	Fixed
	CONF-6211	Including a slash at the end of the personal space URL does not work.	Resolved	Fixed
	CONF-6219	Spelling error in "page conflict" messages	Closed	Fixed
	CONF-6221	Javascript error "captchaTextField has no properties" when saving edits and CAPTCHA is disabled.	Closed	Fixed

	CONF-5979	CAPTCHA gets confused when you're editing two pages at once.	Closed	Fixed
	CONF-6082	Exception when deleting group with members	Closed	Fixed
	CONF-6116	Single page PDF export fails	Resolved	Not a bug
	CONF-6158	Personal spaces appear on recently updated, even when the personal space is not accessible	Closed	Fixed
	CONF-6251	Error thown user email value is not set	Resolved	Fixed
	CONF-6185	$item in daily summary mail as author of anonymous comment	Resolved	Fixed
	CONF-2547	Paths to internal images are incorrect on export.	Resolved	Fixed
	CONF-6262	License partner not checked on upgrade	Closed	Fixed
	CONF-4529	Generated PDFs are not searchable	Closed	Fixed
	CONF-6053	Permission guide link gives garbage screen	Resolved	Fixed
	CONF-5631	SMTP is broken in the stand-alone config	Closed	Fixed
	CONF-6118	IllegalStateException when creating a space that already exists	Closed	Cannot Reproduce
	CONF-6132	Editing a page with the Rich Text Editor breaks picturized links.	Resolved	Fixed
	CONF-6237	Restore loses defaults to "global" in the field "spacetype" in the "space" table	Closed	Fixed

Release Notes 2.2.3

Confluence 2.2.3 is a maintenance release that resolves some issues users may have encountered using Confluence 2.2, including issues related to LDAP user management, installation, backup/restore and more. This release also includes a patch for a serious security issue regarding global permissions.

2.2.3 is a free upgrade for all customers who purchased their Confluence license after June 8th, 2005.

Who should upgrade?

Confluence 2.2.3 resolves a security bug related to unauthorised modification of global permissions. As such **this release is a recommended upgrade for all customers**.

For users of older versions of Confluence who do not wish to upgrade, a patch for this issue will be available shortly.

Upgrade Procedure

Upgrading Confluence should be pretty easy. If you are upgrading from Confluence 2.2, 2.2.1 or 2.2.2, you can find instructions here. *We strongly recommend that you backup your `confluence.home` directory and database before upgrading!*

If you are upgrading from Confluence 2.1.5 or earlier, please check the release-notes of the other major Confluence releases:

- Release Notes 2.2
- Release Notes 2.1
- Release Notes 2.0
- Release Notes 1.4
- Release Notes 1.3
- Release Notes 1.2
- Release Notes 1.1

Changes in 2.2.3

As part of a new process for pushing out bug fix releases more quickly, this release follows just a week after the release of 2.2.2. Six issues were resolved between Confluence 2.2.2 and 2.2.3. For a complete list, see: Issues resolved for 2.2.3

Global Permissions

- A security problem was identified where unauthorised users could modify global permissions. For more information, including details on how to patch previous Confluence versions, please see the security advisory

Backup/Restore

- An issue was resolved where attachments were not imported from space exports. Existing backup files are still valid, the error was only during the re-importing process – CONF-6286

External user management

- Queries to LDAP repositories are now correctly escaped, so user and group names may contain commas – USER-106

Other

- (UI) Fixed parent link in preview of new page – CONF-6275
- (authentication) Fixed incorrect redirect from login when base URL contained an extra slash – CONF-6261
- (installation) Improved error handling during Confluence setup – CONF-6276

Issues resolved for 2.2.3

Errors were reported by the JIRA trusted connection.

- APP_UNKNOWN; Unknown Application: {0}; ["confluence:4557196"]

JIRA Issues (9 issues)				
Type	**Key**	**Summary**	**Status**	**Resolution**
	CONF-6286	Importing post-2.2 space export does not correctly import attachments	Resolved	Fixed
	CONF-6070	sAMAccountName vs display names in confluence LDAP integration	Closed	Fixed
	CONF-6558	LDAP queries do not escape special characters correctly	Resolved	Fixed
	CONF-6276	Setup ignores exceptions thrown by LDAP group manager getGroup()	Resolved	Fixed
	CONF-6275	Links in preview mode do not work when creating a page	Resolved	Fixed
	CONF-6261	LoginAction does not generate correct original URL from Referer if base URL ends with a slash	Resolved	Fixed
	CONF-6288	Welcome text not dispayed on homepage	Closed	Fixed
	CONF-6350	Welcome message resetting to default	Resolved	Duplicate
	CONF-6331	Security problem in permission editing	Resolved	Fixed

Release Notes 2.2.4

Confluence 2.2.4 is a maintenance release that resolves an assortment of issues users may have encountered using Confluence 2.2.

2.2.4 is a free upgrade for all customers who purchased their Confluence license after June 22nd, 2005.

Who should upgrade?

Confluence 2.2.4 is a minor bugfix release. Customers should consult the list of issues resolved in this release to determine if it is worth their while upgrading.

Customers still running Confluence 2.2.2 or earlier are recommended to upgrade, as a significant security vulnerability was resolved in Confluence 2.2.3

Upgrade Procedure

Upgrading Confluence should be pretty easy. If you are upgrading from Confluence 2.2 - 2.2.3, you can find instructions here. *We strongly recommend that you backup your `confluence.home` directory and database before upgrading!*

If you are upgrading from Confluence 2.1.5 or earlier, please check the release-notes of the other major Confluence releases:

- Release Notes 2.2
- Release Notes 2.1
- Release Notes 2.0
- Release Notes 1.4
- Release Notes 1.3
- Release Notes 1.2
- Release Notes 1.1

Changes in 2.2.4

We are currently trialling a process of more frequent bugfix releases, to make fixes available as soon as possible to those people who want them. 2.2.4 resolves six issues:

See also: Issues Resolved for 2.2.4

User Interface

- The "Space Admin" tab is no longer displayed to non-administrators. (Even though the tab was visible, non-administrators could still not access any of the administrative functions) CONF-6385
- Personal spaces are no longer listed in the space picker in the feed builder CONF-6830
- Attempts to create news items with invalid characters now fail with the correct error message CONF-6358

Performance

- The velocity manager is no longer loaded twice, saving memory. CONF-6355
- The velocity cache now expires templates that have not been recently used, potentially saving memory. CONF-6339

Developer API

- Putting pages into, and removing them from the trash now trigger the correct events within Confluence CONF-6353

Issues Resolved for 2.2.4

Errors were reported by the JIRA trusted connection.

- APP_UNKNOWN; Unknown Application: {0}; ["confluence:4557196"]

JIRA Issues (8 issues)				
Type	**Key**	**Summary**	**Status**	**Resolution**
	CONF-6353	Add trashed and restore events for pages and news	Closed	Fixed
	CONF-6355	Wasting memory - Velocity is loaded twice	Closed	Fixed
	CONF-6380	Personal spaces appear in space picker (feed builder)	Closed	Fixed
	CONF-6385	Space admin tab is visible to non-space admins	Resolved	Fixed
	CONF-6358	Error creating news	Resolved	Fixed
	CONF-6339	Velocity cache never gets cleared	Resolved	Fixed
	CONF-6449	No email address in LDAP causes page editing/movement to fail if user is watching	Resolved	Fixed
	CONF-6879	An error occurs when creating an RSS Feed.	Resolved	Fixed

Release Notes 2.2.5

⚠ Release notes are to follow ...

Issues Resolved

Errors were reported by the JIRA trusted connection.

- APP_UNKNOWN; Unknown Application: {0}; ["confluence:4557196"]

JIRA Issues (18 issues)										
Type	**Key**	**Summary**	**Assignee**	**Reporter**	**Priority**	**Status**	**Resolution**	**Created**	**Updated**	**Due**
	CONF-6851	Removing permissions to upload attachments still leaves ability to upload attachments form the Page Editor Area.	Don Brown [Atlassian]	Brendan Patterson [Atlassian]		Closed	Duplicate	Aug 24, 2006	Sep 07, 2006	
	CONF-6486	Exception when attaching new image in rich text editor insert image popup	Dave Loeng [Atlassian]	Matt Ryall [Atlassian]		Resolved	Fixed	Jun 30, 2006	Sep 07, 2006	

	CONF-6423	Merge algorithm reliably destroys Umlauts	Unassigned	Thomas Wana		Resolved	Fixed	Jun 22, 2006	Jun 27, 2006	
	CONF-6419	Base URL derived from request includes unnecessary port 443	Unassigned	Matt Ryall [Atlassian]		Resolved	Fixed	Jun 22, 2006	Jun 22, 2006	
	CONF-6410	Jiraportlet URL munging doesn't work properly for links with single quotes	Unassigned	Matt Ryall [Atlassian]		Resolved	Fixed	Jun 21, 2006	Jun 21, 2006	
	CONF-6401	excluded mail count no longer works	Unassigned	Tom Davies [Atlassian]		Closed	Fixed	Jun 20, 2006	Jun 21, 2006	
	CONF-6382	Make more indexing language options available in General Configuration	Unassigned	Charles Miller [old account, do not assign issues]		Closed	Fixed	Jun 18, 2006	Oct 21, 2008	
	CONF-6379	Search for Japanese strings should not include partial matches	Unassigned	Neeraj Jhanji		Closed	Fixed	Jun 18, 2006	Aug 22, 2007	
	CONF-6301	Add Attachment link shows without permission	Matt Ryall [Atlassian]	Mike Cannon-Brookes [Atlassian]		Resolved	Fixed	Jun 05, 2006	Jun 22, 2006	
	CONF-6228	Global custom color scheme is not used by spaces	Unassigned	Jens Schumacher [Atlassian]		Resolved	Fixed	May 23, 2006	Jul 20, 2006	
	CONF-6184	Trying to navigate past the 10th page in Manager Users throws an exception	Christopher Owen [Atlassian]	Dave Loeng [Atlassian]		Closed	Fixed	May 17, 2006	Jun 27, 2006	
	CONF-6170	Delete Icon (trash can) functionality not immediately obvious to end users	Unassigned	Noam Field		Resolved	Fixed	May 16, 2006	Sep 06, 2006	
	CONF-6148	System error viewing drafts	Unassigned	Bob Swift		Resolved	Fixed	May 12, 2006	Jul 11, 2006	
	CONF-5848	blog-posts + label that doesn't exist results in all news shown	Agnes Ro [Atlassian]	Paul Pavlidis		Resolved	Fixed	Apr 03, 2006	Jul 06, 2008	
	CONF-5437	Text +like this+ should use the sematic <ins> rather than <u>.	Unassigned	Matt Ryall [Atlassian]		Closed	Won't Fix	Feb 07, 2006	Jun 25, 2006	
	CONF-4434	IMAP/Exchange support in mail archive	Unassigned	Robert Christenson		Closed	Fixed	Oct 27, 2005	Dec 30, 2008	
	CONF-3907	Support accessing the POP account via SSL	Unassigned	Daniel Ostermeier		Closed	Fixed	Sep 01, 2005	Jul 10, 2006	
	CONF-2258	Enable IMAPS and POPS	Unassigned	Mike Cannon-Brookes [Atlassian]		Closed	Fixed	Nov 25, 2004	Nov 16, 2006	

Release Notes 2.2.6a

Confluence 2.2.6a is a maintenance release that resolves an assortment of issues users may have encountered using Confluence 2.2.

2.2.6a is a free upgrade for all customers who purchased their Confluence license after June 22nd, 2005.

Who should upgrade?

Confluence 2.2.6a is a minor bugfix release. Customers should consult the list of issues resolved in this release to determine if it is worth their while upgrading. Customers running **Confluence 2.2.5 should upgrade**, as 2.2.6a fixes a security problem which could allow an unauthorised user to modify space mail box configuration. Only version 2.2.5 is affected by this vulnerability. Customers running 2.2.5 who don't wish to upgrade can follow these instructions to patch the vulnerability.

Customers still running Confluence 2.2.2 or earlier are recommended to upgrade, as a significant security vulnerability was resolved in Confluence 2.2.3

Upgrade Procedure

Upgrading Confluence should be pretty easy. If you are upgrading from Confluence 2.2 - 2.2.3, you can find instructions here. *We strongly recommend that you backup your `confluence.home` directory and database before upgrading!*

If you are upgrading from Confluence 2.1.5 or earlier, please check the release-notes of the other major Confluence releases:

- Release Notes 2.2
- Release Notes 2.1
- Release Notes 2.0
- Release Notes 1.4
- Release Notes 1.3
- Release Notes 1.2
- Release Notes 1.1

Changes in 2.2.6a

We are currently trialling a process of more frequent bugfix releases, to make fixes available as soon as possible to those people who want them. 2.2.6a resolves 20 issues.

A few of the issues are noted below, see Issues Resolved for 2.2.6a for a complete list.

User Interface

- The rich text editor allows underlining to be removed, and lists in tables are better behaved. CONF-6450, CONF-6508
- Users with capital letters in their names can watch pages and spaces. CONF-6489

LDAP

- LDAP users can be removed from local Confluence groups. CONF-6131
- The 'Change Password' link is no longer shown when LDAP is enabled. CONF-6360

Performance

- Indexing performance has been improved. CONF-6465

2.2.5 Security Patch

This patch fixes a minor security vulnerability regarding the administration of space mail accounts. It affects only Confluence 2.2.5.

To install this patch please download the three files attached to this page, stop Confluence and copy them to your `WEB-INF/classes/com/atlassian/confluence/mail/actions` directory, where they will replace the existing files. Then start Confluence.

Name	Size	Creator	Creation Date	Comment
RemoveMailAccountAction.class	1 kB	Tom Davies	Jul 12, 2006 23:35	
EditMailAccountAction.class	6 kB	Tom Davies	Jul 12, 2006 23:35	
AddMailAccountAction.class	5 kB	Tom Davies	Jul 12, 2006 23:35	

Issues Resolved for 2.2.6a

Errors were reported by the JIRA trusted connection.

- APP_UNKNOWN; Unknown Application: {0}; ["confluence:4557196"]

JIRA Issues (25 issues)				
Type	**Key**	**Summary**	**Status**	**Resolution**
	CONF-6493	Invalid keyword or missing delimiter error	Closed	Fixed
	CONF-5548	Concurrent use of inner groups and group based on LDAP (based on AtlassianUser)	Closed	Fixed
	CONF-6562	Increase the number of groups shown on the Manage Groups screen	Resolved	Fixed
	CONF-3034	Confluence renderer should render markup within words (w_o_rds)	Resolved	Fixed
	CONF-6200	Please allow editing of links and email addresses included in the ConfluenceActionSupport.properties file	Resolved	Fixed

	CONF-6519	google maps plugin fails in IE	Closed	Fixed
	CONF-6563	Referrers not showing on info page, even when referrers turned on	Resolved	Fixed
	CONF-6489	User is not marked as watching a page or space if their username contains capital letters	Closed	Fixed
	CONF-6518	Space list macro does not update user interface state when requested tab is not available	Resolved	Fixed
	CONF-6499	Site Homepage Default incorrectly set	Closed	Fixed
	CONF-6292	Cannot login when a user belongs to more than 100 groups	Resolved	Fixed
	CONF-6473	Redirect to default homepage of a space after signup	Closed	Fixed
	CONF-6475	Confluence should upgrade to the latest version of the CombinedCachingServlet once JRA-10504 has been resolved	Resolved	Fixed
	CONF-6476	Decimal Number Format in General Config cannot be updated	Resolved	Fixed
	CONF-6131	Cannot remove LDAP user from local confluence group	Resolved	Fixed
	CONF-6508	WYSIWYG Editor problem when editing bulleted list within a table cell	Resolved	Fixed
	CONF-5333	Export to Word doesn't work with Japanese	Resolved	Fixed
	CONF-6616	attachment Error in Linking Window using Insert/Edit Link button while edit page	Closed	Fixed
	CONF-6709	Character is not shown...	Resolved	Duplicate
	CONF-7093	EmbeddedRenderer incorrectly	Resolved	Fixed
	CONF-6450	Unable to remove underlining via the Rich text editor.	Closed	Fixed
	CONF-4721	Image and link insertion dialogs are fixed size and don't scroll	Resolved	Fixed
	CONF-6497	BeanInfo introspection cache not being cleared on webapp shutdown.	Resolved	Fixed
	CONF-6465	Major reindexing performance regression	Resolved	Fixed
	CONF-6515	plugin manager should ignore/log unknown module types	Closed	Fixed

Release Notes 2.2.7

Confluence 2.2.7 is a maintenance release that resolves an assortment of issues users may have encountered using Confluence 2.2.

2.2.7 is a free upgrade for all customers who purchased their Confluence license after July 28th, 2005.

Who should upgrade?

Confluence 2.2.7 is a minor bugfix release. Customers should consult the list of issues resolved in this release to determine if it is worth their while upgrading.

Customers still running Confluence 2.2.2 or earlier are recommended to upgrade, as a significant security vulnerability was resolved in Confluence 2.2.3

Upgrade Procedure

Upgrading Confluence should be pretty easy. If you are upgrading from Confluence 2.2 - 2.2.3, you can find instructions here. *We strongly recommend that you backup your `confluence.home` directory and database before upgrading!*

If you are upgrading from Confluence 2.1.5 or earlier, please check the release-notes of the other major Confluence releases:

- Release Notes 2.2
- Release Notes 2.1
- Release Notes 2.0
- Release Notes 1.4
- Release Notes 1.3
- Release Notes 1.2
- Release Notes 1.1

Changes in 2.2.7

We are currently trialling a process of more frequent bugfix releases, to make fixes available as soon as possible to those people who want them. 2.2.7 resolves 23 issues.

A few of the issues are noted below, see Issues Resolved for 2.2.7 for a complete list.

Task List Macro

- Tasks with the same name in different lists are now distinguished. CONF-5809

User Management

- The manage groups page can now display an unlimited number of groups. CONF-6458
- Users that belong to more than 100 groups are now always able to login. CONF-6292

PDF Export

- Exporting pages to PDF which include other pages will now have correct links and not lose images. CONF-1155

Internationalisation

- Tree view of page hierarchy will not be truncated when using multi-byte character sets. CONF-5872

Issues Resolved for 2.2.7

Errors were reported by the JIRA trusted connection.

- APP_UNKNOWN; Unknown Application: {0}; ["confluence:4557196"]

JIRA Issues (22 issues)				
Type	**Key**	**Summary**	**Status**	**Resolution**
	CONF-6458	MAX_GROUP in EditUserGroupsAction Class is bad	Resolved	Fixed
	CONF-5809	Tasks with same name in different lists aren't distinguished	Resolved	Fixed
	CONF-6610	The personal / favourite labels paging actions are broken (or missing)	Resolved	Fixed
	CONF-5815	Jiraissues mappings don't support status with umlaut (internationalised text)	Resolved	Fixed
	CONF-6598	Empty profiles in search results repeat the previous search excerpt	Resolved	Fixed
	CONF-6632	Display labels in search results	Closed	Fixed
	CONF-6615	Plus (+) in page name is stripped	Resolved	Fixed
	CONF-7093	EmbeddedRenderer incorrectly	Resolved	Fixed
	CONF-6569	Users with view page permission but not edit page permission can remove view page restriction of a page	Closed	Fixed
	CONF-6601	German Umlauts in page titles	Resolved	Fixed
	CONF-5872	In Tree View of List Pages, page name is not displayed when we use Korean characters.	Resolved	Fixed

	CONF-6607	DB2 error when BLOB data too large (Confluence page)	Closed	Fixed
	CONF-6650	Attachments Tab in Browse Space only displays first 100 attachments to a space	Resolved	Fixed
	CONF-6573	group names should be URL-encoded when used in links...	Closed	Fixed
	CONF-6580	Curly-brace tags with nothing between them get mangled when switching to WYSIWYG editing.	Closed	Fixed
	CONF-6596	Posting comments with new images fails silently without permission to upload files	Resolved	Fixed
	CONF-6365	Postgres error in alphabetical page views on confluence.atlassian.com	Resolved	Fixed
	CONF-6590	Max attachment size affects uploadplugin.action	Closed	Fixed
	CONF-6641	Edit Space Permissions footer layout broken	Resolved	Fixed
	CONF-6644	Page renaming should only trigger minor edits in referring pages	Resolved	Fixed
	CONF-6031	Following the "Include these Matches" link in a Site Search displays a page which includes "Searched for blah in $inSpace"	Resolved	Fixed
	CONF-6599	Blog/news posts should show labels	Resolved	Fixed

Release Notes 2.2.8

Confluence 2.2.8 is a maintenance release that resolves an assortment of issues users may have encountered using Confluence 2.2.

> ⚠ Confluence 2.2.8 has slower indexing than previous versions. There is a fix for that regression here: CONF-6908

2.2.8 is a free upgrade for all customers who purchased their Confluence license after August 8th, 2005.

Who should upgrade?

Confluence 2.2.8 is a minor bugfix release. Customers should consult the list of issues resolved in this release to determine if it is worth their while upgrading.

Customers still running Confluence 2.2.2 or earlier are recommended to upgrade, as a significant security vulnerability was resolved in Confluence 2.2.3

Upgrade Procedure

Upgrading Confluence should be pretty easy. If you are upgrading from Confluence 2.2 - 2.2.3, you can find instructions here. *We strongly recommend that you backup your `confluence.home` directory and database before upgrading!*

If you are upgrading from Confluence 2.1.5 or earlier, please check the release-notes of the other major Confluence releases:

- Release Notes 2.2
- Release Notes 2.1
- Release Notes 2.0
- Release Notes 1.4
- Release Notes 1.3
- Release Notes 1.2
- Release Notes 1.1

Changes in 2.2.8

We are currently trialling a process of more frequent bugfix releases, to make fixes available as soon as possible to those people who want them. 2.2.8 resolves 22 issues.

A few of the issues are noted below, see Issues Resolved for 2.2.8 for a complete list.

- Exporting a space as HTML now creates an index page which includes all the pages in the export. CONF-6670
- Timeouts can be set for HTTP requests made from Confluence to RSS feeds and JIRA instances. CONF-6697
- RSS Autodiscovery works in Safari. CONF-4607

- When a group is deleted, space permissions for the group are now correctly removed. CONF-6733
- Change default background colour from grey to white. CONF-6733
- Non-breaking spaces are handled correctly when indexing. CONF-6685
- Rich text editor respects leading spaces in noformat macro. CONF-6527

Issues Resolved for 2.2.8

Errors were reported by the JIRA trusted connection.

- APP_UNKNOWN; Unknown Application: {0}; ["confluence:4557196"]

JIRA Issues (25 issues)				
Type	**Key**	**Summary**	**Status**	**Resolution**
	CONF-6697	Allow configuration of timeouts for any http request made by Confluence	Resolved	Fixed
	CONF-6733	Error when editing permissions for a space	Resolved	Fixed
	CONF-6689	Broken link in View License Page	Resolved	Fixed
	CONF-6527	Rich text editor loses leading spaces in first line after 'noformat' tag.	Resolved	Fixed
	CONF-4607	RSS Autodiscovery doesn't show up on Safari RSS	Resolved	Fixed
	CONF-6668	Contact administrators page broken $getText("no.admins") instead of message appears	Resolved	Fixed
	CONF-6685	Index Tokenizer should treat ' ' as space	Resolved	Fixed
	CONF-6686	User not redirected to restricted page they were trying to view	Closed	Fixed
	CONF-6694	RSS feed macro doesn't show date/time on its items while its help menu show them.	Resolved	Fixed
	CONF-6701	Quotes around image parameters produce invalid HTML	Resolved	Fixed
	CONF-6700	Javascript errors on global permissions and space permissions pages	Resolved	Fixed
	CONF-6727	Site has grey background	Closed	Fixed
	CONF-6740	LDAP user with CN equal to user name hides following group members in group display	Closed	Fixed
	CONF-6741	BucketPropertySetItem belonging to deleted page breaks space import	Resolved	Fixed
	CONF-6731	Dashes in links incorrectly rendered	Resolved	Fixed
	CONF-6675	Deadlock during user creation	Resolved	Fixed
	CONF-7093	EmbeddedRenderer incorrectly	Resolved	Fixed
	CONF-6776	Viewing LDAP groups in Manage Groups	Closed	Fixed
	CONF-6699	Improve 'Too many users' message on License page	Closed	Fixed
	CONF-5821	NullPointerException in UserLister macro when LDAP is used	Resolved	Fixed
	CONF-6690	Blog calendar for July 2006 doesn't show enough rows	Resolved	Fixed

	CONF-6122	When turning on atlassian-user logging Passwords get logged in plaintext	Resolved	Fixed
	CONF-6730	Breadcrumb wrapping disrupts left navigation layout in small windows	Resolved	Fixed
	CONF-6742	Encoding Test fails on MS SQL Server	Resolved	Fixed
	CONF-7054	Use of Draft results in Configuration Exception - "There is no Action mapped for namespace /pages and action name edit$parameters.draftType"	Resolved	Fixed

Release Notes 2.2.9

Confluence 2.2.9 is a maintenance release that resolves an assortment of issues users may have encountered using Confluence 2.2. 2.2.9 resolves an indexing performance regression introduced in 2.2.8, and includes fixes for external user management, and MaxDB and Sybase compatibility.

2.2.8 is a free upgrade for all customers who purchased their Confluence license after September 9th, 2005.

Who should upgrade?

Confluence 2.2.9 is a minor bugfix release. Customers should consult the list of issues resolved in this release to determine if it is worth their while upgrading.

If you are running Confluence 2.2.8, but do not wish to upgrade to 2.2.9, we strongly recommend installing the patch attached to CONF-6908 instead to resolve an issue related to re-indexing performance.

Customers still running Confluence 2.2.2 or earlier are recommended to upgrade, as a significant security vulnerability was resolved in Confluence 2.2.3

Upgrade Procedure

Upgrading Confluence should be pretty easy. If you are upgrading from Confluence 2.2 - 2.2.3, you can find instructions here. *We strongly recommend that you backup your `confluence.home` directory and database before upgrading!*

If you are upgrading from Confluence 2.1.5 or earlier, please check the release-notes of the other major Confluence releases:

- Release Notes 2.2
- Release Notes 2.1
- Release Notes 2.0
- Release Notes 1.4
- Release Notes 1.3
- Release Notes 1.2
- Release Notes 1.1

Changes in 2.2.9

2.2.8 resolves approximately 25 issues. See Issues Resolved for 2.2.9 for a complete list. A few of the issues are noted below.

External user management:

- LDAP users are no longer given the option to change their passwords in Confluence – CONF-6144
- External users belonging to more than 100 groups can log in successfully – CONF-6292
- External users can be removed from local groups – CONF-6832

Database Compatibility

- 'Review restricted pages' report now succeeds on MaxDB – CONF-6849
- Fixed Sybase error when removing a space – CONF-6715

Indexing

- A re-indexing performance regression introduced in 2.2.8 has been resolved – CONF-6908
- Advanced search queries for specific index fields now work as expected – CONF-6714
- Errors in text extractors no longer cause index queue flushing to fail – CONF-6857

Issues Resolved for 2.2.9

Errors were reported by the JIRA trusted connection.

- APP_UNKNOWN; Unknown Application: {0}; ["confluence:4557196"]

JIRA Issues (25 issues)				
Type	**Key**	**Summary**	**Status**	**Resolution**

	CONF-6773	Improve validation of character encoding	Resolved	Fixed
	CONF-6820	Reindexing fails if temp directory doesn't exist	Closed	Fixed
	CONF-6908	Rebuilding index gets progressively slower or fails	Resolved	Fixed
	CONF-5852	Can't select a page as a link via the search	Resolved	Fixed
	CONF-5869	Search result paging is broken for 'Add Link' dialog in Rich Text editor	Resolved	Fixed
	CONF-6881	Possible deadlock in DefaultDraftManager	Resolved	Fixed
	CONF-6292	Cannot login when a user belongs to more than 100 groups	Resolved	Fixed
	CONF-6144	Change password screen should be disabled for LDAP users (or read only users)	Resolved	Fixed
	CONF-6889	Changing permissions on the page changes last modifier	Resolved	Fixed
	CONF-6844	Cannot resume editing of news posts	Resolved	Fixed
	CONF-6874	Content author without matching user record throws NullPointerException in DefaultFeedBuilder	Resolved	Fixed
	CONF-6496	Null Pointer Exception when uploading images	Closed	Fixed
	CONF-6921	Velocity macro not expanded	Resolved	Fixed
	CONF-6914	Catch the zip-too-large-for-VM error and explain it	Resolved	Fixed
	CONF-6741	BucketPropertySetItem belonging to deleted page breaks space import	Resolved	Fixed
	CONF-6783	BODYCONTENT clob column creation fails in DB2	Resolved	Fixed
	CONF-6695	RSS feed for non-existing space throws exceptions	Resolved	Fixed
	CONF-6832	Cannot remove LDAP user from local confluence group	Resolved	Fixed
	CONF-6714	Confluence Lucence Search Terms broken	Closed	Fixed
	CONF-6963	Cannot start Confluence after configuring LDAPDynamicGroupAdaptor	Resolved	Fixed
	CONF-7093	EmbeddedRenderer incorrectly	Resolved	Fixed
	CONF-6717	Column widths specified without % render as pixel length in IE but percent in Firefox	Resolved	Fixed
	CONF-6715	Sybase error when removing space	Resolved	Fixed
	CONF-6779	Blog navigation calendar malfunctions on date with more than one post	Resolved	Fixed
	CONF-6825	Error on Admin main when license ended.	Resolved	Fixed

Release Notes 2.2.10

Confluence 2.2.10 is a maintenance release that resolves an assortment of issues users may have encountered using Confluence 2.2. 2.2.10 resolves an assortment of issues that were reported by customers.

2.2.10 is a free upgrade for all customers who purchased their Confluence license after November 30th, 2005.

Who should upgrade?

Confluence 2.2.10 is a minor bugfix release. Customers should consult the list of issues resolved in this release to determine if it is worth their while upgrading.

If you are running Confluence 2.2.8, but do not wish to upgrade to 2.2.10, we strongly recommend installing the patch attached to CONF-6908 to resolve an issue related to re-indexing performance.

Customers still running Confluence 2.2.2 or earlier are recommended to upgrade, as a significant security vulnerability was resolved in Confluence 2.2.3

Upgrade Procedure

Upgrading Confluence should be pretty easy. If you are upgrading from another Confluence 2.2.x version, you can find instructions here. *We strongly recommend that you backup your `confluence.home` directory and database before upgrading!*

If you are upgrading from Confluence 2.1.5 or earlier, please check the release-notes of the other major Confluence releases:

- Release Notes 2.2
- Release Notes 2.1
- Release Notes 2.0
- Release Notes 1.4
- Release Notes 1.3
- Release Notes 1.2
- Release Notes 1.1

Changes in 2.2.10

2.2.10 resolves over 40 issues. All these issues are listed below:

Errors were reported by the JIRA trusted connection.

- APP_UNKNOWN; Unknown Application: {0}; ["confluence:4557196"]

JIRA Issues (44 issues)

Type	Key	Summary	Status	Resolution
	CONF-5489	Trackbacks are broken	Resolved	Fixed
	CONF-7092	Section/Column is broken for fixed-width columns	Resolved	Fixed
	CONF-6853	Page and Space decorators need a catch all mode for $context and $mode	Resolved	Fixed
	CONF-8916	Create patch to CONF-8877 for Confluence 2.5.4	Resolved	Fixed
	CONF-7387	Possible to create user from XML RPC with null fullname	Closed	Fixed
	CONF-7181	The link table can have rows with spurious space keys inserted	Resolved	Fixed
	CONF-6741	BucketPropertySetItem belonging to deleted page breaks space import	Resolved	Fixed
	CONF-2921	Daily emails list user login names not full names	Resolved	Fixed
	CONF-6995	Rich text editor inserts images with no space between text and '!'	Resolved	Fixed
	CONF-7035	Draft form can be submitted with multiple space keys	Closed	Fixed
	CONF-7038	User with null email address breaks daily report job	Resolved	Fixed
	CONF-7050	labelString attribute doesn't restrict RSS feeds	Resolved	Fixed
	CONF-7061	Some digest notification links don't include the base URL	Resolved	Fixed

	CONF-4801	include cancel button for comments	Resolved	Fixed
	CONF-7134	groupSearchAllDepths and userSearchAllDepths not respected	Resolved	Fixed
	CONF-7132	Clicking 'add' on the permissions administration screen when there is nothing to add should result in a validation error	Resolved	Fixed
	CONF-7282	"Create Personal Space" can create Space with NULL name	Closed	Fixed
	CONF-7305	For comments, first CAPTCHA word fails	Resolved	Fixed
	CONF-7285	Moving attachment to another page clobbers authors of previous versions	Resolved	Fixed
	CONF-7322	Jiraissues macro doesn't show icons or timestamps	Resolved	Fixed
	CONF-6913	Images don't export to PDF if capitalization is wrong	Resolved	Fixed
	CONF-7163	SQL problem with deleting users	Resolved	Fixed
	CONF-7030	contentbylabel macro throws Java errors when used in BLOG or COMMENT objects	Resolved	Fixed
	CONF-7034	problem with embedding mpeg and mp4 within page	Resolved	Fixed
	CONF-7189	Export to PDF & broken {attachment} macro links	Resolved	Fixed
	CONF-6835	wrong mouse tool tip text	Resolved	Fixed
	CONF-6950	Export to PDF using icon=false in templates results in PDF with each word on it's own line	Resolved	Fixed
	CONF-6969	Favorites listing on Personal Labels page -- paging broken	Resolved	Fixed
	CONF-6991	Can't link to image attachments stored on other pages from within comments (but it works in preview of the comment)	Resolved	Fixed
	CONF-7000	Mail sender shows as null in RSS feeds	Resolved	Fixed
	CONF-7058	{dynamictasklist} macro: Problem deleting last item	Resolved	Fixed
	CONF-7101	Draft merge failure logging is too verbose	Resolved	Fixed
	CONF-7016	Restricting a page to an invalid group shows misleading error	Resolved	Fixed
	CONF-7145	Pagination of user labels fails when context path is used	Resolved	Fixed
	CONF-7249	DefaultMailAccountManager.getStore error throws object reference instead of useful toString	Closed	Fixed
	CONF-7178	"Recently Updated" list size on dashboard resets to 10 on each visit to dashboard	Resolved	Fixed
	CONF-7270	Update blog post remotely doesn't save history	Resolved	Fixed
	CONF-7366	Cancelling create page workflow	Resolved	Fixed
	CONF-8272	Confluence - What's This? link near page restrictions note brings up error	Resolved	Fixed

	CONF-7006	After un-installing plugin notification is incorrect	Resolved	Fixed
	CONF-6983	Global Statistics Plugin - incorrect SQL statement	Resolved	Fixed
	CONF-6936	ConfigureSpaceLogoAction affected by attachment size	Resolved	Fixed
	CONF-7085	Editing lists in tables with WYSIWYG adds extra indent asterisk (*) to list	Resolved	Fixed
	CONF-6833	hardcoded http request in {dynamictasklist} macro	Resolved	Fixed

Release Notes 2.3

After much hard toil from the Confluence developers we are happy, nay ecstatic, to announce the availability of Confluence 2.3 (known affectionately as Snowy).Snowy is the eighth major update to Confluence. It supports clustered deployment as 'Confluence Massive', and introduces a people directory, activity statistics plugin, personal timezone preferences, and the ability to access Confluence via the Metaweblog and WebDAV APIs.
Confluence 2.3 is a **free upgrade** for any customer who purchased Confluence after January 4th, 2006.

Upgrading from Confluence 2.2.x

Upgrading Confluence should be fairly straightforward: you can find instructions here. *We strongly recommend that you backup your `confluence.home` directory and database before upgrading!*

Upgrading from Confluence 2.1 and earlier

Users upgrading directly from 2.1 or earlier should also read the 2.2 Release Notes for caveats regarding the 2.1 -> 2.2 upgrade.

Installation Notes

Known Issues and Patches
Please read the Known Issues and Patches section of the release-notes before installing Confluence 2.3.
These issues were resolved with a new maintenance release of Confluence 2.3.1

Contents

1. New Features
2. Improvements

See also: Issues Resolved for 2.3

New Features

Confluence Massive

Confluence is now a clusterable application. This means that it is possible to run multiple Confluence servers behind a load-balancer, to provide high availability, and to scale Confluence beyond the capacity of a single server.

Confluence Massive uses Tangosol Coherence to share data between nodes (and many other things besides).

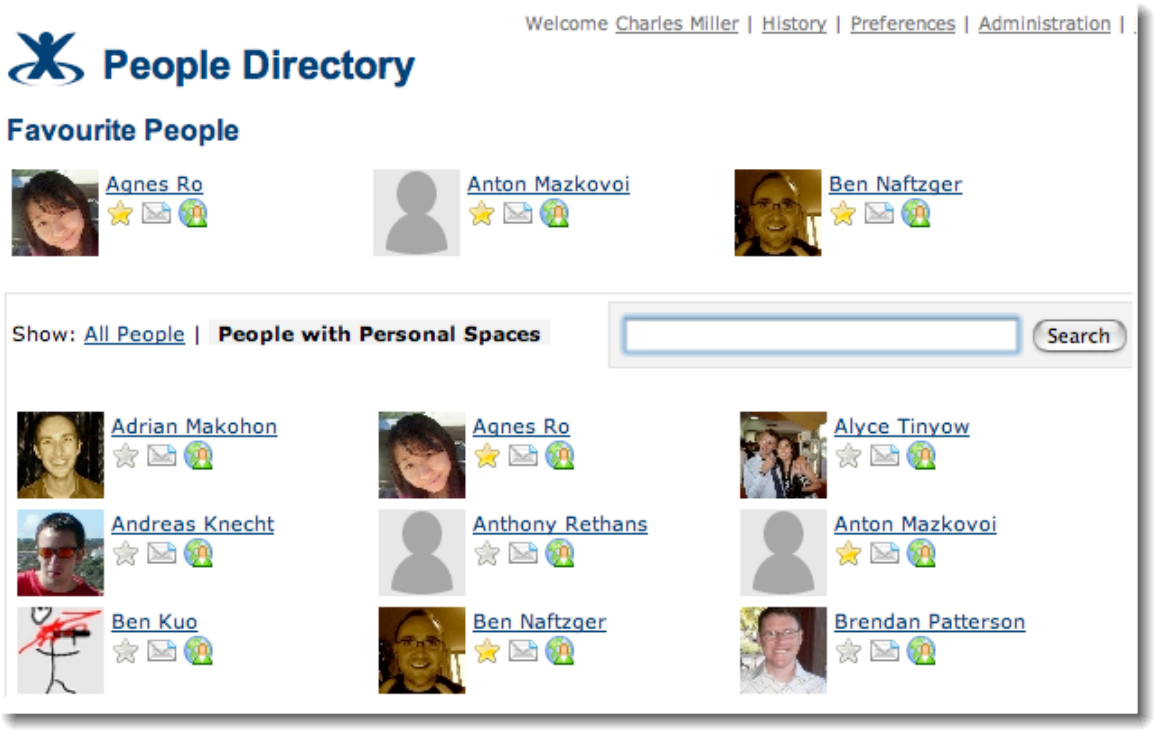
If you are thinking of running Confluence in a cluster, and need to know what is required and how it works, you can read Technical Overview of Clustering in Confluence.

You can find instructions for installing a Confluence cluster here: Confluence Cluster Installation.

Cluster Licensing
Confluence Massive clustering is only enabled if you have a clustered licence. For information on purchasing clustered Confluence licences, please check our pricing page, or contact our friendly sales team.

The People Directory

In version 2.2 (Shoalhaven), we added personal spaces to Confluence. Because of the potential for the proliferation of personal spaces we kept them off the Dashboard and search results by default, but this made it quite hard to find people in a Confluence site.



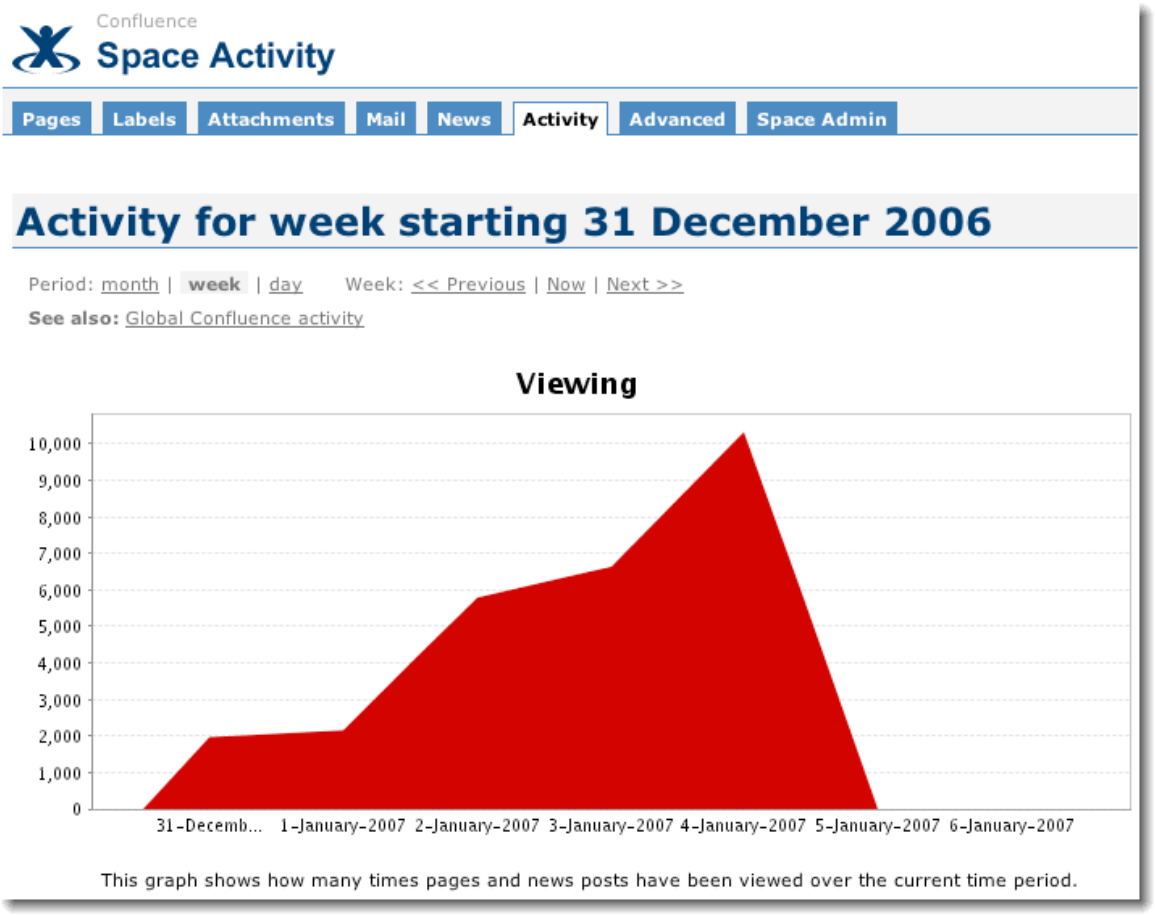
The People Directory, which you can find linked from the bottom of the Dashboard, gives you a way to browse through the other people who use a Confluence instance, their profiles and personal spaces.

If you find people whose personal spaces you want to follow, you can mark them as favourites from the people directory. This will add their personal spaces to your 'My' tab on the Dashboard, and will also make sure they're included when you search or make RSS feeds from your favourite spaces.



Activity Tracking

The activity plugin lets you know what's happening in Confluence: how many pages are being visited or edited in each space or across the whole site, which spaces or pages are the busiest, who are the most prolific editors.



You can view activity statistics for a space in Browse Space, or globally from the administration console.

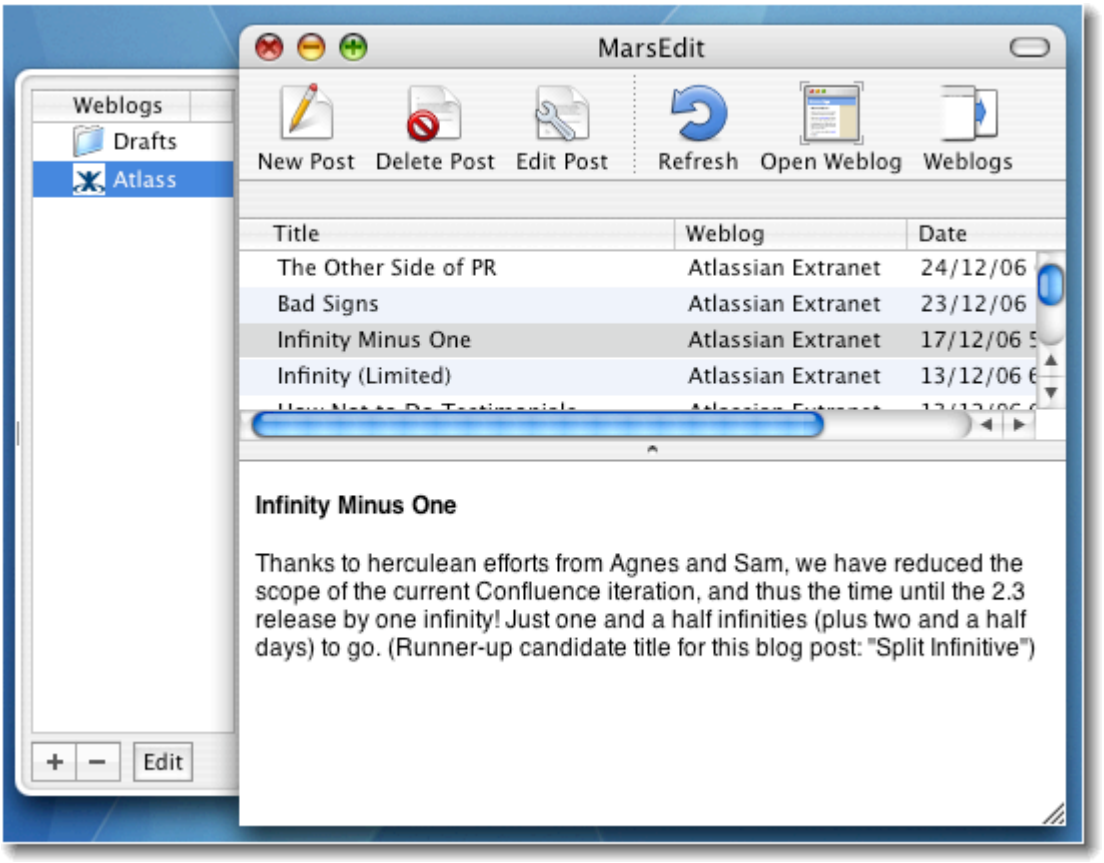
The plugin also provides macros that allow you to embed usage data into a Confluence page: {usage}, {popular} and {topusers}.



Activity tracking does not work in a cluster, and will be disabled for clustered deployments. We're working on making the activity tracker clusterable in a future release. You can follow this issue in JIRA: CONF-7520

Blogger and MetaWeblog API Support

Confluence 2.3 bundles the Blogging RPC Plugin. This allows users to manage their News in Confluence using one of the many available blogger-compatible desktop clients.

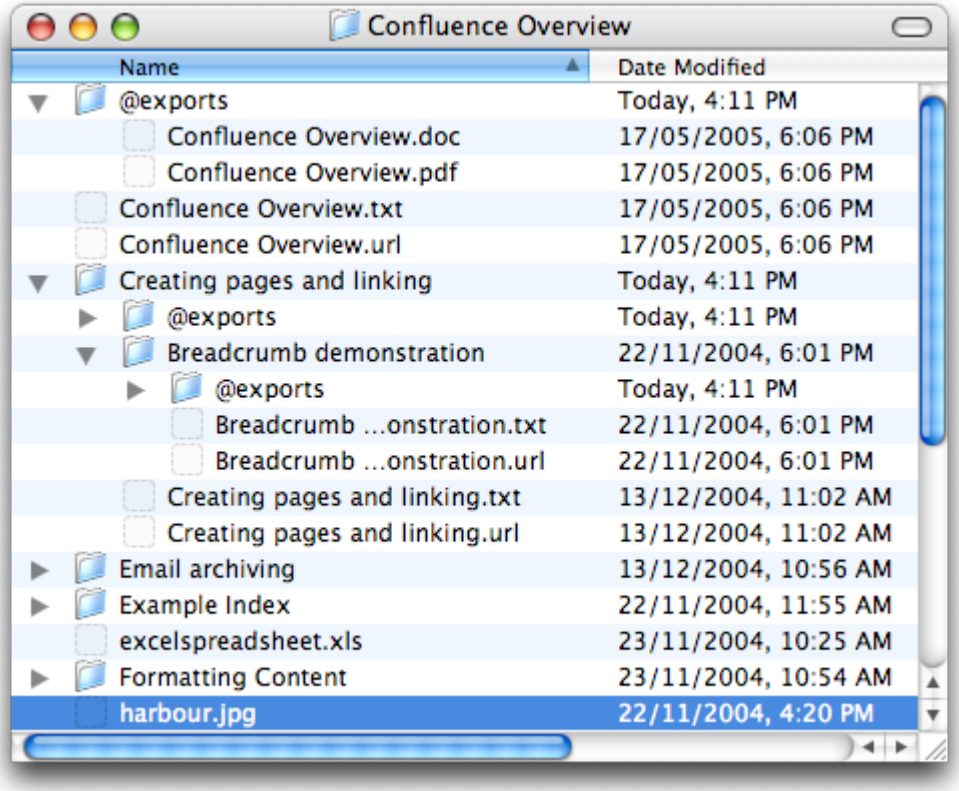


You can read more about this plugin, including instructions for setting up various blogging clients for use with Confluence, on the plugin information page.

WebDAV Client Support

The Confluence WebDAV plugin allows users to mount Confluence as a shared drive, using the WebDAV protocol.

This provides a familiar interface for anyone who uses Confluence as a repository for files: you can browse your wiki straight from Windows Explorer or the Mac Finder; view Word or PDF versions of pages; upload attachments and edit attachments in place; create, edit and move pages.



For more information on the WebDAV plugin, check out: WebDAV Plugin

WebDAV is supported natively in Windows XP (as "Web Folders") and Mac OS X, although there are third-party clients that may provide different functionality or compatibility. We have a compatibility matrix if you're interested.

> **Experimental**
> WebDAV client support is currently experimental, and is disabled by default. You can turn it on from the Plugin management page in the global administration console.

Improvements

Plugins

- Adaptivist's fantastic Plugin Repository Client is bundled with Confluence.
- User Macro Module allow you to share user macros more easily.
- Lifecycle Module allow you to hook into Confluence startup and shutdown.
- Events are now produced when themes are applied, or colour-schemes modified.
- Confluence now ships with 'bundled' plugins that are installed when Confluence is installed, but are not hard-coded into the application (so they can be upgraded without upgrading Confluence).
- Plugin resources now set Last-Modified and ETag headers properly, so they can be conditionally retrieved by browsers.
- Confluence is now built entirely using Maven 2. More on this soon.

Other Changes

- Users can now choose to have dates and times displayed in their own timezone instead of the server's - CONF-1026.
- Most Confluence configuration data has been moved from the filesystem into the database.
- The Clickr Theme is bundled with Confluence.
- HTML exports are more reliable for sites with non-ASCII page titles - CONF-4862.
- You can now get users from multiple LDAP repositories at the same time (or different search paths in the same LDAP repository) - CONF-6034.
- Shortcut links can now have the substitution string in the middle of the URL (CONF-3246), and have custom display text and titles (CONF-514).
- User macros are more flexible: you can configure how macro bodies are pre-rendered (CONF-2293), and you can also specify that a user macro generates wiki markup instead of HTML (CONF-3780).
- Pages with large numbers of attached images no longer eat up too many database connections - CONF-6393.
- When creating an RSS feed with the feed builder, you can choose whether to see full content, diffs, or both - CONF-6321.
- Searches now default to AND rather than OR for combining search terms - CONF-5874.
- The login page no longer prompts you to log in when you're already logged in. - CONF-6843.
- Support for the Nintendo Wii.
- Various improvements to performance that should result in a faster display of pages, and fewer problems for servers with large numbers of space permissions.
- Many static resources in Confluence are cached more effectively, so browsers have to retrieve CSS and Javascript files significantly

less often - CONF-7212.

Known Issues and Patches

- Confluence 2.3 was originally accidentally shipped with three **testing-only language packs**. Since Sunday January 14th, the language packs have been removed from the distribution, but if you downloaded Confluence before that date they will be present. These language packs are machine translations, and are not intended for use on production systems (unless you're a native German speaker and really want a good laugh). If you downloaded an affected copy of Confluence and want to remove these packs, delete `de_DE-1.0.jar`, `ja_JP-1.0.jar` and `ru_RU-1.0.jar` from `confluence/WEB-INF/lib`.
- You need to apply this patch to enable **LDAP user integration** : CONF-7585
- If you have problems **restoring backups** please apply the patch included in this issue: CONF-7584
- Due to class name changes, if you're **upgrading from an existing system** running:
  1. MySQL - you need to edit `confluence.home/confluence.cfg.xml` and replace **`bucket.dialect.MySQLDialect`** with **`com.atlassian.hibernate.dialect.MySQLDialect`**. More information in this issue: CONF-7579
  2. DB2 - you need to edit `confluence.home/confluence.cfg.xml` and replace **`bucket.dialect.DB2Dialect`** with **`com.atlassian.hibernate.dialect.DB2Dialect`**. More information in this issue: CONF-7690
- Log file `atlassian-confluence.log` is not presently generated. Please see these instructions for enabling it.

The Confluence 2.3 Team

Development
Tom Davies
Samuel Le Berrigaud
David Loeng
Charles Miller
Christopher Owen
Agnes Ro
Matt Ryall
Jens Schumacher
Don Willis

Maven Mavens
Michael Mekaail
Tony Truong

Oversight & ~~Mis~~ management
Mike Cannon-Brookes
Scott Farquhar

> *And one was there, a stripling on a small and weedy beast,*
> *He was something like a racehorse undersized,*
> *With a touch of Timor pony — three parts thoroughbred at least —*
> *And such as are by mountain horsemen prized.*
> *He was hard and tough and wiry — just the sort that won't say die —*
> *There was courage in his quick impatient tread;*
> *And he bore the badge of gameness in his bright and fiery eye,*
> *And the proud and lofty carriage of his head.*
>
> - The Man from Snowy River, *A. B. 'Banjo' Paterson*

FileAppender log4j.properties

The 2.3 log4j.properties file doesn't specify a file for the FileAppender (because the necessary lines haven't been uncommented). Hence the log files normally generated via the file appender, `logs/atlassian-confluence.log`, is presently missing?

If you experience this situation and see errors in the logs such as:

```
log4j:WARN File option not set for appender [confluencelog].
log4j:WARN Are you using FileAppender instead of ConsoleAppender?
log4j:ERROR No output stream or file set for the appender named [confluencelog].
```

- Please uncomment the following lines in the **`log4j.properties`** file located under `confluence/WEB-INF/classes` and restart Confluence
  (ie remove the #'s)

```
#log4j.appender.confluencelog.File=${catalina.home}/logs/atlassian-confluence.log
#log4j.appender.confluencelog.MaxFileSize=20480KB
#log4j.appender.confluencelog.MaxBackupIndex=5
```

Relevant resources

Confluence 2.3 Release Notes

Issues Resolved for 2.3

Errors were reported by the JIRA trusted connection.

- APP_UNKNOWN; Unknown Application: {0}; ["confluence:4557196"]

JIRA Issues (170 issues)				
Type	**Key**	**Summary**	**Status**	**Resolution**
	CONF-6106	Document cluster configuration	Resolved	Fixed
	CONF-6406	Ensure upgrade tasks for decorators are working.	Closed	Fixed
	CONF-6595	Update CGLIB to 2.0.2	Resolved	Fixed
	CONF-6035	User and Group caches need to be aware of the repository they are caching for	Closed	Fixed
	CONF-7488	excerpt-include doesn't work with news pages	Resolved	Fixed
	CONF-6104	Cluster admin screen	Resolved	Fixed
	CONF-6102	Make upgrade system work properly in a cluster	Resolved	Fixed
	CONF-6271	Create new plugin bundling system that will allow for user upgrades	Resolved	Fixed
	CONF-1033	Confluence Usage Statistics?	Resolved	Fixed
	CONF-1554	Make Confluence clusterable	Closed	Fixed
	CONF-7196	Add events for look and feel changes	Resolved	Fixed
	CONF-6710	User macro plugin module type	Resolved	Fixed
	CONF-3431	Confluence as a WebDav Server	Resolved	Fixed
	CONF-6321	Allow Feedbuilder to choose to see content, diff, both, or none.	Resolved	Fixed
	CONF-1540	Centralised blogs page	Resolved	Fixed
	CONF-7083	AccessLogFilter should log remote host	Resolved	Fixed
	CONF-8726	User Event Logging and Reporting	Closed	Fixed
	CONF-928	Microsoft Word import	Resolved	Fixed
	CONF-1026	Display times in user's timezone	Resolved	Fixed

	CONF-4281	Conditional-get for plugin resources	Resolved	Fixed
	CONF-3780	User macros should be able to call macro plugins and user macros	Resolved	Fixed
	CONF-847	Error Pages should refer to SITE ADMINSTRATORS	Resolved	Fixed
	CONF-5354	Left hand menu should remember what's open	Resolved	Fixed
	CONF-6564	Referral queue holding on to Hibernate sessions	Resolved	Fixed
	CONF-4862	Use a page's content ID as the filename when exporting as HTML	Resolved	Fixed
	CONF-7702	Leftnav theme should not add background colours to headers h2 to h4	Resolved	Fixed
	CONF-6034	Multiple LDAP repositories	Closed	Fixed
	CONF-6503	Improve and simplify cache statistics page	Closed	Fixed
	CONF-3337	Allow configuration of default search	Resolved	Duplicate
	CONF-3519	Support "Edit attachment via WebDAV"	Closed	Fixed
	CONF-6100	Add caching to database-backed Bandana	Resolved	Fixed
	CONF-514	Shortcut Links should have title & display values	Resolved	Fixed
	CONF-6323	Add warning to Setup Wizard's "Create Empty Database" when there's existing data	Resolved	Fixed
	CONF-6375	Remove usernames from people directory	Resolved	Fixed
	CONF-6393	Attachment loading scalability improvement	Resolved	Fixed
	CONF-6420	Uninstall bad plugins, enforce module load order, enable/disable modules themselves	Closed	Fixed
	CONF-5874	Search should use AND logic by default	Resolved	Fixed
	CONF-6319	Make sure Change Comment is shown in RSS view.	Resolved	Fixed
	CONF-4955	Confluence users should inherit permissions from the anonymous user	Resolved	Fixed
	CONF-3888	Blog post macro: add support for showing blog posts from multiple spaces.	Resolved	Fixed
	CONF-6588	css: auto overflow for .preformatted	Closed	Duplicate
	CONF-7212	Improve caching of static resources	Resolved	Fixed
	CONF-7552	typo in RSS feed screen - 'Attachmends' should be 'Attachments'	Resolved	Fixed
	CONF-6688	Bundle the Confluence Repo Client created by Dan Hardiker to offer a plugin download&install interface	Resolved	Fixed
	CONF-6726	Move default-formatting.properties from filesystem to database	Resolved	Fixed
	CONF-6968	Add tab for personal spaces to dashboard		Fixed

			Resolved	
	CONF-7101	Draft merge failure logging is too verbose	Resolved	Fixed
	CONF-6663	Make the "Visit page outside Confluence" tooltip more user customisable	Resolved	Fixed
	CONF-7251	Plugins need to link stylesheet manually if space-specific colour scheme is to be used	Resolved	Fixed
	CONF-3246	Shortcut links, append-only limitation	Resolved	Fixed
	CONF-6678	Tab from comment writing should go to Post button	Resolved	Fixed
	CONF-7314	Error page (500page.jsp) should suggest Confluence admin to create support case	Resolved	Fixed
	CONF-4638	Uploaded plugin classes are inaccessible from other plugins	Resolved	Fixed
	CONF-7528	Error page should ask users to contact Confluence administrator	Resolved	Fixed
	CONF-6805	Increase Standalone performance with platform-specific Apache Tomcat Portable Runtime Project library	Resolved	Won't Fix
	CONF-7321	Comments should be enabled by default for personal spaces	Resolved	Fixed
	CONF-4845	Bundle different web.xml files with the distribution to support Resin 3.x out of the box	Resolved	Fixed
	CONF-4565	Add overflow control to pre and code blocks via CSS	Closed	Fixed
	CONF-6344	Page restrictions based on group names should be case-insensitive	Resolved	Fixed
	CONF-7097	Rename permission Administrate Confluence to Administer Confluence	Resolved	Fixed
	CONF-4207	Allow plugins to use the full Spring config XML	Resolved	Fixed
	CONF-6103	Move all non-bootstrap configuration out of confluence.cfg.xml	Resolved	Fixed
	CONF-3281	HTML export of a space - links on pages always resolve locally, even if the linked page was not exported	Resolved	Fixed
	CONF-4671	web.xml 2.4/Resin 3.x Schema Violation	Resolved	Fixed
	CONF-4322	Icons missing in HTML-Export of space	Resolved	Fixed
	CONF-4958	HTML export fails to include all attachments	Resolved	Fixed
	CONF-5556	Can not add groups with commas in the name to space permissions	Resolved	Fixed
	CONF-7386	LoginFilter does not redirect to absolute destinations correctly	Resolved	Fixed
	CONF-6517	NullPointerException in AbstractUserProfileAction.getPersonalSpaceKey	Resolved	Fixed
	CONF-6527	Rich text editor loses leading spaces in first line after 'noformat' tag.	Resolved	Fixed
	CONF-7392	Non labelable content inherits labels from previous hit in search results	Resolved	Fixed
	CONF-5781	Certain PNG images in pages cause corrupt PDF exports for pre Java 1.5		Duplicate

			Closed	
	CONF-1881	Default search behaviour should be "AND" for multi-term search	Resolved	Fixed
	CONF-7407	System Error: Confluence cannot be started on Vista with Java 6 RC	Resolved	Fixed
	CONF-7408	Personal spaces listed under "global spaces" in search space drop-down	Closed	Fixed
	CONF-5930	Restore a backup from the filesystem: bucket.core.InfrastructureException: java.lang.OutOfMemoryError	Resolved	Fixed
	CONF-7358	Plugin resource downloads fail if they use a plugin key in the URL	Resolved	Fixed
	CONF-5489	Trackbacks are broken	Resolved	Fixed
	CONF-1155	Export of page which includes other pages loses images	Resolved	Fixed
	CONF-5475	The resources of language packs are not recognised when uploaded through the plugin interface	Resolved	Fixed
	CONF-5284	User Macros not restored after a full restore	Resolved	Fixed
	CONF-6059	Confluence breaks with cglib error on JDK 1.6	Resolved	Fixed
	CONF-7420	Livesearch doesnt work any more and throws exception	Resolved	Duplicate
	CONF-6105	Fix import/export of database-backed Bandana settings	Resolved	Cannot Reproduce
	CONF-5919	Setup should display error when database user does not have permissions to create tables	Resolved	Fixed
	CONF-7424	Typos on excerpt macro in Notation Guide	Closed	Fixed
	CONF-6926	SQL Macro does not work on the extranet	Resolved	Fixed
	CONF-7428	Livesearch macro throws exception when spacekey not specified	Resolved	Fixed
	CONF-6366	People with personal spaces do not appear in people directory search	Closed	Fixed
	CONF-6345	Seemingly random people listed as favourites in the people directory	Closed	Fixed
	CONF-6385	Space admin tab is visible to non-space admins	Resolved	Fixed
	CONF-6408	Make plugin modules state aware	Resolved	Fixed
	CONF-3331	Export process is not able to export thumbnails.	Resolved	Fixed
	CONF-7433	Indenting does not work for bullet lists in rich text editor	Resolved	Duplicate
	CONF-6339	Velocity cache never gets cleared	Resolved	Fixed
	CONF-6431	Tree view in browse space fails with NPE	Resolved	Fixed
	CONF-6418	People Directory number of found users does not match the number of actual users displayed	Closed	Fixed
	CONF-6432	Dynamic tasklist atlassian-plugin.xml lists components in the wrong order		Fixed

			Closed	
	CONF-5955	Old version of cglib 2.0 creates problems with Java 2 security in WebSphere 5.1	Resolved	Fixed
	CONF-1852	We don't index user details	Resolved	Fixed
	CONF-6317	WebUI plugins are unable to properly display the I18N-value of the link name	Resolved	Fixed
	CONF-6876	Temp directory cleanup job should be separate to BackupJob	Resolved	Fixed
	CONF-6657	HTML Export: Duplicate attached images	Resolved	Fixed
	CONF-4773	Long running task view should display a red bar when the task fails, not green bar.	Resolved	Fixed
	CONF-7181	The link table can have rows with spurious space keys inserted	Resolved	Fixed
	CONF-7189	Export to PDF & broken {attachment} macro links	Resolved	Fixed
	CONF-7195	CAPTCHA uses a predictable temp file	Resolved	Fixed
	CONF-7363	Clicking on a news attachment from recently updated list takes you to an incorrect page	Resolved	Fixed
	CONF-4219	Can not render wiki content as inline text using user macros	Resolved	Fixed
	CONF-6701	Quotes around image parameters produce invalid HTML	Resolved	Fixed
	CONF-6728	Confluence app server restart throws away Date/Time format settings	Resolved	Fixed
	CONF-6773	Improve validation of character encoding	Resolved	Fixed
	CONF-6675	Deadlock during user creation	Resolved	Fixed
	CONF-6695	RSS feed for non-existing space throws exceptions	Resolved	Fixed
	CONF-6817	Missing localization strings for Notation Guide	Closed	Fixed
	CONF-6935	Remote API method getPermissions() only returns "modify" for space admins	Resolved	Fixed
	CONF-6973	Thumbnail does not get downloaded on the first time it is viewed	Resolved	Fixed
	CONF-6974	Users who can not view a page due to page level permissions can still see the edit tab	Resolved	Fixed
	CONF-1956	Umlaut in space title breaks PDF export	Resolved	Fixed
	CONF-6995	Rich text editor inserts images with no space between text and '!'	Resolved	Fixed
	CONF-6999	Search not finding specific page	Closed	Duplicate
	CONF-7035	Draft form can be submitted with multiple space keys	Closed	Fixed
	CONF-7038	User with null email address breaks daily report job	Resolved	Fixed
	CONF-7014	XStream introspection cache not cleared when plugin upgraded		Fixed

			Resolved	
	CONF-7050	labelString attribute doesn't restrict RSS feeds	Resolved	Fixed
	CONF-7061	Some digest notification links don't include the base URL	Resolved	Fixed
	CONF-7059	If you install a plugin compiled against the wrong java version, you can't uninstall it.	Resolved	Fixed
	CONF-7093	EmbeddedRenderer incorrectly	Resolved	Fixed
	CONF-7015	Bandana table not found on upgrade to Confluence 2.3-dev	Resolved	Fixed
	CONF-7126	Oops on EAC staff home info page	Resolved	Fixed
	CONF-6084	Captcha is not shown on reply-to comment form	Resolved	Fixed
	CONF-7134	groupSearchAllDepths and userSearchAllDepths not respected	Resolved	Fixed
	CONF-7117	not all LDAP groups shown	Resolved	Fixed
	CONF-7218	"View Conflict" link on drafts page results in ClassCastException	Closed	Cannot Reproduce
	CONF-7467	Confluence slow - looking up space permissions	Resolved	Handled by Support
	CONF-7305	For comments, first CAPTCHA word fails	Resolved	Fixed
	CONF-7315	Confluence HTML renders does not recognise styles applies to links	Closed	Fixed
	CONF-7319	Cannot cancel space removal	Resolved	Fixed
	CONF-7323	Content properties not removed when associated content is deleted	Resolved	Fixed
	CONF-7018	Cannot delete space mysql	Resolved	Fixed
	CONF-7336	Space import fails on content properties without associated content	Resolved	Fixed
	CONF-7359	Clickr Theme Missing some CAPTCHA fields	Resolved	Fixed
	CONF-7322	Jiraissues macro doesn't show icons or timestamps	Resolved	Fixed
	CONF-7337	Clickr plugin: ClassCastException changing tabs from attachments to edit on a blog post.	Resolved	Fixed
	CONF-7286	RSS Feeds listed under the Advanced > RSS Feeds section don't work	Resolved	Fixed
	CONF-7197	No difference between modified and created RSS feed status	Resolved	Fixed
	CONF-7352	Attachment versions not working	Resolved	Fixed
	CONF-4679	Scheduled tasks should _not_ be started before or during an upgrade	Resolved	Fixed
	CONF-5336	HTML Export fails to redirect URL attachments to the "locally" exported directory structure	Resolved	Fixed
	CONF-5642	Incorrect html links in Export Space		Fixed

			Resolved	
	CONF-6665	Sort order broken on space attachments page	Resolved	Fixed
	CONF-2293	Markup disallowed in macros?	Resolved	Fixed
	CONF-2416	Content Sorting in exported space PDF	Resolved	Fixed
	CONF-7005	Flush All Caches link broken	Resolved	Fixed
	CONF-7373	Cannot rename a page with a link to itself	Resolved	Fixed
	CONF-4721	Image and link insertion dialogs are fixed size and don't scroll	Resolved	Fixed
	CONF-5964	Thumbnail etags and last modified data derived from related image	Resolved	Fixed
	CONF-5143	Html Export doens't include images when the image isn't stored within the page itself	Resolved	Fixed
	CONF-5496	Rich text link edit dialog has problems with links where the link text is different to the link markup when there is no alias	Resolved	Fixed
	CONF-6133	WYSIWYG mode is turnig space:page link like test:Döner to emoticon and a corrupted link	Resolved	Fixed
	CONF-6402	Thumbnails are not regenerated when attachments change	Resolved	Fixed
	CONF-7419	Daily notification emails blank from confluence.atlassian.com	Resolved	Fixed
	CONF-7341	BaseAttachmentContentExtractor missing or moved in 2.3-dr2	Resolved	Fixed
	CONF-7060	Plugin Repo doesn't place nicely with clustering	Resolved	Fixed
	CONF-6998	Related matches in other spaces displaying illogical count	Resolved	Fixed
	CONF-6655	Cannot serialise custom objects defined in plugins using Bandana	Resolved	Fixed
	CONF-6843	Login page should tell you clearly if you're already logged in	Resolved	Fixed
	CONF-7381	CLONE -Export of page which includes other pages loses images	Resolved	Fixed
	CONF-7413	Popular labels macro shows empty bulleted list	Resolved	Fixed
	CONF-6891	Formatting Content on Confluence demonstration space grammer mistake	Resolved	Fixed

Release Notes 2.3.1

Confluence 2.3.1 is a maintenance release that resolves an assortment of issues users may have encountered using Confluence 2.3.

2.3.1 is a free upgrade for all customers who purchased their Confluence license after January 23rd, 2006.

Who should upgrade?

Confluence 2.3.1 is a bugfix release which resolves some significant issues in Confluence 2.3. Customers running Confluence 2.3 should review the list of resolved issues below, and upgrade if this release fixes any problems with their current Confluence installation.

Customers running releases older than 2.3 and wishing to upgrade should use this version also.

Significant issues for customers upgrading to Confluence 2.3

- **LDAP integration**: Confluence 2.3 shipped with a broken user migration script for customers wishing to integrate LDAP servers with their Confluence installation. 2.3.1 ships with a fixed version of this script which can also be obtained from CONF-7585
- **Upgrading with MySQL**: Customers using Confluence with MySQL may have had difficulty upgrading their instances to Confluence 2.3. This issue has been resolved in 2.3.1.

Significant issues for existing Confluence 2.3 users

- **Backup import**: Confluence 2.3 users wishing to import backups into their Confluence instance should upgrade to 2.3.1 or apply the patch listed against CONF-7584

Upgrade Procedure

Upgrading Confluence should be pretty easy. If you are upgrading from another Confluence 2.2.x or Confluence 2.3 version, you can find instructions here. *We strongly recommend that you backup your `confluence.home` directory and database before upgrading!*

Upgrading from Confluence 2.1 and earlier

Users upgrading directly from 2.1 or earlier should also read the 2.2 Release Notes for caveats regarding the 2.1 -> 2.2 upgrade.

Changes in 2.3.1

2.3.1 resolves 20 issues. All these issues are listed below:

Errors were reported by the JIRA trusted connection.

- APP_UNKNOWN; Unknown Application: {0}; ["confluence:4557196"]

JIRA Issues (22 issues)

Type	Key	Summary	Status	Resolution
	CONF-1643	Comments to pages should be wrapper to make sure they fit in the page	Resolved	Fixed
	CONF-7222	Create user does not trim() usernames	Resolved	Fixed
	CONF-7483	Mail notification templates contain un-internationalised text	Resolved	Fixed
	CONF-7489	Uninstalling a bundled plugins causes stack trace	Resolved	Fixed
	CONF-7534	Cluster setup workflow could be improved	Resolved	Fixed
	CONF-7546	Page Comparison information shown is wrong / missing	Resolved	Fixed
	CONF-7556	Javascript in create space page disables checkbox without un-checking it	Resolved	Fixed
	CONF-7562	Login link on forgot password page redirects to an error after logging in	Resolved	Fixed
	CONF-7573	Dashboard Exception with Space List macro	Closed	Fixed
	CONF-7574	User Preferences results in an empty screen	Resolved	Fixed
	CONF-7575	Confluence 2.3 - "Vor kurzem Aktualisiert" instead of "Recently Updated"	Resolved	Fixed
	CONF-7578	Missing text ${generalUtil.buildDateString}	Resolved	Fixed
	CONF-7580	Testing language packs made it into the final build... again.	Closed	Fixed
	CONF-7581	Missing Europe/London timezome causes runtime errors	Resolved	Fixed
	CONF-7584	Restore backup fails with "Import failed. null" error	Closed	Fixed
	CONF-7585	osuser2atluser.jsp throws JasperException	Closed	Fixed
	CONF-7589	Upgrade to 2.3 with MySQL causes error - Dialect class not found: bucket.dialect.MySQLDialect	Resolved	Fixed

	CONF-7597	View changes not showing on 2.3	Resolved	Fixed
	CONF-7611	Send a notification to any page / space watchers when an attachment, blogpost, comment, page is removed	Resolved	Fixed
	CONF-7630	Cannot edit layout decorators more than once	Resolved	Fixed
	CONF-7679	Jiraissues macro does not display dates	Resolved	Fixed
	CONF-9266	code Macro not supported in PDF Export	Resolved	Fixed

Release Notes 2.3.2

Confluence 2.3.2 is a maintenance release that resolves an assortment of issues users may have encountered using Confluence 2.3.1

Warning for LDAP users
Confluence 2.3.2 has a serious bug that prevents it from working with LDAP servers. Please use Confluence 2.3.3 instead.

2.3.2 is a free upgrade for all customers who purchased their Confluence license after February 12th, 2006.

Who should upgrade?

Confluence 2.3.2 is a bugfix release which resolves some significant issues in Confluence 2.3.1 Customers running Confluence 2.3.1 and 2.3 should review the list of resolved issues below, and upgrade if this release fixes any problems with their current Confluence installation.

Customers running releases older than 2.3 and wishing to upgrade should use this version also.

Upgrade Procedure

Upgrading Confluence should be pretty easy. If you are upgrading from another Confluence 2.2.x or Confluence 2.3 version, you can find instructions here. *We strongly recommend that you backup your `confluence.home` directory and database before upgrading!*

Upgrading from Confluence 2.1 and earlier

Users upgrading directly from 2.1 or earlier should also read the 2.2 Release Notes for caveats regarding the 2.1 -> 2.2 upgrade.

Changes in 2.3.2

2.3.2 resolves 22 issues. All these issues are listed below:

Errors were reported by the JIRA trusted connection.

- APP_UNKNOWN; Unknown Application: {0}; ["confluence:4557196"]

JIRA Issues (22 issues)				
Type	**Key**	**Summary**	**Status**	**Resolution**
	CONF-6371	New line not persistent between Rich TextT and Wiki Markup	Resolved	Fixed
	CONF-6494	Comment display with the ' (apostrophe) when adding an attachment.	Closed	Fixed
	CONF-7335	livesearch macro spaceKey parameter value is case sensitive	Resolved	Fixed
	CONF-7410	Calendar plugin doesn't work in Clickr theme	Resolved	Fixed
	CONF-7608	Can't 'upgrade' bundled plugins with the Plugin Repository	Resolved	Fixed
	CONF-7654	Custom colour scheme for spaces will not apply	Resolved	Fixed
	CONF-7658	Updates to global colour scheme are not shown in spaces until restart	Resolved	Fixed
	CONF-7663	Searches return no results after initial index of upgraded data, until content is updated	Resolved	Fixed
	CONF-7688	Confluence should handle users without full names (from external user management)	Closed	Fixed
	CONF-7690	Upgrade fails for DB2 due to incorrect dialect	Resolved	Fixed
	CONF-7751	Url Encode spacekey	Resolved	Fixed

	CONF-7753	Duplication of invalid links while saving a page in Rich Text Editor	Resolved	Fixed
	CONF-7798	More non-internationalized text on blogposts macro	Resolved	Fixed
	CONF-7799	Page restrictions don't always show up in brose space, page permissions.	Resolved	Fixed
	CONF-7802	LDAP support doesn't work in cluster	Resolved	Fixed
	CONF-7807	{code} macro sometimes loses whitespace in Wysiwyg	Resolved	Fixed
	CONF-7808	Threads getting stuck in HitTrackingLocalCache, possible synchronisation problem	Resolved	Fixed
	CONF-7813	People Directory throws NullPointerException	Resolved	Fixed
	CONF-7815	Some DBs incorrectly use Postgres lower casing	Resolved	Fixed
	CONF-7834	Remove the people directory link from dashboard when share mode is on	Resolved	Fixed
	CONF-7843	Restoring a site backup can set cluster nodes to use file system attachment storage	Closed	Fixed
	CONF-7858	LDAP pooling defaults are not appropriate	Resolved	Fixed

Release Notes 2.3.3

Confluence 2.3.3 is a maintenance release that resolves an assortment of issues users may have encountered using Confluence 2.3.2

2.3.3 is a free upgrade for all customers who purchased their Confluence license after February 14th, 2006.

Who should upgrade?

Confluence 2.3.3 is a bugfix release which resolves some significant issues in Confluence 2.3.2 Customers running Confluence 2.3.2, 2.3.1 and 2.3 should review the list of resolved issues below, and upgrade if this release fixes any problems with their current Confluence installation.

Customers running releases older than 2.3 and wishing to upgrade should use this version also.

Upgrade Procedure

Upgrading Confluence should be pretty easy. If you are upgrading from another Confluence 2.2.x or Confluence 2.3 version, you can find instructions here. *We strongly recommend that you backup your `confluence.home` directory and database before upgrading!*

Upgrading from Confluence 2.1 and earlier

Users upgrading directly from 2.1 or earlier should also read the 2.2 Release Notes for caveats regarding the 2.1 -> 2.2 upgrade.

Changes in 2.3.3

2.3.3 resolves 4 issues. All these issues are listed below:

Errors were reported by the JIRA trusted connection.

- APP_UNKNOWN; Unknown Application: {0}; ["confluence:4557196"]

JIRA Issues (5 issues)				
Type	**Key**	**Summary**	**Status**	**Resolution**
	CONF-5701	RSS macro ignores 'nonProxyHosts' JVM property	Resolved	Fixed
	CONF-6986	Most Plugin Themes require name and description i18n	Resolved	Fixed
	CONF-7855	Previewing a copied page (before saving) doesn't work	Resolved	Fixed
	CONF-7862	Illegal value for java.naming.referral property	Closed	Fixed
	CONF-8210	Build source release for Confluence 2.3.3	Resolved	Fixed

Release Notes 2.4

The Atlassian Confluence team is proud to present to you Confluence 2.4.2! What happened to 2.4 and 2.4.1? Let's just say we gave them a good workout.

Confluence 2.4 is the ninth release of Confluence. It introduces much awaited support for editable comments, the ability to mail a page to groups of users and includes the usual slew of fixes since the last stable release.

Confluence 2.4 is the first in a series of frequent, small releases planned for the first half or 2007, that will help us get the functionality you want from Confluence, faster. For more information on our new release schedule, you might want to read this blog post.

Confluence 2.4.2 is a **free upgrade** for any customer who purchased Confluence after March 13th, 2006.

Upgrading from Confluence 2.2.x and 2.3.x

Upgrading Confluence should be fairly straightforward: you can find instructions here. *We strongly recommend that you backup your `confluence.home` directory and database before upgrading!*

Upgrading from Confluence 2.1 and earlier

Users upgrading directly from 2.1 or earlier should also read the 2.2 Release Notes for caveats regarding the 2.1 -> 2.2 upgrade.

Contents

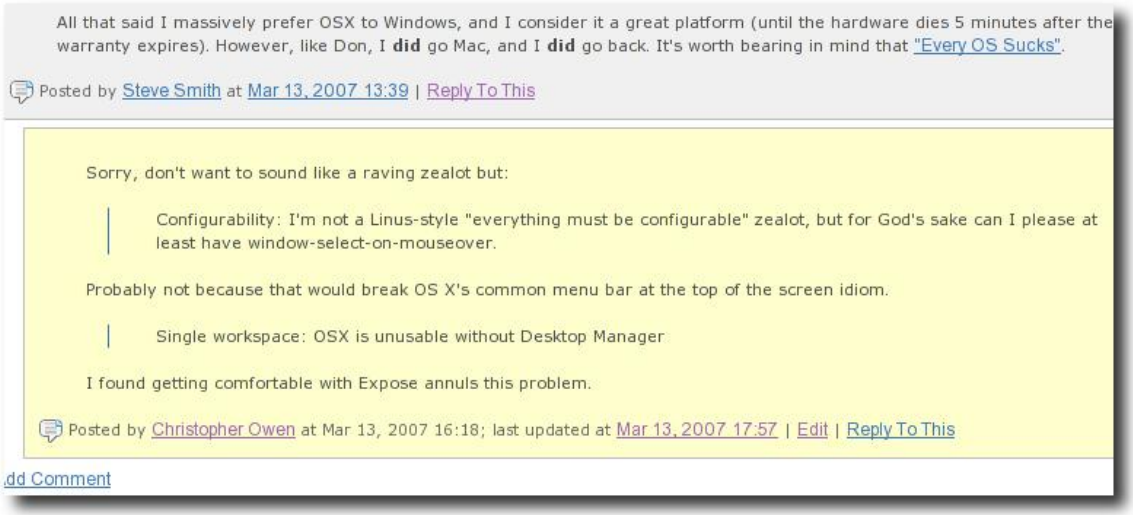
1. New Features
2. Improvements

See also: Issues Resolved for 2.4.2

New Features

Editable Comments

Say *so long!* to those annoying typos and ill-structured sentences as comments in Confluence are now editable. Comments may be edited by the original author or the administrator for the space in which the comment was posted.



Confluence has a ten minute grace period after the comment is posted in which updates will not be flagged in the user interface. This is in acknowledgment of the fact that most edits within this period are for correcting minor typos and formatting which will not impact the flow of conversation.

For more information, see the Confluence user guide: Editing a Comment

Page Mailing

Confluence 2.4 ships with the Mail Page plugin allowing users to conveniently send the contents of a Confluence page to other interested parties. You can easily specify who you want the page mailed to by using a combination of:

- Email addresses
- Confluence user names
- Confluence group names

For security reasons, the Mail Page plugin is disabled by default. If you want to use this feature, an administrator will have to enable it. There must also be a mail server configured in the Confluence instance for this operation to be available.



This operation is reached via a page's Info tab.



For more information, see the Confluence user guide: E-mailing a Page

Improvements

- The People Directory no longer shows inactive users (CONF-7771)
- Log messages now have more context such as the URL of the page being displayed and the name of the logged in user (CONF-7878)

The Confluence 2.4 Team

Development
Tom Davies
Matthew Jensen
Samuel Le Berrigaud
David Loeng
Charles Miller
Christopher Owen
Agnes Ro
Matt Ryall
Don Willis

Oversight & ~~Mis~~ management
Mike Cannon-Brookes
Scott Farquhar

Changes to the Page Permission API in Confluence 2.4

Confluence 2.4 contains changes to the Confluence API that will affect any plugins that modify Page Permissions (aka ContentPermissions). Some methods have been deprecated, some new methods have been created, and one method now behaves differently from before.

Confluence 2.4 does not support having multiple view and edit permissions on a page. The Content Permission API allows such a state, but it should be avoided. For example, if a Page has multiple permissions set on it, these permissions will not be displayed properly via the user interface, and the pages may not be returned appropriately in a search.

Summary of Changes

Deprecated methods

```
ContentPermissionManager.getInheritedViewContentPermissions(Page)
ContentPermissionManager.getInheritedContentPermissions(ContentEntityObject
contentEntityObject)

ContentEntityObject.getPermissions()
ContentEntityObject.getContentPermission(String permissionType)
```

Changed methods

```
ContentPermissionManager.addContentPermission(ContentPermission permission,
ContentEntityObject content)
```

Added methods

```
ContentPermissionManager.setContentPermissions(List contentPermissions, ContentEntityObject
content, String type)
ContentPermissionManager.getInheritedContentPermissionSets(ContentEntityObject
contentEntityObject)

ContentEntityObject.getContentPermissionSet(String type)

ContentEntityObject.hasPermissions(String type)
ContentEntityObject.removeContentPermissionSet(ContentPermissionSet set)
```

Removed methods

```
ContentPermissionManager.saveContentPermission(ContentPermission permission)
ContentPermissionManager.getContentUserPermission(ContentEntityObject contentEntityObject)
ContentPermissionManager.onContentContextChanged(ContentEntityObject contentEntityObject)

ContentEntityObject.setPermissions(List permissions)
ContentEntityObject.clearPermissions(ContentPermissionManager permissionManager)
```

The change most likely to cause problems is:

```
ContentPermissionManager.addContentPermission(ContentPermission permission, ContentEntityObject content)
```

This method now adds a ContentPermission to the ContentEntityObject as one would expect. Previously if a ContentPermission already existed on the ContentEntityObject, then any existing ContentPermission would be removed. That is, the method really behaved as one would expect "setContentPermission" to behave. As a result, any code that used addContentPermission in Confluence 2.3 or earlier, should now use setContentPermissions with a single entry list.

Example conversion of old addContentPermission usage

```
1.replace
2.ContentPermissionManager.addContentPermission(p, c)
3.with
4.ContentPermissionManager.setContentPermissions(java.util.Collections.singletonList(p), c,
p.getType())
```

Most methods that used to return a ContentPermission have been replaced with methods that return a ContentPermissionSet. In some places we have used deprecation to point to the new method and reimplemented the old method to extract the first ContentPermission from the ContentPermissionSet.

As usual, the methods on the ContentPermissionManager should be used rather than those on the ContentEntityObject. The only ContentEntityObject method that should be used is getContentPermissionSet. Even that method should be used only to display the existing Permissions. All writes to the ContentPermissionSets should be performed via the ContentPermissionManager.

The onContentContextChanged method of the ContentPermissionManager was previously used to notify a ContentPermissionManager that cached permissions on a ContentEntityObject may be invalid. This is now accomplished by publishing a ContentPermissionEvent. That event is published automatically when changing ContentPermissions via the ContentPermissionManager.

No changes have been made yet to the the XML RPC API. It does not currently allow modification of ContentPermissions.

Release Notes 2.4.1

The Atlassian Confluence team is proud to present to you Confluence 2.4.2! What happened to 2.4 and 2.4.1? Let's just say we gave them a good workout.

Confluence 2.4 is the ninth release of Confluence. It introduces much awaited support for editable comments, the ability to mail a page to groups of users and includes the usual slew of fixes since the last stable release.

Confluence 2.4 is the first in a series of frequent, small releases planned for the first half or 2007, that will help us get the functionality you want from Confluence, faster. For more information on our new release schedule, you might want to read this blog post.

Confluence 2.4.2 is a **free upgrade** for any customer who purchased Confluence after March 13th, 2006.

Upgrading from Confluence 2.2.x and 2.3.x

Upgrading Confluence should be fairly straightforward: you can find instructions here. *We strongly recommend that you backup your `confluence.home` directory and database before upgrading!*

Upgrading from Confluence 2.1 and earlier

Users upgrading directly from 2.1 or earlier should also read the 2.2 Release Notes for caveats regarding the 2.1 -> 2.2 upgrade.

Contents

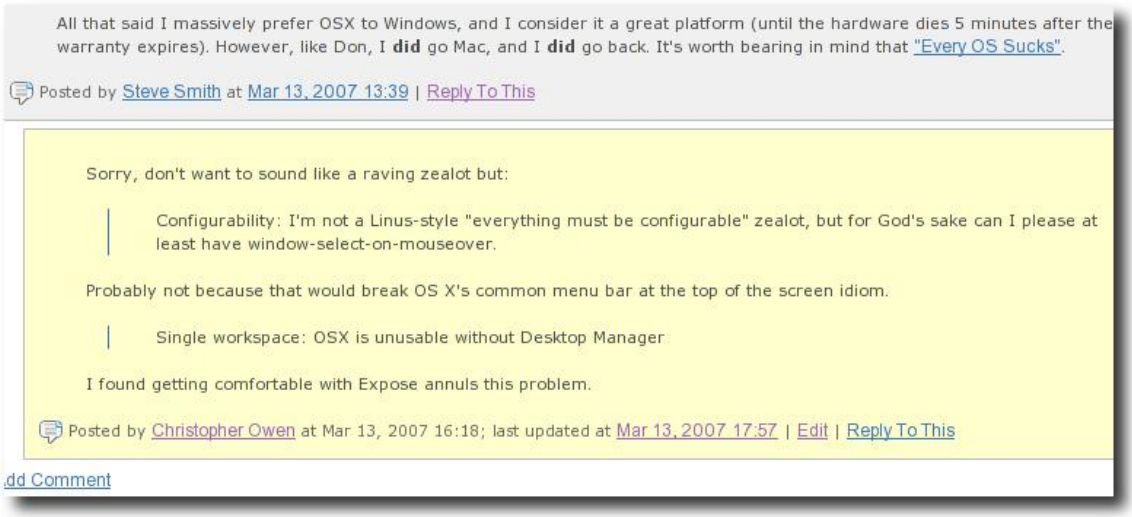
1. New Features
2. Improvements

See also: Issues Resolved for 2.4.2

New Features

Editable Comments

Say *so long!* to those annoying typos and ill-structured sentences as comments in Confluence are now editable. Comments may be edited by the original author or the administrator for the space in which the comment was posted.



Confluence has a ten minute grace period after the comment is posted in which updates will not be flagged in the user interface. This is in acknowledgment of the fact that most edits within this period are for correcting minor typos and formatting which will not impact the flow of conversation.

For more information, see the Confluence user guide: Editing a Comment

Page Mailing

Confluence 2.4 ships with the Mail Page plugin allowing users to conveniently send the contents of a Confluence page to other interested parties. You can easily specify who you want the page mailed to by using a combination of:

- Email addresses
- Confluence user names
- Confluence group names

For security reasons, the Mail Page plugin is disabled by default. If you want to use this feature, an administrator will have to enable it. There must also be a mail server configured in the Confluence instance for this operation to be available.



This operation is reached via a page's Info tab.



For more information, see the Confluence user guide: E-mailing a Page

Improvements

- The People Directory no longer shows inactive users (CONF-7771)
- Log messages now have more context such as the URL of the page being displayed and the name of the logged in user (CONF-7878)

The Confluence 2.4 Team

Development
Tom Davies
Matthew Jensen
Samuel Le Berrigaud
David Loeng
Charles Miller
Christopher Owen
Agnes Ro
Matt Ryall
Don Willis

Oversight & ~~Mis~~ management
Mike Cannon-Brookes
Scott Farquhar

Release Notes 2.4.2

The Atlassian Confluence team is proud to present to you Confluence 2.4.2! What happened to 2.4 and 2.4.1? Let's just say we gave them a good workout.

Confluence 2.4 is the ninth release of Confluence. It introduces much awaited support for editable comments, the ability to mail a page to groups of users and includes the usual slew of fixes since the last stable release.

Confluence 2.4 is the first in a series of frequent, small releases planned for the first half or 2007, that will help us get the functionality you want from Confluence, faster. For more information on our new release schedule, you might want to read this blog post.

Confluence 2.4.2 is a **free upgrade** for any customer who purchased Confluence after March 13th, 2006.

Upgrading from Confluence 2.2.x and 2.3.x

Upgrading Confluence should be fairly straightforward: you can find instructions here. *We strongly recommend that you backup your* `confluence.home` *directory and database before upgrading!*

Upgrading from Confluence 2.1 and earlier

Users upgrading directly from 2.1 or earlier should also read the 2.2 Release Notes for caveats regarding the 2.1 -> 2.2 upgrade.

Contents

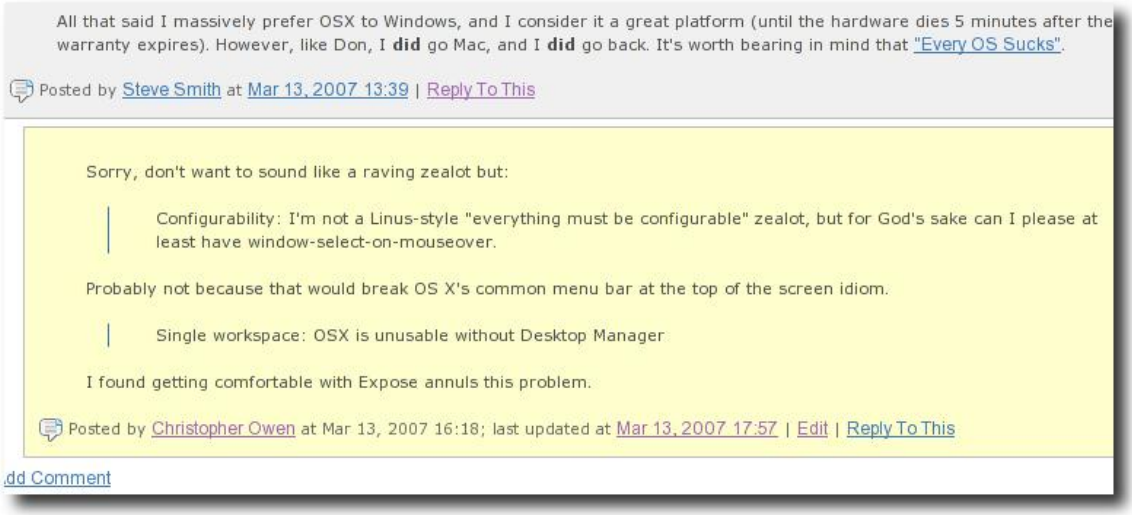
1. New Features
2. Improvements

See also: Issues Resolved for 2.4.2

New Features

Editable Comments

Say *so long!* to those annoying typos and ill-structured sentences as comments in Confluence are now editable. Comments may be edited by the original author or the administrator for the space in which the comment was posted.



Confluence has a ten minute grace period after the comment is posted in which updates will not be flagged in the user interface. This is in acknowledgment of the fact that most edits within this period are for correcting minor typos and formatting which will not impact the flow of conversation.

For more information, see the Confluence user guide: Editing a Comment

Page Mailing

Confluence 2.4 ships with the Mail Page plugin allowing users to conveniently send the contents of a Confluence page to other interested parties. You can easily specify who you want the page mailed to by using a combination of:

- Email addresses
- Confluence user names
- Confluence group names

For security reasons, the Mail Page plugin is disabled by default. If you want to use this feature, an administrator will have to enable it. There must also be a mail server configured in the Confluence instance for this operation to be available.



This operation is reached via a page's Info tab.



For more information, see the Confluence user guide: E-mailing a Page

Improvements

- The People Directory no longer shows inactive users (CONF-7771)
- Log messages now have more context such as the URL of the page being displayed and the name of the logged in user (CONF-7878)

The Confluence 2.4 Team

Development
Tom Davies
Matthew Jensen
Samuel Le Berrigaud
David Loeng
Charles Miller
Christopher Owen
Agnes Ro
Matt Ryall
Don Willis

Oversight & ~~Mis~~ management
Mike Cannon-Brookes
Scott Farquhar

Issues Resolved for 2.4.2

Errors were reported by the JIRA trusted connection.

- APP_UNKNOWN; Unknown Application: {0}; ["confluence:4557196"]

JIRA Issues (38 issues)				
Type	**Key**	**Summary**	**Status**	**Resolution**
	CONF-8164	When exporting a space to restore on another system, change history is lost	Resolved	Fixed
	CONF-8045	Can't delete template with previous version	Closed	Fixed
	CONF-8032	Use HTML id instead of empty named anchors in headings	Resolved	Fixed

	CONF-8029	Outdated event listener interface warning should be logged at a lower priority	Resolved	Fixed
	CONF-8018	Page list template renders broken page title when greater than 64 characters	Resolved	Fixed
	CONF-8014	To provide easier configuration between Crowd and Confluence the attached crowd-ehcache.xml file will need to be added to the confluence release	Resolved	Fixed
	CONF-7989	Fix display for my favourite labels - $webwork.htmlEncode($textUtils.trimToEndingChar($page.realTitle, 60))	Resolved	Fixed
	CONF-7953	CSS and Javascript not loaded in Websphere 6.1.0.5	Resolved	Fixed
	CONF-7939	Confluence atlassian-user.xml has typo for Crowd integration.	Resolved	Fixed
	CONF-7926	Template Lists Fail with Non-English Characters	Closed	Fixed
	CONF-7925	SQLException on Sybase and SQL Server - Invalid column name 'creationDate'	Closed	Fixed
	CONF-7895	Feed builder only builds private feeds	Resolved	Fixed
	CONF-7878	Add RenderContext information to exceptions that filter through the Wiki Renderer	Resolved	Fixed
	CONF-7854	Error deleting template that has been edited	Resolved	Fixed
	CONF-7843	Restoring a site backup can set cluster nodes to use file system attachment storage	Closed	Fixed
	CONF-7815	Some DBs incorrectly use Postgres lower casing	Resolved	Fixed
	CONF-7788	Insert link dialog doesn't search properly	Resolved	Fixed
	CONF-7786	Exclude space group from space export	Resolved	Fixed
	CONF-7783	Don't include user profiles in daily changed reports in shared mode	Resolved	Fixed
	CONF-7770	Add ability to select a space group in the create space form	Resolved	Fixed
	CONF-7769	Update atlassian-extras dependency to 0.7.32	Resolved	Fixed
	CONF-7768	Add "Groups" tab in space list macro on dashboard for Space Groups	Resolved	Fixed
	CONF-7759	Truncate email subjects longer than 255 characters	Closed	Fixed
	CONF-7757	Add importSpace method to RPC interfaces	Resolved	Fixed
	CONF-7739	Some content migrated from 2.0 to 2.3 fails to render (due to:- system error: java.lang.String java.lang.ClassClastException: java.langString)	Resolved	Fixed
	CONF-7690	Upgrade fails for DB2 due to incorrect dialect	Resolved	Fixed
	CONF-7678	Lazy init exception checking attachment permissions	Resolved	Fixed
	CONF-7663	Searches return no results after initial index of upgraded data, until content is updated	Resolved	Fixed
	CONF-7661	README files in confluence-2.3 have URLs that refer to non-existent pages	Resolved	Fixed

	CONF-7608	Can't 'upgrade' bundled plugins with the Plugin Repository	Resolved	Fixed
	CONF-7456	Intermittent problem loading stylesheets and javascript on extranet	Resolved	Fixed
	CONF-7095	Text for the key "licenseexpired.explanation" not being displayed	Resolved	Fixed
	CONF-5391	Links in news items to attachments in normal pages do not work	Closed	Fixed
	CONF-5194	Provide a way to mail to User Groups or Individual from a wiki page	Resolved	Fixed
	CONF-4550	Macro that can render Mathematical Formulae	Resolved	Fixed
	CONF-2492	Email wiki page to a "friend"	Resolved	Fixed
	CONF-1062	Email page	Resolved	Fixed
	CONF-830	Editable comments	Resolved	Fixed

Release Notes 2.4.3

Confluence 2.4.3 is a maintenance release that resolves an assortment of issues users may have encountered using Confluence 2.4.2

2.4.3 is a free upgrade for all customers who purchased their Confluence license after March 22nd 2006.

Who should upgrade?

Confluence 2.4.3 is a recommended upgrade release which resolves some significant issues in Confluence 2.4.2. Customers running older versions of Confluence should review the list of resolved issues below, and upgrade if this release fixes any problems with their current Confluence installation.

Customers running releases older than 2.4 and wishing to upgrade should use this version also.

Upgrade Procedure

Upgrading Confluence should be pretty easy. If you are upgrading from another Confluence 2.2.x or Confluence 2.3 version, you can find instructions here.

We strongly recommend that you backup your `confluence.home` directory and database before upgrading!

Upgrading from Confluence 2.1 and earlier

Users upgrading directly from 2.1 or earlier should also read the 2.2 Release Notes for caveats regarding the 2.1 -> 2.2 upgrade.

Changes in 2.4.3

2.4.3 resolves 15 issues. All these issues are listed below:

Errors were reported by the JIRA trusted connection.

- APP_UNKNOWN; Unknown Application: {0}; ["confluence:4557196"]

JIRA Issues (15 issues)				
Type	**Key**	**Summary**	**Status**	**Resolution**
	CONF-6058	Group picker for page restrictions silently hides non-member groups, may need clarifying sentence added	Resolved	Fixed
	CONF-6505	Comment index entries do not inherit permissions correctly	Closed	Fixed
	CONF-6907	New license types footers	Resolved	Fixed
	CONF-7416	LazyInitializationException on page create preview when page contains un-resolveable embedded image	Closed	Fixed

	CONF-7564	'Advanced' space tab should show 'personal space' icon instead of 'user profile' icon	Closed	Fixed
	CONF-7717	Rich text editor breaks shortcut links with custom titles	Resolved	Fixed
	CONF-7725	IllegalStateException "zip file closed" occurring in plugins	Closed	Fixed
	CONF-7898	AttachmentsSoapService getAttachmentData doesn't close inputStream	Resolved	Fixed
	CONF-7922	Personal Edition - "Evaluate Confluence" Link Broken At Bottom Of All Pages	Closed	Fixed
	CONF-8001	When using LDAP with a cluster, a user cannot log into the second node if they previously logged into the first node	Resolved	Fixed
	CONF-8039	Create Space form gives option of anonymous access, even when not enabled globally	Resolved	Fixed
	CONF-8048	REGRESSION: "View Change" link missing	Closed	Fixed
	CONF-8050	zip_src from tiny mce served without caching headers on extranet	Resolved	Fixed
	CONF-8056	Disabled upgraded bundled plugins are reenabled on startup	Resolved	Fixed
	CONF-8077	Attachment file names not HTML encoded in attachment list	Resolved	Fixed

Release Notes 2.4.4

Confluence 2.4.4 is a recommended upgrade release that resolves issues users may have encountered using Confluence 2.4.3 or earlier.

2.4.4 is a free upgrade for all customers who purchased their Confluence license after March 30th 2006.

Who should upgrade?

Confluence 2.4.4 is a recommended upgrade release that resolves issues users may have encountered using Confluence 2.4.3 or earlier. Customers running older versions of Confluence should review the list of resolved issues below, and upgrade if this release fixes any problems with their current Confluence installation.

Customers running releases older than 2.4 and wishing to upgrade should upgrade directly to 2.4.4.

Upgrade Procedure

Upgrading Confluence should be pretty easy. If you are upgrading from another Confluence 2.2.x or Confluence 2.3 version, you can find instructions here.

We strongly recommend that you backup your `confluence.home` directory and database before upgrading!

Upgrading from Confluence 2.1 and earlier

Users upgrading directly from 2.1 or earlier should also read the 2.2 Release Notes for caveats regarding the 2.1 -> 2.2 upgrade.

Source Release

Please note that the source release for 2.4.4 is currently unavailable. We have scheduled a fix for this which should be included in 2.4.5. Please see CONF-8007 for more details.

Changes in 2.4.4

2.4.4 resolves 17 issues. All these issues are listed below:

Errors were reported by the JIRA trusted connection.

- APP_UNKNOWN; Unknown Application: {0}; ["confluence:4557196"]

JIRA Issues (19 issues)				
Type	**Key**	**Summary**	**Status**	**Resolution**

	CONF-6490	PDF export breaks with angle brackets inside comments	Resolved	Fixed
	CONF-7711	Only latest page history comment is shown in page history	Resolved	Fixed
	CONF-7774	Global Activity link on Space Activity page is not correct	Resolved	Fixed
	CONF-7775	Cross site scripting - action name not escaped in group picker	Resolved	Fixed
	CONF-7776	Cross site scripting - space key not escaped in listpages-alphaview	Resolved	Fixed
	CONF-7777	Cross site scripting - destination not escaped on login page error message	Resolved	Fixed
	CONF-7778	Cross site scripting - on POST, title not escaped in createpage-entervariables	Resolved	Fixed
	CONF-8065	Cross Site Scripting issue when integration RSS feeds	Resolved	Fixed
	CONF-8101	XSS on "Site Search"	Resolved	Fixed
	CONF-8108	Attempted Space Removal caused a database exception.	Resolved	Fixed
	CONF-8124	"Space Activity" XSS hole and Exception throwing	Closed	Fixed
	CONF-8128	Confluence Repository Client broke with 2.4.3 upgrade	Resolved	Fixed
	CONF-8143	Avatar upload - no HTML tags encoding in filenames	Resolved	Fixed
	CONF-8144	XSS on User Search	Resolved	Fixed
	CONF-8145	View mail thread icon link navigates to blank page	Resolved	Fixed
	CONF-8159	Enable comment content layout to be editable in admin	Resolved	Fixed
	CONF-8162	Rich text editor fails to load because ConfluenceTinyMCEServlet doesn't notice base URL changes	Closed	Fixed
	CONF-8167	Previous page versions view lists current revision twice	Resolved	Fixed
	CONF-8178	Download links broken for Confluence 2.4.4	Closed	Fixed

Release Notes 2.4.5

Confluence 2.4.5 is a maintenance release that resolves issues users may have encountered using Confluence 2.4.4 or earlier.

2.4.5 is a free upgrade for all customers who purchased their Confluence license after April 12th 2006.

Who should upgrade?

Confluence 2.4.5 is a recommended upgrade release which resolves some significant issues in Confluence 2.4.4. Customers running older versions of Confluence should review the list of resolved issues below, and upgrade if this release fixes any problems with their current Confluence installation.

Customers running releases older than 2.4 and wishing to upgrade should upgrade directly to 2.4.5.

Upgrade Procedure

Upgrading Confluence should be pretty easy. If you are upgrading from another Confluence 2.2.x or Confluence 2.3 version, you can find instructions here.

We strongly recommend that you backup your `confluence.home` directory and database before upgrading!

Upgrading from Confluence 2.1 and earlier

Users upgrading directly from 2.1 or earlier should also read the 2.2 Release Notes for caveats regarding the 2.1 -> 2.2 upgrade.

Changes in 2.4.5

2.4.5 resolves these issues:

Errors were reported by the JIRA trusted connection.

- APP_UNKNOWN; Unknown Application: {0}; ["confluence:4557196"]

JIRA Issues (14 issues)				
Type	Key	Summary	Status	Resolution
	CONF-5756	After moving a page with attachments to a new Space, the attachments are listed in the old Space not on the new Space	Resolved	Fixed
	CONF-6068	So when/how do you fix the problem?	Closed	Fixed
	CONF-6461	Can see other Users Profile/Personal Spaces hybrid in Recently Updated even if you dont have permissions to view them.	Resolved	Fixed
	CONF-7606	Incorrect associated <label> tags on administration page	Resolved	Fixed
	CONF-7832	Documenation on adding a template to a space is not up to date for 2.3	Resolved	Fixed
	CONF-8007	Fix confluence source release	Resolved	Fixed
	CONF-8078	Error page if new password doesn't match Crowd password validation	Resolved	Fixed
	CONF-8125	References to scriptaculous files are wrong	Resolved	Fixed
	CONF-8185	Cached TinyMCE Servlet doesn't handle BaseURL Changes or Multi-homed environment	Resolved	Fixed
	CONF-8228	On/off list components have inappropriate HTML IDs	Resolved	Fixed
	CONF-8393	Plugin repository client 2.0.2 fails to update plugins	Resolved	Fixed
	CONF-8426	Clicking on Preferences->Edit Profile->Email->Cancel leads to a blank screen	Resolved	Fixed
	CONF-8501	Unable to load plugins through the plugin repository - Status Code 302	Resolved	Fixed
	CONF-8916	Create patch to CONF-8877 for Confluence 2.5.4	Resolved	Fixed

Release Notes 2.5

The Atlassian Confluence team is proud to present Confluence 2.5. Confluence 2.5 is the tenth release of Confluence. It introduces more flexible page restrictions, and also includes a number of minor fixes.

Confluence 2.5 is a **free upgrade** for any customer who purchased Confluence after April 27th, 2006.

Upgrading from Confluence 2.2 and later

Upgrading Confluence should be fairly straightforward: you can find instructions here. *We strongly recommend that you backup your `confluence.home` directory and database before upgrading!*

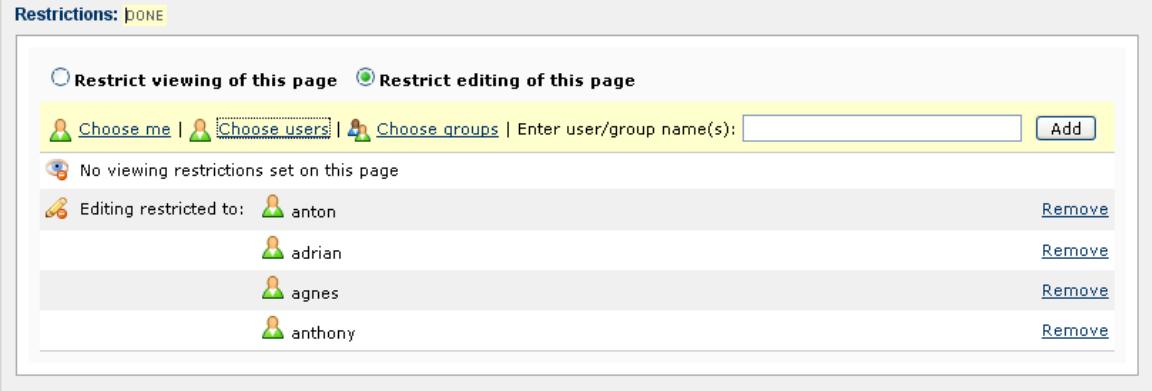
Upgrading from Confluence 2.1 and earlier

Users upgrading directly from 2.1 or earlier should also read the 2.2 Release Notes for caveats regarding the 2.1 -> 2.2 upgrade.

Introducing flexible page restrictions

Page restrictions give you the ability to choose who may read and/or edit any given page.

By popular request, page restrictions have been enhanced and can now be assigned to multiple people and/or groups:



Other significant issues resolved

Dynamic task list JRE incompatibilities

Many customers have been prevented from upgrading the Java runtime hosting their Confluence instance due to CONF-4082, an issue which would prevent existing dynamic task lists from working with the new JRE. This issue has been resolved in Confluence 2.5 allowing dynamic task list users to upgrade and downgrade between Java versions without experiencing conversion exceptions.

`contentbylabel` macro supports AND condition

You may now use the `operator=AND` parameter with the `contentbylabel` macro to select pages that have all of the supplied labels. (CONF-4969)

The Confluence 2.5 Team

Development
Tom Davies
Matthew Jensen
Samuel Le Berrigaud
David Loeng
Charles Miller
Christopher Owen
Agnes Ro
Matt Ryall
Don Willis

Oversight & ~~Mis~~ management
Mike Cannon-Brookes
Scott Farquhar

Issues Resolved for 2.5

Errors were reported by the JIRA trusted connection.

- APP_UNKNOWN; Unknown Application: {0}; ["confluence:4557196"]

JIRA Issues (48 issues)				
Type	**Key**	**Summary**	**Status**	**Resolution**
	CONF-8810	Deadlock when working with Bandana table, database locks and thread locks	Resolved	Fixed
	CONF-8426	Clicking on Preferences->Edit Profile->Email->Cancel leads to a blank screen	Resolved	Fixed
	CONF-8403	The code macro inserts a new line at its end	Resolved	Fixed
	CONF-8393	Plugin repository client 2.0.2 fails to update plugins	Resolved	Fixed
	CONF-8330	user X moving a page created by user Y triggers a notification for each child in heirarchy, attributed to user Y	Resolved	Fixed
	CONF-8321	Implement setting page permissions thru XML-RPC and SOAP	Resolved	Fixed

	CONF-8301	recently-updated throws NullPointerException when no pages with label found	Resolved	Fixed
	CONF-8299	Remove group or user should remove content permissions that are assigned to them	Resolved	Fixed
	CONF-8285	HTML Blogpost navigation have a trailing '}' in the space link.	Resolved	Fixed
	CONF-8283	remove comment notification wrongly claims that the comment author is the comment remover	Resolved	Fixed
	CONF-8282	Newline collapsing between horizontal rule and the following element screws things up	Resolved	Fixed
	CONF-8278	Create Space Button disables embedded images	Resolved	Fixed
	CONF-8275	Info macros help not internationalized	Resolved	Fixed
	CONF-8271	Confluence tries to index non text files as text, e.g. .wnk files crunch search	Resolved	Fixed
	CONF-8269	Bundled plugins without internationalized help text.	Resolved	Fixed
	CONF-8259	Anonymous user should not be allowed to set page permissions	Closed	Fixed
	CONF-8255	Links of PDF export from Confluence Space are all web links instead of local links	Resolved	Fixed
	CONF-8250	JIRA issues macro should use nofollow for refresh and main title link	Resolved	Fixed
	CONF-8242	create-space-button macro throws NPE when using preview	Resolved	Fixed
	CONF-8221	Using page mailing, page link is not ok for news	Resolved	Fixed
	CONF-8206	Confluence jar shipped in standalone lib directory	Resolved	Fixed
	CONF-8205	Custom SQL query on ancestors table breaks Firebird	Resolved	Fixed
	CONF-8150	Email search is case sensitive, makes search difficult.	Resolved	Fixed
	CONF-8111	Livesearch throws exception if search term contains spaces	Resolved	Fixed
	CONF-8043	Better Crowd Integration (from a user's perspective)	Resolved	Fixed
	CONF-7974	Anonymous additons are attributed to $page.creatorName in text update notifications	Resolved	Fixed
	CONF-7966	'add comment' and 'remove comment' notification do not link to the page that the comment was attached to	Resolved	Fixed
	CONF-7912	Rich Text Editor: Bullets and text in a table cell	Resolved	Fixed
	CONF-7893	Link to Plugin Repository is 'plugin.repository.link'	Resolved	Fixed
	CONF-7683	Upgrade the bundled Blogging RPC plugin to version 1.1.1	Resolved	Fixed
	CONF-7628	Rich Text Editor changes spacing around user-macros	Resolved	Fixed
	CONF-7333	Blank line between bullets is lost	Resolved	Fixed
	CONF-6196	Copying a page with an image on it causes an error		Fixed

			Resolved	
	CONF-6057	Users can manually restrict pages operations to custom groups of which they are not members	Resolved	Fixed
	CONF-5783	Markup with emoticon should have whitespace around it	Resolved	Fixed
	CONF-5682	User level permission on pages	Resolved	Fixed
	CONF-5677	Javascript exception NS_ERROR_NOT_AVAILABLE when changing editor tabs in Firefox	Resolved	Fixed
	CONF-5334	Hovering over the tabs in the editor view causes page to scroll on IE	Resolved	Fixed
	CONF-5159	Draft saving incorrectly encodes characters in Safari	Resolved	Fixed
	CONF-5134	Blogpost macro - order by created date only	Resolved	Fixed
	CONF-5080	Administrators cannot restrict View/Edit of a page to a group that the administrator does not belong	Resolved	Fixed
	CONF-4986	Pages with permissions are not being listed under 'Restricted Pages' in the space admin screens.	Resolved	Fixed
	CONF-4969	contentbylabel macro should support AND condition	Resolved	Fixed
	CONF-4082	ConversionException on dynamic tasklist after JDK version change	Resolved	Fixed
	CONF-4005	Page permission info is misleading	Resolved	Fixed
	CONF-3908	Can not set page permissions via Remote API	Resolved	Fixed
	CONF-3759	The login.action page should forward to the homepage when logged in.	Closed	Fixed
	CONF-3701	Allow selection of multiple users and multiple groups (or both) for page level permissions	Resolved	Fixed

Release Notes 2.5.1

Confluence 2.5.1 is a maintenance release that resolves issues users may have encountered using Confluence 2.5 or earlier.

2.5.1 is a free upgrade for all customers who purchased their Confluence license after May 7th 2006.

Who should upgrade?

Confluence 2.5.1 is a bugfix release which resolves some significant issues in Confluence 2.5. Customers running older versions of Confluence should review the list of resolved issues below, and upgrade if this release fixes any problems with their current Confluence installation.

Customers running releases older than 2.5 and wishing to upgrade should upgrade directly to 2.5.1.

Upgrade Procedure

Upgrading Confluence should be pretty easy. If you are upgrading from another Confluence 2.2.x or a later version, you can find instructions here.

We strongly recommend that you backup your `confluence.home` directory and database before upgrading!

Upgrading from Confluence 2.1 and earlier

Users upgrading directly from 2.1 or earlier should also read the 2.2 Release Notes for caveats regarding the 2.1 -> 2.2 upgrade.

Changes in 2.5.1

2.5.1 resolves these issues:

Errors were reported by the JIRA trusted connection.

- APP_UNKNOWN; Unknown Application: {0}; ["confluence:4557196"]

JIRA Issues (26 issues)				
Type	**Key**	**Summary**	**Status**	**Resolution**
	CONF-3345	Password Reminder will change password even though the email was not send	Resolved	Fixed
	CONF-3426	AND search doesn't match if part of match in title and part in body.	Resolved	Fixed
	CONF-3453	{include} of a page in another space which has {blog-posts} renders error in blog list	Closed	Fixed
	CONF-6049	Export of Documentation Space to PDF and XML restoration are broken	Resolved	Fixed
	CONF-6625	Time elapsed keeps counting after reindexing is complete	Resolved	Fixed
	CONF-7367	Only one livesearch macro functions per page	Resolved	Fixed
	CONF-7621	Activity Plugin macro notation doco needs to be created	Resolved	Fixed
	CONF-7985	Rich Text Editor - Improper handling of Line Feed in {code} parts	Resolved	Fixed
	CONF-8202	Make 'Anonymous' reserve key word for username	Resolved	Fixed
	CONF-8238	Header anchors do not work in Firefox with non-ASCII characters	Resolved	Fixed
	CONF-8249	Rename 'Maximum Attachments per Form' setting to "... per Upload"	Resolved	Fixed
	CONF-8331	Wiki to HTML Conversion is Slow	Resolved	Fixed
	CONF-8334	Create space AJAX validation doesn't use context path	Resolved	Fixed
	CONF-8337	anchor link broken when moving a page acroos spaces	Resolved	Fixed
	CONF-8342	Leftnav, Plain Website and Tabless themes have broken link to view entire email thread	Resolved	Fixed
	CONF-8343	Poor quality of Thumbnails	Resolved	Fixed
	CONF-8380	Activity plugin report page includes unescaped angle brackets	Resolved	Fixed
	CONF-8392	Sort favourite spaces alphabetically in search drop-down	Resolved	Fixed
	CONF-8402	README.txt contains out of date information about support	Resolved	Fixed
	CONF-8411	listlabels macro behaves like recently-used-labels macro if no space key is provided	Resolved	Fixed
	CONF-8417	Missing plugin-info knocks Confluence over	Resolved	Fixed
	CONF-8423	NullPointerException when editing a group from Manage Groups	Closed	Fixed
	CONF-8424	Fix CONF-6733 for RPC method removeGroup()	Resolved	Fixed
	CONF-8431	ConfigurationException when attachments are uploaded from the "insert link" icon on rich-text editor	Resolved	Fixed
	CONF-8434	Exporting spaces with pages containing a ContentPermission may throw an exception		Fixed

			Resolved	
	CONF-8859	{toc} macro not working for Japanese links	Resolved	Fixed

Release Notes 2.5.2

Confluence 2.5.2 is a maintenance release that resolves issues users may have encountered using Confluence 2.5 or earlier.

2.5.2 is a free upgrade for all customers who purchased their Confluence license after May 22nd 2006.

Who should upgrade?

Confluence 2.5.2 is a bugfix release which resolves some significant issues in Confluence 2.5. Customers running older versions of Confluence should review the list of resolved issues below, and upgrade if this release fixes any problems with their current Confluence installation.

Customers running releases older than 2.5 and wishing to upgrade should upgrade directly to 2.5.2.

Upgrade Procedure

Upgrading Confluence should be pretty easy. If you are upgrading from another Confluence 2.2.x or a later version, you can find instructions here.

We strongly recommend that you backup your `confluence.home` directory and database before upgrading!

Upgrading from Confluence 2.1 and earlier

Users upgrading directly from 2.1 or earlier should also read the 2.2 Release Notes for caveats regarding the 2.1 -> 2.2 upgrade.

Changes in 2.5.2

2.5.2 resolves these issues:

Errors were reported by the JIRA trusted connection.

- APP_UNKNOWN; Unknown Application: {0}; ["confluence:4557196"]

JIRA Issues (18 issues)				
Type	**Key**	**Summary**	**Status**	**Resolution**
	CONF-5527	Redirect to parent page or space homepage after user deletes page	Resolved	Fixed
	CONF-5907	When using the 'copy page' feature, page labels from the original page are not copied to the new page.	Resolved	Fixed
	CONF-7924	Out-of-date Confluence upgrade fails with no useful explanation	Resolved	Fixed
	CONF-8152	Link to results in other spaces given when searching all spaces	Resolved	Fixed
	CONF-8170	A link within a comment breaks when the linked page is renamed	Resolved	Fixed
	CONF-8297	Attachments macro doesn't list previous versions of attachments though "old' property is set to true	Resolved	Fixed
	CONF-8332	Update 'latest' javadoc symlink	Resolved	Fixed
	CONF-8407	"Restore backup" does not detect invalid ZIP files, exceptions are rendered	Resolved	Fixed
	CONF-8450	On the user management page, if you click "Search" without entering results you get a java exception	Resolved	Fixed
	CONF-8451	Duplicate javamail / mail JARs	Resolved	Fixed
	CONF-8454	Hyperlinks containing round brackets are broken	Resolved	Fixed

	CONF-8462	Links produced by the {index} macro aren't correct when a space is exported to HTML	Resolved	Fixed
	CONF-8468	Members of groups in Group Management are printed in ugly technical way	Resolved	Fixed
	CONF-8473	Daily backup files not overwritten if daily backup file name is static	Resolved	Fixed
	CONF-8482	Remote editing causes disappearance and appearance of the content of the excerpt	Resolved	Fixed
	CONF-8491	A small typo on the mailto link on the Administration view user profile.	Resolved	Fixed
	CONF-8506	Preformatted text from the Rich Text Editor removes link properties	Resolved	Fixed
	CONF-8533	Rich-Text-Editor failed to load for some users	Resolved	Fixed

Release Notes 2.5.3

Confluence 2.5.3 is a maintenance release that resolves issues users may have encountered using Confluence 2.5 or earlier.

2.5.3 is a free upgrade for all customers who purchased their Confluence license after May 30th 2006.

Who should upgrade?

Confluence 2.5.3 is a bugfix release which resolves some significant issues in Confluence 2.5. Customers running older versions of Confluence should review the list of resolved issues below, and upgrade if this release fixes any problems with their current Confluence installation.

In particular, 2.5.3 fixes a problem which could cause indexing to fail when extracting text from unprintable encrypted PDF files.

Customers running releases older than 2.5 and wishing to upgrade should upgrade directly to 2.5.3.

Shared Mode Removed

This release removes the 'Shared Mode' setting from General Configuration. If you have shared mode enabled you should disable it before upgrading.

Upgrade Procedure

Upgrading Confluence should be pretty easy. If you are upgrading from another Confluence 2.2.x or a later version, you can find instructions here.

We strongly recommend that you backup your `confluence.home` directory and database before upgrading!

Upgrading from Confluence 2.1 and earlier

Users upgrading directly from 2.1 or earlier should also read the 2.2 Release Notes for caveats regarding the 2.1 -> 2.2 upgrade.

Changes in 2.5.3

2.5.3 resolves these issues:

Errors were reported by the JIRA trusted connection.

- APP_UNKNOWN; Unknown Application: {0}; ["confluence:4557196"]

JIRA Issues (21 issues)				
Type	**Key**	**Summary**	**Status**	**Resolution**
	CONF-3892	Body of mail is going missing.	Resolved	Fixed
	CONF-3913	HTML Export: Images with link are lost	Resolved	Fixed
	CONF-4303	Mail Archiver displays certain emails with missing bodies or content	Resolved	Fixed
	CONF-5954	ConfluenceSoapServiceImpl.changeMyPassword() incorrectly delegates to SoapServiceDelegator.changeUserPassword	Resolved	Fixed

	CONF-6342	body of an excerpt macro is not rendered in a page but rendered when included in other macros such as excerpt-include	Resolved	Fixed
	CONF-7080	Unable to link to attachments on blog posts with the embedded resource renderer	Resolved	Fixed
	CONF-8055	Document process for moving from evaluation to commercial cluster license	Resolved	Fixed
	CONF-8163	Space creation broken when running Confluence in Shared Mode and not having space groups	Resolved	Fixed
	CONF-8481	when changing parent page info of a page, page's 'last edited by' info changes username, but not the date?	Resolved	Fixed
	CONF-8489	Dynamic Tasklist does not work with a "-" sign in the Title of the Tasklist	Resolved	Fixed
	CONF-8509	$userName is displayed one or more times when browsing the members of a group	Resolved	Fixed
	CONF-8523	Edit Space Permissions failure in IE7 (works in FireFox 2.0.0.3)	Resolved	Fixed
	CONF-8535	Cannot delete my:favorite labels from pages	Resolved	Fixed
	CONF-8536	Fully support 'my:favorite' and 'my:favourite'	Resolved	Fixed
	CONF-8543	Junit macro broken	Resolved	Fixed
	CONF-8544	Underscores used to work in template values but now don't	Resolved	Fixed
	CONF-8555	Check version of HSQL	Resolved	Fixed
	CONF-8558	Can not see the members of a group	Resolved	Fixed
	CONF-8560	Querying for an LDAP group retrieves all attributes of the group including all its members	Closed	Fixed
	CONF-8580	Indexing unprintable/encrypted PDFs fails	Resolved	Fixed
	CONF-8593	DefaultBreadcrumbsManager generates invalid HTML	Resolved	Fixed

Release Notes 2.5.4

Confluence 2.5.4 is a maintenance release that resolves issues users may have encountered using Confluence 2.5 or earlier.

2.5.4 is a free upgrade for all customers who purchased their Confluence license after June 13th 2006.

Who should upgrade?

Confluence 2.5.4 is a bugfix release which resolves some significant issues in Confluence 2.5. Customers running older versions of Confluence should review the list of resolved issues below, and upgrade if this release fixes any problems with their current Confluence installation.

Customers running releases older than 2.5 and wishing to upgrade should upgrade directly to 2.5.4.

Upgrade Procedure

Upgrading Confluence should be pretty easy. If you are upgrading from another Confluence 2.2.x or a later version, you can find instructions here.

We strongly recommend that you backup your `confluence.home` directory and database before upgrading!

Upgrading from Confluence 2.1 and earlier

Users upgrading directly from 2.1 or earlier should also read the 2.2 Release Notes for caveats regarding the 2.1 -> 2.2 upgrade.

Changes in 2.5.4

2.5.4 resolves these issues:

Errors were reported by the JIRA trusted connection.

- APP_UNKNOWN; Unknown Application: {0}; ["confluence:4557196"]

JIRA Issues (26 issues)

Type	Key	Summary	Status	Resolution
	CONF-4743	HTML space export does not use either global or space layout for index page	Resolved	Fixed
	CONF-5931	Fix Confluence for Turkish locale	Resolved	Fixed
	CONF-6250	Access to array items in code macro is recognized as undefined page	Closed	Fixed
	CONF-6284	Pages served over HTTPS that embed Flash movies display a security warning on IE	Resolved	Fixed
	CONF-6745	JavaScript error on Create Page with opened "Labels" section	Resolved	Fixed
	CONF-7643	Searching for pages with a certain label in ONE space returns the pages related to the label in ALL spaces (when clicking on "Next>>" link on first result page)	Resolved	Fixed
	CONF-7732	Shortcuts to with certain chars in them do not resolve correctly	Resolved	Fixed
	CONF-7739	Some content migrated from 2.0 to 2.3 fails to render (due to:- system error: java.lang.String java.lang.ClassClastException: java.langString)	Resolved	Fixed
	CONF-7864	Remove trailing and prefacing empty character in SPACE name	Resolved	Fixed
	CONF-7892	Syntax error in usage-stats.vm in Confluence Usage Stats plugin	Resolved	Fixed
	CONF-7970	Labels that are no longer associated with any content should not be displayed	Resolved	Fixed
	CONF-8112	slashes in paths for @shortcuts links are traslated into %2F	Resolved	Fixed
	CONF-8176	Updates to the Usage-Tracking-Plugin do not appear in the Repostiory Client	Resolved	Fixed
	CONF-8190	Size parameter not trimmed in {create-space-button} macro	Resolved	Fixed
	CONF-8243	Global Activity Title not rendered	Resolved	Fixed
	CONF-8395	Bundle WebDAV 1.1 plugin	Resolved	Fixed
	CONF-8459	Export Layouts don't work for Spaces, only for Site	Resolved	Fixed
	CONF-8554	Misleading errormessage when trying to edit a nonexisting users group (via direct URL-access)	Resolved	Fixed
	CONF-8586	Creating a page on an not authorized space	Resolved	Fixed
	CONF-8625	Shortcut links white spaces are changed to + for file:// links	Resolved	Fixed
	CONF-8630	errors.jsp should return HTTP 500 Server Error instead of HTTP 200 OK when errors are present	Resolved	Fixed
	CONF-8634	When 'AND' is used to search, the label 'and' is matched	Resolved	Fixed
	CONF-8658	In rare cases new users don't get added to the proper group and therefore can't use		Fixed

		Confluence	Resolved	
	CONF-8665	exporting a space to XML incorrectly includes comments, even when 'Include comments' is deselected.	Resolved	Fixed
	CONF-8668	JavaScript error ('null' is null or not an object) when adding (saving) a new page	Resolved	Fixed
	CONF-8682	"Not Permitted" error when I try to use the time sheet template	Resolved	Fixed

Release Notes 2.5.5

Confluence 2.5.5 is a recommended upgrade that resolves issues you may have encountered in Confluence 2.5.4 or earlier. Confluence 2.5.5 resolves two security bugs, related to space permissions and invalid characters in space names and keys.

As such **this release is a recommended upgrade for all customers**.

Confluence 2.5.5 is a free upgrade for all customers who purchased their Confluence license or maintenance renewal after July 26th 2006.

Upgrading to Confluence 2.5.5

Upgrading Confluence should be fairly straightforward. You can find instructions here. *We strongly recommend that you backup your `confluence.home` directory and database before upgrading!*

Changes in 2.5.5

Patch for security issues

Confluence 2.5.5 resolves two security bugs, related to:

- space permissions
- invalid characters in space names and keys.

For more information, please see the security advisory.

Server ID

Starting with release 2.5.5, Confluence will generate a server ID for you. Server ID has replaced License ID on the License Details page. You will find the server ID useful when contacting Atlassian support.

Server ID:

- is generated when you install Confluence for the first time
- exists for the life of the Confluence instance
- survives an upgrade
- is held in the database
- is not bound to a specific licence
- is the same for all servers in a cluster.

Translations for rich text editor now compatible with Confluence language pack plugins

Thank you to everyone who voted for this popular fix.

The rich text editor in Confluence shows text in tooltips, warnings and other messages. If you are using Confluence in a language other than English, you will want to translate these messages as well as the standard Confluence text.

With Confluence 2.5.5:

- The translations for the rich text editor can be part of a Confluence language pack plugin.
- If your language pack does not contain translations for the rich text editor, the text will show in English. (Before 2.5.5, it showed the 'key' value rather than English.)
- Partial translations of rich text editor messages are already included in the professional French and German language packs. Full translation is under construction.

This makes things much simpler!

You will find more information here:

- Overview of language pack translations
- Technical overview of language pack plugins
- Specific information on translating the rich text editor

Other fixes in 2.5.5

2.5.5 resolves these issues:

Errors were reported by the JIRA trusted connection.

- APP_UNKNOWN; Unknown Application: {0}; ["confluence:4557196"]

JIRA Issues (19 issues)

Type	Key	Summary	Assignee	Reporter	Priority	Status	Resolution	Create
	CONF-8992	Confluence system error exporting a page as PDF	Unassigned	Igor Mironov		Closed	Cannot Reproduce	Jul 24, 2007
	CONF-8917	XSS vulnerability: space name and key not validated nor escaped	Unassigned	Samuel Le Berrigaud [Atlassian]		Resolved	Fixed	Jul 16, 2007
	CONF-8916	Create patch to CONF-8877 for Confluence 2.5.4	Samuel Le Berrigaud [Atlassian]	Samuel Le Berrigaud [Atlassian]		Resolved	Fixed	Jul 16, 2007
	CONF-8877	Security issue: user can copy page with only view permissions	Samuel Le Berrigaud [Atlassian]	Jon Nermut		Resolved	Fixed	Jul 10, 2007
	CONF-8868	Userlister plugin causes NullPointerException	Unassigned	Tuomas Jormola		Resolved	Fixed	Jul 09, 2007
	CONF-8864	RSS feed doesn't filter personal labels	Charles Miller [old account, do not assign issues]	Charles Miller [old account, do not assign issues]		Resolved	Fixed	Jul 08, 2007
	CONF-8860	Naming a page with a single dot locks page or space	Charles Miller [old account, do not assign issues]	François Nonnenmacher		Resolved	Fixed	Jul 06, 2007
	CONF-8855	Page title length is not being validated, leading to errorpage for titles longer than 255 characters	Samuel Le Berrigaud [Atlassian]	Igor Minar		Resolved	Fixed	Jul 05, 2007
	CONF-8851	"Next" link appears on the "Manage Groups" screen for duplicate groupnames	Charles Miller [old account, do not assign issues]	Gurleen Anand [Atlassian]		Resolved	Fixed	Jul 05, 2007
	CONF-8846	Input for Labels is not properly validated	Christopher Owen [Atlassian]	Igor Minar		Resolved	Fixed	Jul 05, 2007
	CONF-8830	Stand alone tomcat server.xml has enableLookups="true"	Unassigned	Tom Davies [Atlassian]		Resolved	Fixed	Jul 03, 2007
	CONF-8789	Add Server ID implementation	Christopher Owen [Atlassian]	Christopher Owen [Atlassian]		Resolved	Fixed	Jun 28 2007
	CONF-8770	Email address exposure - email hidding option is ignored in user lookup	Don Willis [Atlassian]	Igor Minar		Resolved	Fixed	Jun 26 2007
	CONF-8703	contentbylabel operator=AND performs like an OR	Unassigned	Fennie Ng [Atlassian]		Resolved	Fixed	Jun 14 2007
	CONF-8657	Rich Text Editor broken with IE6 and French language pack	Unassigned	Don Willis [Atlassian]		Resolved	Fixed	Jun 06 2007
	CONF-8123	Using attachments syntax on space template results in NPE	Charles Miller [old account, do not assign issues]	Lothar Hegebart		Resolved	Fixed	Mar 22 2007
	CONF-8092	Rescan for new plugins in {1} - ConfluenceActionSupport_en_US.properties	Unassigned	Jim Clark		Resolved	Fixed	Mar 19 2007
	CONF-6987	Simplify Rich Text Editor localisation	Unassigned	Don Willis [Atlassian]		Resolved	Fixed	Sep 14 2006
	CONF-6167	Javascript errors when draft/heartbeat is interrupted can cause Internet Explorer to crash	Unassigned	Matt Ryall [Atlassian]		Resolved	Fixed	May 15 2006

Upgrade Guide 2.5.5

Who should upgrade?

Confluence 2.5.5 is a recommended upgrade that resolves issues you may have encountered in Confluence 2.5.4 or earlier. Confluence 2.5.5 resolves two security bugs, related to space permissions and invalid characters in space names and keys.

As such **this release is a recommended upgrade for all customers**.

Confluence 2.5.5 is a free upgrade for all customers who purchased their Confluence license or maintenance renewal after July 26th 2006.

If you are running a release older than 2.5 and you want to upgrade, you should upgrade directly to 2.5.5. Refer to the upgrade procedure below.

Upgrade Procedure

Upgrading Confluence should be pretty easy. *We strongly recommend that you backup your `confluence.home` directory and database before upgrading!*

If you are upgrading from Confluence 2.2.x or a later version, you can find instructions here.

If you are upgrading directly from 2.1 or earlier, you should also read the 2.2 Release Notes for warnings about the 2.1 -> 2.2 upgrade.

Release Notes 2.5.6

Confluence 2.5.6 is a recommended upgrade that resolves a number of security bugs and other issues you may have encountered in Confluence 2.5.5 or earlier.

As such **this release is a recommended upgrade for all customers**.

Confluence 2.5.6 can be downloaded from http://www.atlassian.com/software/confluence/ConfluenceDownloadCenter.jspa, and is a free upgrade for all customers who purchased their Confluence license or maintenance renewal after August 8th 2006.

Upgrading to Confluence 2.5.6

Upgrading Confluence should be fairly straightforward. You can find instructions here. *We strongly recommend that you backup your `confluence.home` directory and database before upgrading!*

Changes in 2.5.6

- For details about the security fixes, please see the security advisory.
- CONF-8944 resolves a Crowd integration issue for Confluence 2.5.6 and later.
- The Crowd integration fix will be ported to previous Confluence versions in the near future - please see CONF-9122.

Here's a complete list of the bug fixes in Confluence 2.5.6:

Errors were reported by the JIRA trusted connection.

- APP_UNKNOWN; Unknown Application: {0}; ["confluence:4557196"]

JIRA Issues (19 issues)

Type	Key	Summary	Priority	Status	Resolution
	CONF-9073	Changes to Crowd and Confluence integration instructions		Resolved	Fixed
	CONF-9067	Division by zero in SnipSnapImporter		Resolved	Fixed
	CONF-9060	Missing text in breadcrumbs when viewing changes since last login		Resolved	Fixed
	CONF-9000	OutOfMemoryError's during indexing		Resolved	Fixed
	CONF-8993	Reflected XSS Vulnerability in the Feed Builder		Resolved	Fixed
	CONF-8989	Index rebuild tasks involving an index optimization produce an OutOfMemoryError when there are many large textual attachments		Closed	Incomplete
	CONF-8980	XSS vulnerability at "Edit Space Permissions"		Resolved	Fixed
	CONF-8979	Vulnerability against DoS attack at permission setting		Resolved	Fixed

	CONF-8978	Vulnerability against DoS attack via labels		Resolved	Fixed
	CONF-8956	stored XSS vulnerability in app/themes/leftnavigation/configuretheme.action		Resolved	Fixed
	CONF-8952	XSS vulnerability in app/pages/listpages-alphaview.action		Resolved	Fixed
	CONF-8951	XSS vulnerability in app/spaces/editspace.action		Closed	Duplicate
	CONF-8950	XSS vulnerability in app/spaces/listattachmentforspace.action		Resolved	Fixed
	CONF-8944	Profile settings are lost when using Crowd with Confluence		Closed	Fixed
	CONF-8914	"Save" button mentioned as "Update" on change password screen		Resolved	Fixed
	CONF-8869	JavaScript error occurs when editing a page without having permission to restrict it		Resolved	Fixed
	CONF-8623	Improve and speed up plugin resource loading		Resolved	Fixed
	CONF-7414	Officially support mySQL 5.0		Resolved	Fixed
	CONF-3673	Page restrictions do not restrict Space Administrators		Resolved	Fixed

Upgrade Guide 2.5.6

Who should upgrade?

Confluence 2.5.6 is a recommended upgrade that resolves a number of security bugs and other issues you may have encountered in Confluence 2.5.5 or earlier.

As such **this release is a recommended upgrade for all customers**.

Confluence 2.5.6 can be downloaded from http://www.atlassian.com/software/confluence/ConfluenceDownloadCenter.jspa, and is a free upgrade for all customers who purchased their Confluence license or maintenance renewal after August 8th 2006.

If you are running a release older than 2.5 and you want to upgrade, you should upgrade directly to 2.5.6. Refer to the upgrade procedure below.

Upgrade Procedure

Upgrading Confluence should be fairly straightforward. *We strongly recommend that you backup your `confluence.home` directory and database before upgrading!*

If you are upgrading from Confluence 2.2.x or a later version, please see the Confluence Upgrade Instructions.

If you are upgrading directly from 2.1 or earlier, you should also read the 2.2 Release Notes for warnings about the 2.1 -> 2.2 upgrade.

Release Notes 2.5.7

Confluence 2.5.7 is a recommended upgrade that resolves two issues you may have encountered in Confluence 2.5.6:

- Indexing errors claiming 'too many open files', caused by duplicate libraries in the Confluence web application.
- Caching should be enabled by default for the LDAP configuration in `atlassian-user.xml`.

Confluence 2.5.7 can be downloaded from http://www.atlassian.com/software/confluence/ConfluenceDownloadCenter.jspa, and is a free upgrade for all customers who purchased their Confluence license or maintenance renewal after 30 August 2006.

If you don't wish to upgrade to 2.5.7, but do want a fix for issue CONF-9251 ("Too many open files" error during index operations), you can manually remove the duplicate libraries shipped in the Confluence web application. Please follow the instructions on the JIRA issue.

Upgrading to Confluence 2.5.7

Upgrading Confluence should be fairly straightforward. You can find instructions here. *We strongly recommend that you backup your `confluence.home` directory and database before upgrading!*

Changes in 2.5.7

Here's a complete list of the bug fixes in Confluence 2.5.7:

Errors were reported by the JIRA trusted connection.

- APP_UNKNOWN; Unknown Application: {0}; ["confluence:4557196"]

JIRA Issues (3 issues)					
Type	Key	Summary	Priority	Status	Resolution
	CONF-9303	Caching is not enabled by default for the hibernate repository		Resolved	Fixed
	CONF-9251	"Too many open files" error during index operations		Resolved	Fixed
	CONF-9249	Duplicate dependencies in WEB-INF/lib/		Resolved	Fixed

Upgrade Guide 2.5.7

Who should upgrade?

Confluence 2.5.7 is a recommended upgrade that resolves two issues you may have encountered in Confluence 2.5.6:

- Indexing errors claiming 'too many open files', caused by duplicate libraries in the Confluence web application.
- Caching should be enabled by default for the LDAP configuration in `atlassian-user.xml`.

Confluence 2.5.7 can be downloaded from http://www.atlassian.com/software/confluence/ConfluenceDownloadCenter.jspa, and is a free upgrade for all customers who purchased their Confluence license or maintenance renewal after 30 August 2006.

If you are running a release older than 2.5 and you want to upgrade, you should upgrade directly to 2.5.7. Refer to the upgrade procedure below.

Upgrade Procedure

Upgrading Confluence should be fairly straightforward. *We strongly recommend that you backup your `confluence.home` directory and database before upgrading!*

If you are upgrading from Confluence 2.2.x or a later version, please use the Confluence Upgrade Instructions.

If you are upgrading directly from 2.1 or earlier, you should also read the 2.2 Release Notes for warnings about the 2.1 -> 2.2 upgrade.

Release Notes 2.5.8

Confluence 2.5.8 is a **highly recommended** upgrade that resolves some issues in the user management framework and a memory leak in Confluence 2.5.x.

You can download Confluence 2.5.8 from http://www.atlassian.com/software/confluence/ConfluenceDownloadCenter.jspa. This upgrade is free for all customers with active Confluence software maintenance as at 3 October 2007.

This release has been made available after the Confluence 2.6 release date. This 2.5.8 version is for customers who want to fix the issues listed below, but do not want all the new functionality and theme changes in Confluence 2.6 yet.

Upgrading to Confluence 2.5.8

Upgrading Confluence should be fairly straightforward. You can find instructions here. *We strongly recommend that you backup your `confluence.home` directory and database before upgrading!*

Changes in 2.5.8

Here's a complete list of the bug fixes in Confluence 2.5.8:

Errors were reported by the JIRA trusted connection.

- APP_UNKNOWN; Unknown Application: {0}; ["confluence:4557196"]

JIRA Issues (5 issues)					
Type	Key	Summary	Priority	Status	Resolution
	CONF-9521	TinyMceServlet has a memory leak		Closed	Fixed
	CONF-9434	Enabling Caching for Hibernate Repository causes			Fixed

		net.sf.hibernate.LazyInitializationException error		Resolved	
	CONF-9307	Errors during user migration process		Resolved	Fixed
	CONF-9195	Confluence 2.5.6 ldap configuration failing on osuser2atluser.jsp migration		Resolved	Fixed
	CONF-9117	Confluence API supports adding user with null password, but users will null passwords produce NullPointerException when using the osuser to atlassian-user migration utility jsp		Resolved	Fixed

Upgrade Guide 2.5.8

Who should upgrade?

Confluence 2.5.8 is a **highly recommended** upgrade that resolves some issues in the user management framework and a memory leak in Confluence 2.5.x.

You can download Confluence 2.5.8 from http://www.atlassian.com/software/confluence/ConfluenceDownloadCenter.jspa. This upgrade is free for all customers with active Confluence software maintenance as at 3 October 2007.

If you are running a release older than 2.5 and you want to upgrade, you should upgrade directly to 2.5.8, referring to the upgrade procedure below.

Upgrade procedure

Upgrading Confluence should be fairly straightforward. *We strongly recommend that you back up your `confluence.home` directory and database before upgrading!*

If you are upgrading from Confluence 2.2.x or a later version, please use the Confluence Upgrade Instructions.

If you are upgrading directly from 2.1 or earlier, you should also read the 2.2 Release Notes for warnings about the 2.1 -> 2.2 upgrade.

Release Notes_1.0.1

Confluence 1.0.1

Over the course of the last month of supporting Confluence 1.0, there are a number of patches that we have been distributing to fix specific problems our users have encountered. Confluence 1.0.1 is a maintenance release into which all these patches have been integrated.

Who Should Upgrade?

All the issues that were resolved for this release are listed below. We have not snuck in any other changes: what you see here in the release notes is precisely what you will get. If you find something on the list that directly affects you, or that you feel justifies the effort of an upgrade, then do so. Otherwise, feel free to stick with 1.0.

Upgrade Procedure

If You Have Customised `osuser.xml`

If you have customised Confluence's user-management, for example to integrate it with LDAP or JIRA, you will have to integrate your changes to account for the caching OSUser providers we introduced in Confluence 1.0.1. See this document for more details: [Confluence Caching OSUser Provider]. Updated instructions for integrating with JIRA user management are here: Delegate user management to use JIRA logins

Otherwise

To avoid the possibility of data-loss, you should back up your ConfluenceHome directory and your database before upgrading, and perform a full backup from within the application.

Changes in 1.0.1

See also Issues Resolved for 1.0.1

Direct Attachment Links

In response to a loud demand from users, links to attachments using the [^attachment.jpg] syntax will download the attachment file directly, instead of linking to an anchor in the destination page.

Sybase ASE Support

Confluence 1.0.1 resolves the following problems that were causing Confluence not to work with Sybase ASE 12.5.1:

- ORDER-BY mappings failing on Sybase (CONF-1021)
- Can't add users under Sybase (CONF-1022)
- Weird datatype error under Sybase (CONF-1024)
- Sybase doesn't like complicated distinct selects (404 page fails) (CONF-1025)
- Backup import fails under Sybase (CONF-1063)

These bug-fixes may also improve Confluence's compatibility with other untested databases. They will not, however, have any effect on Confluence's operation against PostgreSQL, MySQL or HSQL.

Microsoft SQL Server Dialect in Setup Page

The Microsoft SQL Server database dialect was missing from the database setup page. It has now been added to the list. Confluence has not yet been tested on Microsoft SQL Server, and the usefulness of this option is not yet guaranteed. However, since the Sybase issues above are now resolved and SQL Server belongs to the same family as Sybase, it would be well worth a try.

JIRA User Provider Caching

Users who were linking their user management to JIRA's using the supplied provider were experiencing significant performance problems as a result. 1.0.1 introduces caching to the user provider, which should speed up these installations significantly.

Global Reports Visibility

Under Confluence 1.0, the global "undefined pages" and "orphaned pages" reports did not properly filter out pages that the user could not see. The user could not see the content of any page they did not have access to, but they could learn of the existence of (and names of) pages and spaces they were not permitted to see. This bug is fixed in 1.0.1

Locale-Independent Dates in Backup/Restore

In Confluence 1.0, dates were written into backup files using a localised representation of the month. As such, if you exported Confluence data from a server in one locale it might not import successfully into a server with a different Locale setting. Confluence 1.0.1 still recognises the 1.0 export format, but its own exports will write out dates in a locale-independent format.

As noted above, this means that data exported from Confluence 1.0.1 *can not* be imported successfully into Confluence 1.0.

Fix Browser Crash on Viewing Some Templates

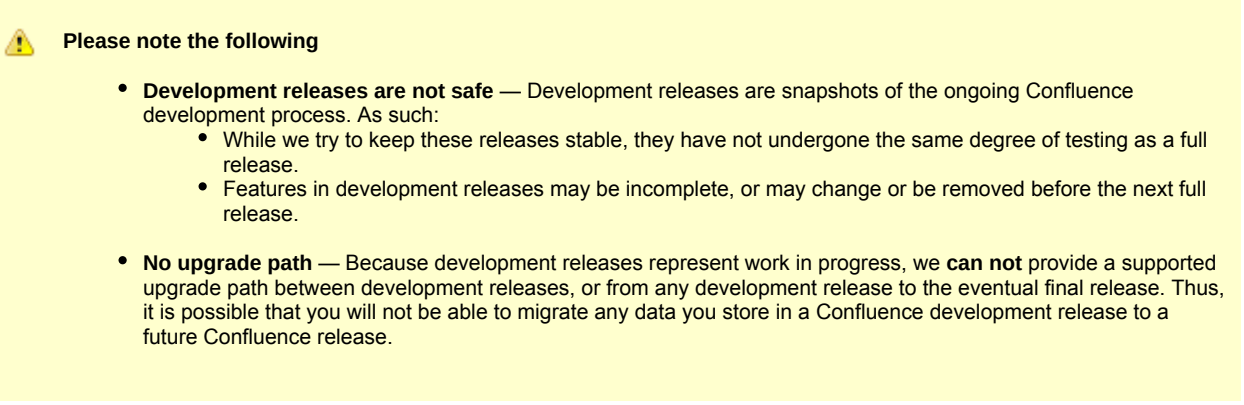
Previously, if you created a template containing no variables, then anyone attempting to preview or use that template would have their browser hang in an infinite Javascript loop. Confluence 1.0.1 fixes this bug.

Typo Fixed on User Group Editing Page

A single-character change from "privilage" to "privilege".

Development Releases

Development Releases are interim builds of Confluence that we make available so that interested customers can try out new features, especially those features that you may have been waiting for, and don't want to wait another month for the next official release.



Current Development Release Cycle

The links below show the improvements made every two weeks only, and will mainly be relevant to plugin developers.

- Release Notes 3.2-beta ("Beta") 1/MAR/2010
- Release Notes 3.2-m4 ("Milestone 4") 22/FEB/2010
- Release Notes 3.2-m3 ("Milestone 3") 5/FEB/2010

Previous Development Release Cycles

3.1 Development

- Release Notes 3.1-rc1("Release Candidate 1") 27/NOV/2009

- Release Notes 3.1-beta2 ("Beta 2") 19/NOV/2009
- Release Notes 3.1-beta1 ("Beta 1") 9/NOV/2009
- Release Notes 3.1-m7 ("Milestone 7") 5/NOV/2009
- Release Notes 3.1-m6 ("Milestone 6") 26/OCT/2009
- Release Notes 3.1-m5 ("Milestone 5") 8/OCT/2009
- Release Notes 3.1-m4 ("Milestone 4") 22/SEPT/2009
- Release Notes 3.1-m3 ("Milestone 3") 10/SEPT/2009
- Release Notes 3.1-m1 ("Milestone 1") 22/JULY/2009

Other Resources for 3.1:

Confluence 3.1 Deprecated Code Cleanup
Prototype REST API
Confluence 3.1 Newly Deprecated Code

3.0 Development

- Release Notes 3.0-beta2 ("Beta 2") 16/MAY/2009
- Release Notes 3.0-m9 ("Milestone 9") 28/APRIL/2009
- Release Notes 3.0-m8 ("Milestone 8") 20/APRIL/2009
- Release Notes 3.0-m7 ("Milestone 7") 31/MARCH/2009
- Release Notes 3.0-m6 ("Milestone 6") 16/MARCH/2009
- Release Notes 3.0-m5 ("Milestone 5") 24/FEBRUARY/2009
- Release Notes 3.0-m3 ("Milestone 3") 29/JANUARY/2009

2.10 Development

- Release Notes 2.10-rc1 ("Release Candidate 1") 23/NOV/2008
- Release Notes 2.10-m8 ("Milestone 8") 13/NOV/2008
- Release Notes 2.10-m7 ("Milestone 7") 12/NOV/2008
- Release Notes 2.10-m5 ("Milestone 5") 23/OCT/2008
- Release Notes 2.10-m4 ("Milestone 4") 2/OCT/2008
- Release Notes 2.10-m3 ("Milestone 3") 18/SEPT/2008
- Release Notes 2.10-m2 ("Milestone 2") 08/SEPT/2008
- Release Notes 2.10-m1 ("Milestone 1") 22/AUGUST/2008

2.9 Development

- Release Notes 2.9-rc1 ("Release Candidate 1") 31/JULY/2008
- Release Notes 2.9-m5 ("Milestone 5") 9/JULY/2008
- Release Notes 2.9-m3 ("Milestone 3") 11/JUNE/2008
- Release Notes 2.9-m2 ("Milestone 2") 27/MAY/2008

2.8 Development

- Release Notes 2.8-m9 ("Milestone 9") 26/MAR/2008
- Release Notes 2.8-m7 ("Milestone 7") 11/MAR/2008
- Release Notes 2.8-m6 ("Milestone 6") 05/MAR/2008
- Release Notes 2.8-m4 ("Milestone 4") 08/FEB/2008
- Release Notes 2.8-m3 ("Milestone 3") 30/JAN/2008
- Release Notes 2.8-m2 ("Milestone 2") 16/JAN/2008

2.7 Development

- Release Notes 2.7-rc1 ("Release Candidate 1") 05/DEC/2007
- Release Notes 2.7-m5 ("Milestone 5") 26/NOV/2007
- Release Notes 2.7-m4 ("Milestone 4") 06/NOV/2007
- Release Notes 2.7-m2 ("Milestone 2") 10/OCT/2007

2.6 Development

- Release Notes 2.6-dr1
- Release Notes 2.6-dr2

2.3 Development

- Release Notes 2.3-DR1
- Release Notes 2.3-DR2

1.5 - 2.0 Development

- 1.5-DR1 was not released
- Release Notes 1.5-DR2
- Release Notes 2.0-RC1
- Release Notes 2.0-RC2

1.4 Development

- Release Notes 1.4-DR1
- Release Notes 1.4-DR2
- Release Notes 1.4-DR3

- Release Notes 1.4-DR4
- 1.4-DR5 was not released
- Release Notes 1.4-DR6
- Release Notes 1.4-DR7
- 1.4-RC1 was not released
- Release Notes 1.4-RC2

1.3 Development

- Release Notes 1.3-DR1
- Release Notes 1.3-DR2
- Release Notes 1.3-DR3
- Release Notes 1.3-DR4
- Release Notes 1.3-final

Confluence 3.1 Newly Deprecated Code

This page should show all the code that was deprecated (not deleted!) during the Confluence 3.1 release cycle.

This is work in progress, and may not be a complete summary. As always, the truth is in the code. This page is more for explaining and discussing our decisions.

Development Release Warnings

Please note the following

- **Development releases are not safe** — Development releases are snapshots of the ongoing Confluence development process. As such:
  - While we try to keep these releases stable, they have not undergone the same degree of testing as a full release.
  - Features in development releases may be incomplete, or may change or be removed before the next full release.
- **No upgrade path** — Because development releases represent work in progress, we **can not** provide a supported upgrade path between development releases, or from any development release to the eventual final release. Thus, it is possible that you will not be able to migrate any data you store in a Confluence development release to a future Confluence release.

Milestone release advisory

This page is to be included into all our milestone release notes for the Confluence 3.2 release cycle

Do not use this release to upgrade your production systems.

For all production use and testing of Confluence, please use the latest official release.
This release is a public **development release** ('milestone') leading up to Confluence 3.2, which will probably ship in mid March 2010.
Development releases are a snapshot of our work in progress, allowing our customers and especially plugin developers to see what we're up to.

Who should upgrade?

Please note the following

- **Development releases are not safe** — Development releases are snapshots of the ongoing Confluence development process. As such:
  - While we try to keep these releases stable, they have not undergone the same degree of testing as a full release.
  - Features in development releases may be incomplete, or may change or be removed before the next full release.
- **No upgrade path** — Because development releases represent work in progress, we **can not** provide a supported upgrade path between development releases, or from any development release to the eventual final release. Thus, it is possible that you will not be able to migrate any data you store in a Confluence development release to a future Confluence release.

In supplying milestone releases, our aim is to provide plugin developers with an opportunity to see the latest changes in the code.

Each milestone release has passed all our automatic tests and has been used for one week on our official internal Confluence server. Most of the issues solved have been reviewed too, and all of our milestone releases even have been performance tested for a while.

However, since our milestones releases are timeboxed (i.e. they get released every two weeks, no matter how far we have come implementing features and bugfixes), there is always a chance that we have new known bugs that are scheduled to be fixed in the next milestone, or completely new bugs unknown even to us.

Additionally, we have not completed our performance testing and compatibility testing for databases and application servers. So, for example, a milestone release may behave well on a small installation but show severe problems when subjected to many users.

Upgrade Procedure

Follow the normal upgrade instructions to upgrade from Confluence 3.1.x to this release. *We strongly recommend that you backup your Confluence home directory and database before upgrading!*

Downloads

All development releases are available from the development releases page on the Atlassian website.

Beta Release Advisory



This page is to be included into all our Beta release notes for the Confluence 3.1 release cycle

Both "Milestone" and "Beta" versions of Confluence are development releases, which are preliminary releases leading up to the official release of a major Confluence version. They are a snapshot of our work in progress and provide an advance preview of new features to our customers and the general public. Confluence plugin developers can also use development releases to test and fix their plugins in advance of an official release.

The main distinction between a beta and a milestone release is that milestone releases typically acquire new features with each subsequent milestone version, whereas beta releases are predominantly feature-complete. Beta releases still undergo bug fixing and occasionally, existing features may be enhanced or added in subsequent beta versions.



Do not use in production

Beta releases should not be used in production environments as they are not officially supported.

For all production use and testing of Confluence, please use the latest official release.

Who should try this out?

With beta releases, the Confluence development team aims to provide plugin developers with an opportunity to see the latest changes in the code.

Furthermore, if you are a Confluence customer who is eager to see the new features and provide us with feedback on our upcoming major release, we encourage you to try out our beta releases.



Please note the following

- **Development releases are not safe** — Development releases are snapshots of the ongoing Confluence development process. As such:
  - While we try to keep these releases stable, they have not undergone the same degree of testing as a full release.
  - Features in development releases may be incomplete, or may change or be removed before the next full release.
- **No upgrade path** — Because development releases represent work in progress, we **can not** provide a supported upgrade path between development releases, or from any development release to the eventual final release. Thus, it is possible that you will not be able to migrate any data you store in a Confluence development release to a future Confluence release.

Each beta release has passed all our automated tests, has undergone some performance testing and has been used for one week on our official internal Confluence server. Furthermore, most of the solved issues have been reviewed.

Be aware that our beta releases are still undergoing final performance and compatibility testing for databases and application servers. Hence, we recommend that you use beta releases on installations with small (as opposed to full production-level) user bases.

Upgrade Procedure

If you wish to upgrade your existing Confluence installation with this version, ensure you have created a separate copy of your current

Confluence production installation *first* and using that copy, follow the normal upgrade instructions to upgrade it to this beta release. If you have also implemented customised site- or space-specific layouts, you will need to re-implement them after the upgrade. Otherwise, some of the new features in Confluence (or possibly existing features) may not function correctly.

Downloads

All development releases are available from Development Releases on the Atlassian website.

Release Notes 1.3-DR1

Confluence 1.3-DR1 is the first development release in the cycle leading up to Confluence 1.3. In it, we have rewritten the Space Summary/Space Administration pages to be more useable, adn implemented a much-requested feature: the ability to undelete pages.

Who should upgrade?

Development releases are snapshots of the ongoing Confluence development process. We make them available for customers who are willing to risk an unpolished release in order to have early access to new features.

If you want to be running the most stable and most reliable version of Confluence, you should stick with the official, numbered releases.

Upgrade Procedure

Upgrading Confluence should be pretty easy. If you are upgrading from Confluence 1.2 or higher, you can find instructions here. *We strongly recommend that you backup your `confluence.home` directory and database before upgrading!*

If you are upgrading from Confluence 1.1.2 or earlier, be sure to read the upgrade instructions in the Confluence 1.2 release notes.

New Features in Confluence 1.3-DR1

Space Summary

The Space Summary page has been the dumping-ground for every space-related operation and report since Confluence's early alpha releases. As such, it had become confusing to navigate and hard to find anything on.

For 1.3-DR1, we have rewritten the space summary to be a suite of tabbed pages, making that whole area of Confluence easier to navigate, more self-explanatory, and much better-looking.

The Trash Can

When a page or blog post is deleted, it is no longer removed completely, but is placed in a space-wide trash can.

The trash can be found in the Administration tab of the newly rearranged Space Summary pages. Space Administrators can choose to restore pages from the trash, or purge them so they are unrecoverable.

New Emoticons

Just because.

Emoticon	Name	Syntax
	add	(+)
	forbidden	(-)
	help	(?)
	idea on	(on)
	idea off	(off)
	star	(*)
	red star	(*r)
	green star	(*g)
	blue star	(*b)
	yellow star	(*y)

Bug Fixes

All the bug-fixes that were part of the Confluence 1.2.2 release have also been incorporated into 1.3-DR1. In addition, we've fixed one or two issues specifically for 1.3-DR1, take a look in JIRA for the full list.

Release Notes 1.3-DR2

Confluence 1.3-DR2 is the second development release in the cycle leading up to Confluence 1.3. For DR2, we have made space-level permissions a lot more fine-grained, and added a new plugin manager which should open the way to making it easier for people to write extensions to Confluence:

Who should upgrade?

Development releases are snapshots of the ongoing Confluence development process. We make them available for customers who are willing to risk an unpolished release in order to have early access to new features.

If you want to be running the most stable and most reliable version of Confluence, you should stick with the official, numbered releases.

Upgrade Procedure

Upgrading Confluence should be pretty easy. If you are upgrading from Confluence 1.2 or higher, you can find instructions here. *We strongly recommend that you backup your `confluence.home` directory and database before upgrading!*

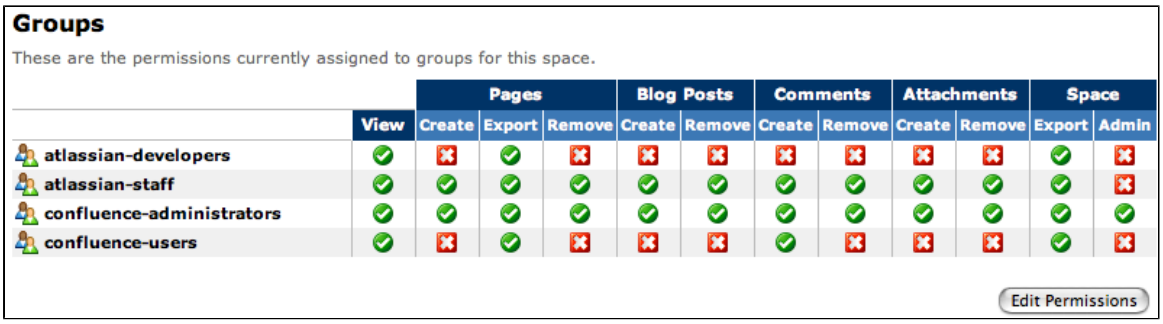
If you are upgrading from Confluence 1.1.2 or earlier, be sure to read the upgrade instructions in the Confluence 1.2 release notes.

Note: Because we have made significant changes to permissions in 1.3-DR2, you should check that your user permissions have been correctly migrated. The upgrade manager should ensure permissions remain consistent between versions, but as with anything related to security, it's best to double-check.

New Features in Confluence 1.3-DR2

New Space-Level Permissions

(CONF-1375, CONF-1764 and their linked issues)



We've been getting a lot of requests for more fine-grained control over who can do what in a space, so as a result, we've added a whole raft of new permissions at the space level:

Object	Permissions
Page	create/edit, remove, export
Blog Post	create/edit, remove
Comment	create, remove
Attachment	create, remove
Space	export, administer

Create space permission is still managed at the global level.

The permissions editing screens have been changed a little, too, taking into consideration the feedback we received from the last DR. All editing functions are now firmly on the edit screen, and we have introduced a user-picker to avoid having to find users in a massive drop-down list.

New Plugin Architecture

(CONF-1877)

Confluence now incorporates the plugin architecture that was written for JIRA 3.0 (which in turn was adapted from Confluence's macro management. Who said code reuse was dead?). Right now, there's not much you can do with it, but plugins pave the way to making Confluence a great deal more flexible and extensible.

See: Writing Confluence Plugins

New Macro Descriptor Format

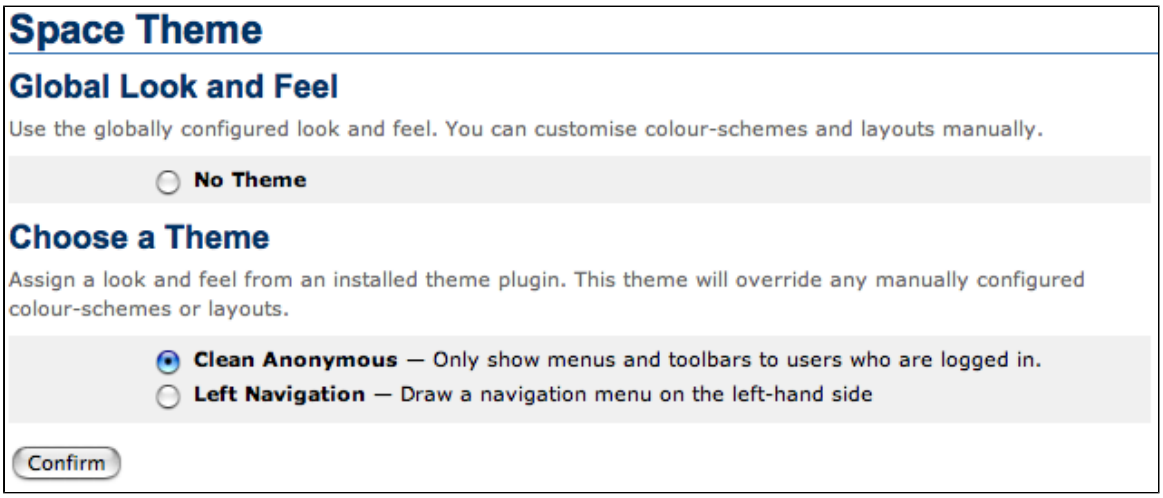
(CONF-1878)

Macro libraries are now implemented as plugins, which means that if you have written a macro library, you will need to convert your `macro-library.xml` files to `atlassian-plugin.xml` files.

Themes

(CONF-1856)

It is now possible to package a colour-scheme and a set of customised decorators into a portable Confluence plugin. Just drop the theme jar into your server's classpath, restart the server and space administrators will be able to select the theme from the space administration console.



DR-2 comes with two very simple themes (look under themes/ in your Confluence directory), but we will have more (and more useful) themes available by the time 1.3.0 is released. In addition, we're hoping that users will continue to share their Confluence modifications in the Confluence Extensions Space, and that themes will help you continue to share neat things you've done with Confluence.

Bugs fixed for 1.3-DR2

All the bug-fixes that were part of the Confluence 1.2.3 release have also been incorporated into 1.3-DR2. In addition, we've fixed a handful of issues specifically for 1.3-DR2, take a look in JIRA for the full list.

Release Notes 1.3-DR3

Confluence 1.3-DR3 is the third development release in the cycle leading up to Confluence 1.3 - and a momentous day in the history of Confluence.

We've finally made it to that special point in every applications life.

Confluence has evolved.

1.3-DR3 is best summed up by Jamie Zawinski's *Law of Software Envelopment*:

> *Every program attempts to expand until it can read mail. Those programs which cannot so expand are replaced by ones which can.*

That's right - Confluence now has *more content* than ever before. It reads, stores and indexes email. Oh, and as a nifty little bonus - it also indexes Word documents, PDF documents, RTF documents, Excel spreadsheets, PowerPoint files, text files, source files etc... *attached* to your mail!

Have we got your interest? How does all this magic happen you ask? Read on to pull back the curtain.

Who should upgrade?

Development releases are snapshots of the ongoing Confluence development process. We make them available for customers who are willing to risk an unpolished release in order to have early access to new features.

If you want to be running the most stable and most reliable version of Confluence, you should stick with the official, numbered releases.

Upgrade Procedure

Upgrading Confluence should be pretty easy. If you are upgrading from Confluence 1.2 or higher, you can find instructions here. *We strongly recommend that you backup your `confluence.home` directory and database before upgrading!*

If you are upgrading from Confluence 1.1.2 or earlier, be sure to read the upgrade instructions in the Confluence 1.2 release notes.

Note: Because we have made significant changes to permissions in 1.3-DR2, you should check that your user permissions have been correctly migrated. The upgrade manager should ensure permissions remain consistent between versions, but as with anything related to security, it's best to double-check.

New Features in Confluence 1.3-DR2

Mail Archiving

(CONF-1965)

Confluence is a collaboration tool. When communication happens through email instead of through Confluence, information will get lost in the depths of someone's mail folders, instead of shared with the team, archived, interlinked and indexed.

While we would love to see all collaboration within a group to happen through Confluence, it's often not possible to change the way people work. So instead of finding new ways to force people to use Confluence instead of email, why not route all the relevant email into Confluence?



It is now possible to archive email within a Confluence space. The support for mail is only just getting started in DR3, but we believe this is an incredibly useful direction for Confluence to go, and will be expanding and improving the mail integration in future releases.

Confluence spaces can retrieve mail periodically from a POP mailbox (this will delete all mail from that POP account, so don't try it on an account you want to keep mail on), or space administrators can import mail directly from a standard mbox-format mail file.

Once mail is imported into Confluence, it can be browsed chronologically from the Content pane of the Space Summary page, and can be searched through Confluence's search interface.

In the future, we will be adding new ways to link to and between email (bringing the email closer to the other content of the wiki), proper representation of email threads, more flexible search, and much, much more.

Read the Mail Archiving FAQ for more information

Want to play with it now?

We love to play with features right now, so here's how you can:

- Testing mail search is a snap. We're indexing our own mailing list on this very space! View archives or search mail.
- Testing mailing content into Confluence is also easy - well, relatively. Simply mail mailtest@atlassian.com and your email will show up here (it only polls the box once an hour, but it will show up eventually!). It will also be indexed.

Improved Indexing Performance

We now make much more efficient use of resources by batching updates to full-text search index. This should lead to improved performance for many tasks within Confluence, but will mean that it may take up to a minute for a change in the site to be reflected in the site's index. (CONF-2029

In addition, we have made a number of improvements to the indexing of large PDFs, including fixing some cases where a PDF might cause indexing to freeze indefinitely. (CONF-1953, CONF=1954)

Library Upgrades

Many of the librares Confluence is dependent on have been upgraded for this release, which should result in improved stability and performance.

Bugs fixed for 1.3-DR3

We've fixed a handful of issues specifically for 1.3-DR3, take a look in JIRA for the full list.

Release Notes 1.3-DR4

Confluence 1.3-DR4 is the fourth development release in the cycle leading up to Confluence 1.3. We're on the home stretch! The final, stable 1.3 isn't far away.

Confluence 1.3-DR4 includes a raft of improvements to mail archiving, a redesigned setup wizard, significant improvements to the way we back up and restore your system configuration, and a truck-load of bug fixes.

Who should upgrade?

Development releases are snapshots of the ongoing Confluence development process. We make them available for customers who are willing to risk an unpolished release in order to have early access to new features.

If you want to be running the most stable and most reliable version of Confluence, you should stick with the official, numbered releases.

Upgrade Procedure

Upgrading Confluence should be pretty easy. If you are upgrading from Confluence 1.2 or higher, you can find instructions here. *We strongly recommend that you backup your `confluence.home` directory and database before upgrading!*

If you are upgrading from Confluence 1.1.2 or earlier, be sure to read the upgrade instructions in the Confluence 1.2 release notes.

Note: You will need to rebuild the search index after you upgrade for certain features (including mail threading) to work properly.

New Features in Confluence 1.3-DR4

Setup Improvements

Confluence's setup wizard was badly in need of an overhaul, so that's exactly what we've done. Among the improvements, we have:

- Improved the ordering of steps. For example, if you are importing straight from a backup, you no longer need a redundant administrative user.
- Added a quick, two-step setup process to get Confluence running straight away
- Made the database setup much more robust, and much better at reporting errors
- Removed those steps that weren't really necessary to get Confluence running

Hopefully this will make it a lot easier to get Confluence up and running, or show off Confluence to your friends and family.

Mail Archive Improvements

(CONF-2050)

We've been working hard to build on the basic mail-archiving features included in 1.3-DR3. Included in the new version are:

- Improved indexing of mail senders, message-ids and subjects
- Improved integration of mail with search
- The ability to delete mail (if you have been assigned the delete mail permission)
- Direct access to mail archives from the dashboard
- Display of mail attachments in the "view mail" page
- The ability to link to a single message by its internal Confluence ID ([\$1234] will link to message 1234)
- The ability to navigate forward and back through mails chronologically
- Email addresses are displayed or masked in accordance with your global preferences

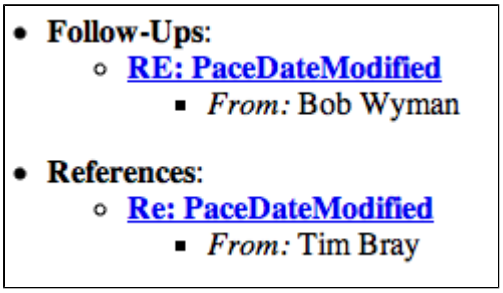
We still have a few things to do (date based views, a calendar), but the mail archive is now very useable. Especially when you take into account...

Mail Threading

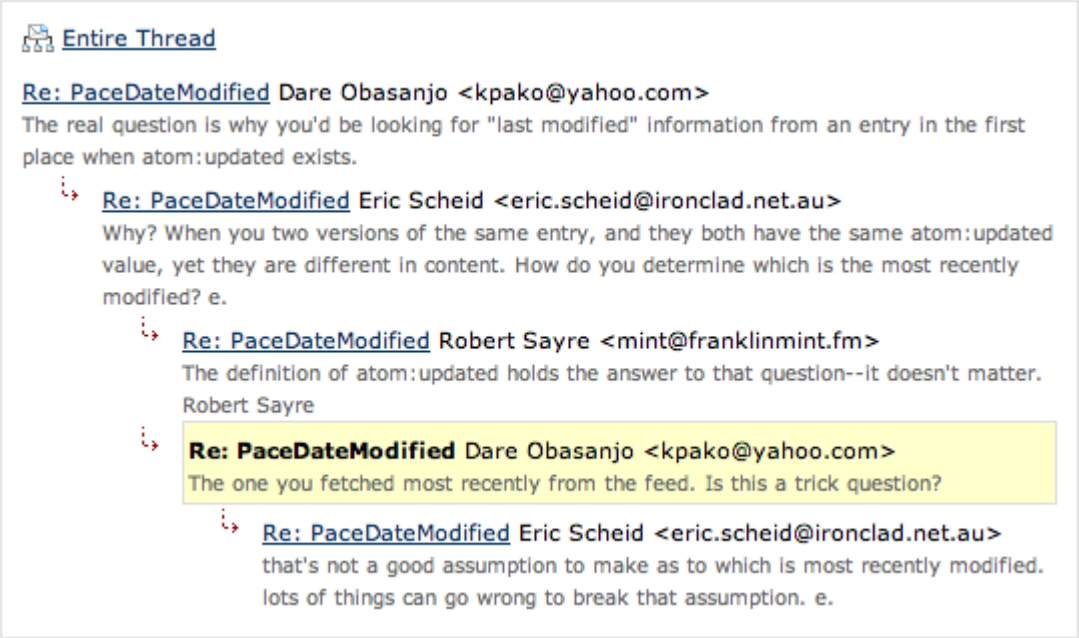
(CONF-2059)

Of course, it's important for your mail archive to be able to represent conversation threads. How you *present* threads, though, is just as important.

Here's how most mail-based web archives present your thread at the bottom of each message:



Not very helpful, is it. Here's how that same bottom-of-the-message view looks in Confluence:



At a glance, you can see exactly where the message sits in a conversation.

Backup Improvements

(CONF-1090 and CONF-2046)

Confluence 1.3-DR4 contains two significant improvements to backing up and restoring the system.

- You can now **exclude attached files from your backups**. Of course, this means you have to back up your attachment directory separately, but if you already have a good backup regime for your filesystem (and can thus restore the attached files separately), it means that your Confluence data backups will take up far less room.
- We now **include important parts system configuration in your backups**. this means that colour-schemes and plugin preferences are now saved, and fixes a large number of issues that users have filed due to the restore process not bringing the system back up in the same state as when they saved it.

Library Upgrades

We have continued the process of upgrading to the most recent stable versions of those libraries Confluence is built on, which should result in improved stability and performance.

Bugs fixed for 1.3-DR4

We've also fixed a lot of bugs, as we ramp up for the stable 1.3 release. Take a look in JIRA for the full list.

Release Notes 1.3-final

Confluence 1.3-final is the stable release of Confluence 1.3. Woohoo! We made it! The full release-notes for Confluence 1.3 are located here , this page documents only the changes made since the 1.3-DR4 development release.

1.3-final contains over 100 improvements over 1.3-DR4, mostly focused around fixing bugs, polishing the interface, and making Confluence ready for a stable release.

Who should upgrade?

Confluence 1.3 is the new stable release of Confluence. It contains a huge raft of enhancements and fixes made since 1.2.3. If you are running Confluence (and not using Oracle, see below), you should upgrade to Confluence 1.3. If you are not running Confluence, you should install 1.3 immediately, regardless of your database.

Current users of Confluence on **Oracle databases** users may wish to delay upgrading.

Upgrade Procedure

Upgrading Confluence should be pretty easy. If you are upgrading from Confluence 1.2 or higher, you can find instructions here. *We strongly recommend that you backup your `confluence.home` directory and database before upgrading!*

If you are upgrading from Confluence 1.1.2 or earlier, be sure to read the upgrade instructions in the Confluence 1.2 release notes.

Note: You will need to rebuild the search index after you upgrade for certain features (including mail threading) to work properly.

New Features in Confluence 1.3-final

New Demonstration Content

As part of the DR4 setup improvements, users were given the option to install an example space, but the demonstration content that was loaded was pretty uninspiring. For the final release, the demonstration space now contains a suite of demonstration pages, examples of what Confluence can do, and a few pictures of the Sydney Opera House. This should make it much easier to get Confluence up and running quickly.

Referrer Performance Improvement

We discovered (from monitoring the http://confluence.atlassian.com site) that our recording of HTTP referrers was causing some serious performance problems for public Confluence sites. Upgrading to Confluence 1.3-final should make Confluence a lot more responsive, especially under heavy load.

Improved Notation Guide

The notation guide has been reorganised to be more user-focused, making it easier to find the markup or macro you are looking for.

It is also now possible for macro plugins to insert themselves into the notation guide. Just write your macro description as a two-column HTML table row, put it in a vm file, and include the following in your `atlassian-plugin.xml` file:

```
1.<resource type="velocity" name="help" location="/path/to/your/helpfile.vm">
2.          <param name="help-section" value="tables"/>
3.</resource>
```

The help section can be one of: texteffects, headings, breaks, links, lists, images, tables, advanced, confluence, external or miscellaneous. If you don't include a help section, it will be put in the 'macros' section.

Improved Search Indexing

We've updated the way we index content within Confluence. A lot of searches that came up empty before will now find something. In addition, you can now set your primary language for indexing under General Configuration, so that the indexer can better optimise itself for non-English content.

Also

- By popular demand (it was our most highly voted-for bug), pages are now exported in alphabetical order, not in order of creation.
- The Remote API can be accessed anonymously, if you wish (anonymous remote access must be turned on in the general configuration).
- When previewing a page, you can continue to edit on that screen, instead of having to go back to the edit page.
- The thread view on a mail page has been improved, and the full-thread view is no longer a pop-up.
- You no longer need to be in the `confluence-admin` group to access the administration pages, you just need global administrative privileges
- You can link to anything in Confluence if you know its ID in the database (currently this is how you must link to email): [$1234]
- You can link to anything relative to the root of the Confluence site (useful for pointing to parts of the site that can't otherwise be linked): [///pages/editpage.action?pageId=1234]
- New macros: {jiraportlet}, {note} {tip}, {information} and {warning}
- A lot more...

Issues Resolved for 1.3-final

In all, over 130 issues were resolved between DR4 and 1.3-final. Unfortunately, merging all the versions together in JIRA means the list of precisely what went into those 130 has been lost, but if you sort this list by last-modification date, you'll get some idea... Issues Resolved for 1.3

Release Notes 1.4-DR1

Confluence 1.4-DR1 is the first development release in the cycle leading up to Confluence 1.4. On the surface, we have implemented one of Confluence's most requested features. Under the hood, there's a whole new event system to play with.

Who should upgrade?

Development releases are snapshots of the ongoing Confluence development process. We make them available for customers who are willing to risk an unpolished release in order to have early access to new features.

If you want to be running the most stable and most reliable version of Confluence, you should stick with the official, numbered releases.

Upgrade Procedure

Upgrading Confluence should be pretty easy. If you are upgrading from Confluence 1.3 or higher, you can find instructions here. *We strongly recommend that you backup your `confluence.home` directory and database before upgrading!*

If you are upgrading from Confluence 1.2.3 or earlier, be sure to read the upgrade instructions in the Confluence 1.3 release notes.

Downloads

All development releases are available from Development Releases on the Atlassian website.

New Features in Confluence 1.4-DR1

Consult JIRA for the full list of issues resolved for 1.4-DR1 (Note, this list doesn't include all the 1.3.1 fixes that were also merged into the development release)

Move Page Hierarchy

When you move a page between spaces, you now have the option to include all children of that page in the move. With 13 votes, this was one of Confluence's most requested features.

(Note, as of DR1, you can not both move a hierarchy of pages and rename the root page at the same time. This bug is filed as CONF-2452)

Attach Multiple Files

Another commonly requested feature, you can now attach multiple files to a page at once. By default, you can attach a maximum of five at a time, but this maximum is configurable.

GZip Content Encoding

Pages are served with GZip content encoding, which means less bandwidth consumption and faster page downloads.

Events and Listeners

We have implemented the beginnings of a pervasive system of events within Confluence.

By the final release of 1.4, all major actions within Confluence will trigger an event. These events can be listened for, and reacted to by a new plugin module type: "listener". This will allow Confluence to be further customised, as extensions can react to anything that happens within the system.

Creating an event listener plugin is insanely easy. Step one, implement this interface:

```
01.package com.atlassian.confluence.event;
02.
03.import com.atlassian.confluence.event.events.ConfluenceEvent;
04.
05./**
06.   * Defines a listener for Confluence events.
07.   */
08.public interface EventListener
09.{
10.          /**
11.             * Perform some action as a response to a Confluence event. The EventManager will
12.             * ensure that this is only called if the class of the event matches one of the
13.             * classes returned by getHandledEventClasses
14.             *
15.             * @param event some event triggered within Confluence
16.             */
17.          void handleEvent(ConfluenceEvent event);
18.
19.          /**
20.             * Determine which event classes this listener is interested in.
21.             *
22.             * The EventManager performs rudimentary filtering of events by their class. If
23.             * you want to receive only a subset of events passing through the system, return
24.             * an array of the Classes you wish to listen for from this method.
25.             *
26.             * For the sake of efficiency, only exact class matches are performed.
Sub/superclassing
27.             * is not taken into account.
28.             *
29.             * Returning an empty array will allow you to receive every event.
30.             *
31.             * @return An array of the event classes that this event listener is interested in,
32.             *         or an empty array if the listener should receive all events. Must not
33.             *         return null.
34.             */
35.          Class[] getHandledEventClasses();
36.}
```

Step two: Package it as a plugin module with a descriptor like this:

```
1.<listener name="My Listener" key="mylistener"
2.             class="com.example.listeners.MyListener">
3.          <description>Listens for stuff.</description>
4.</listener>
```

And that's about it. We'll be adding information about the events that are being produced in Confluence over the next couple of days (We would have had them documented today, but Nick, who wrote most of them, was sick).

Bugs Fixed

All bugs that were fixed in Confluence 1.3.1 are also fixed in Confluence 1.4-DR1.

Release Notes 1.4-DR2

Confluence 1.4-DR2 is the second development release in the cycle leading up to Confluence 1.4. Seeing as the DR2 development fortnight straddled Christmas it's not quite as feature-laden as some previous releases, but we thought we'd give you a few new toys to play with for the New Year.

Who should upgrade?

Development releases are snapshots of the ongoing Confluence development process. We make them available for customers who are willing to risk an unpolished release in order to have early access to new features.

If you want to be running the most stable and most reliable version of Confluence, you should stick with the official, numbered releases.

Upgrade Procedure

Upgrading Confluence should be pretty easy. If you are upgrading from Confluence 1.3 or higher, you can find instructions here. *We strongly recommend that you backup your `confluence.home` directory and database before upgrading!*

If you are upgrading from Confluence 1.2.3 or earlier, be sure to read the upgrade instructions in the Confluence 1.3 release notes.

Downloads

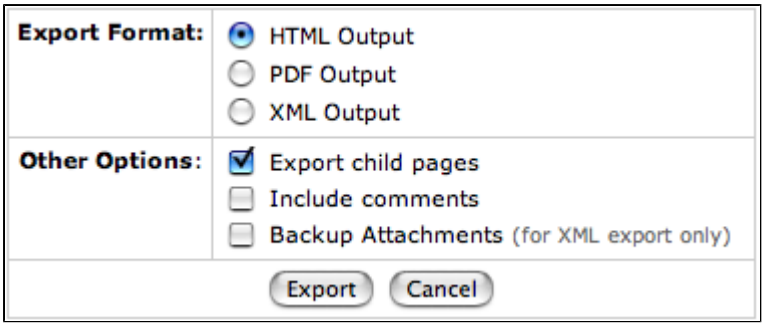
All development releases are available from Development Releases on the Atlassian website.

New Features in Confluence 1.4-DR2

Consult JIRA for the full list of issues resolved for 1.4-DR2

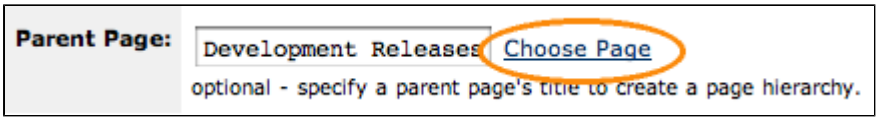
Export Page Hierarchy

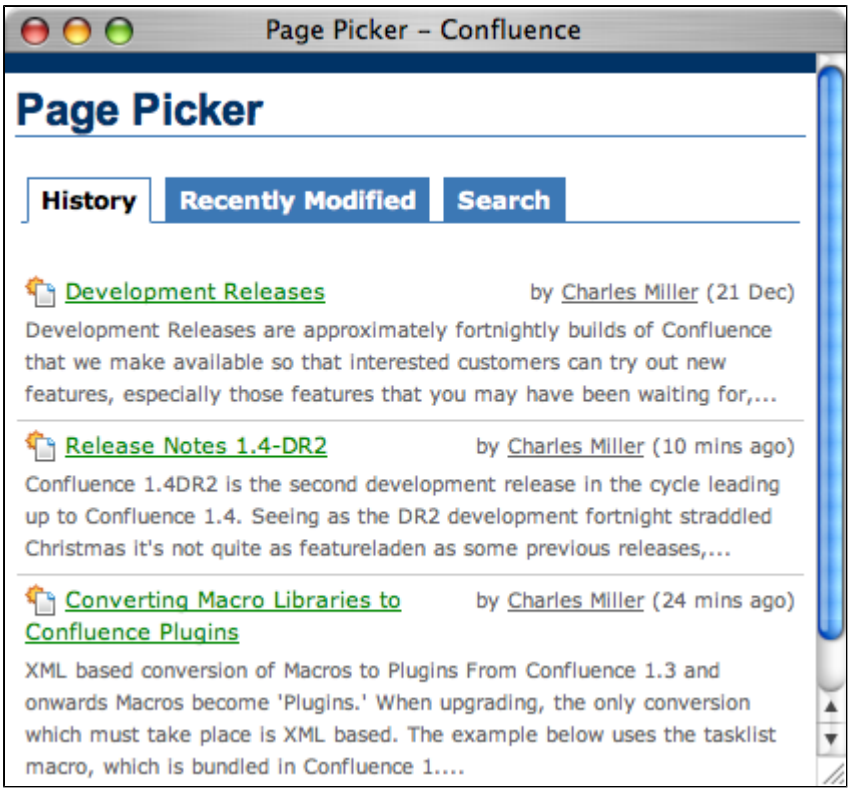
When you export a page to PDF or HTML, you have the option to include its children in the export.



Popup Page Picker

We now have a popup that can be used to select pages for inserting links, or choosing a page's parent. The picker lets you choose from your recently visited pages, the list of pages that link to the current page, or you can perform a search within the picker itself.





New Macros

As part of the process of making Confluence more flexible, we have implemented the various functions of the Confluence dashboard and space summary pages as macros. These macros are currently disabled by default, but if you're interested in using them within regular pages, you can enable them from your global plugin configuration.

This currently has no effect on the dashboard itself - but will become useful in the future when the dashboard becomes customiseable. For example, the {spaces-list} macro will reproduce the list of spaces that appears on the dashboard:

Spaces: My | Team | New (1) | All

- Application Links 1.x
  Documentation for AppLinks version 1.x
- Application Links 2.0
  Documentation for AppLinks 2.0
- Application Links 2.1
  Documentation for the latest version of AppLinks
- Atlassian Developer Network
  For the community of developers modifying and extending JIRA & Confluence.
- Atlassian Development
  Atlassian Developers. Because they're just too good to be kept hidden in an office in Sydney.
- Atlassian Documentation
  Information about and links to the Atlassian product documentation, including downloadable documentation
- Atlassian Events
- Atlassian IDE Connectors
  Documentation for the Atlassian Connectors for Eclipse and IntelliJ IDEA
- Atlassian Integration Guide
  The ways your Atlassian applications work together and how you can make it happen
- Atlassian KnowledgeBase
- Atlassian Partner Wiki
- Atlassian Presentations
- Atlassian Support
- Atlassian Training

- Atlassian t-shirt Competition
- Atlassian User Group
- Atlassian User Interface (AUI)
- Atlassian Webinars
- Bamboo 1.0
  Documentation for Bamboo 1.0
- Bamboo 1.1
  Documentation for Bamboo 1.1
- Bamboo 1.2
  Documentation for Bamboo 1.2
- Bamboo 2.0
  Documentation for Bamboo 2.0
- Bamboo 2.1
  Documentation for Bamboo 2.1
- Bamboo 2.2
  Documentation for Bamboo 2.2
- Bamboo 2.3
  Documentation for Bamboo 2.3
- Bamboo 2.4
  Documentation for Bamboo 2.4
- Bamboo 2.5
  Documentation for Bamboo 2.5
- Bamboo Extensions
- Bamboo Knowledge Base
  Troubleshooting and support tips for Bamboo
- Clover 2.0
  Documentation for Clover 2.0
- Clover 2.1
  Documentation for Clover 2.1
- Clover 2.3
  Documentation for Clover 2.3
- Clover 2.4
  Documentation for Clover 2.4
- Clover 2.5
  Documentation for Clover 2.5
- Clover 2.6
  Documentation for the latest version of Clover
- Clover 2.6 Archive
  Archived documentation for the latest version of Clover
- Clover Knowledge Base
  Troubleshooting and support tips for Clover
- Codegeist
  Atlassian's Codegeist! :)
- Confluence 1.4 User Guide
  User Guide for Confluence 1.4
- Confluence 2.0
  User Guide for Confluence version 2
- Confluence 2.5
  Complete documentation for Confluence versions 2.0 to 2.5.3.
- Confluence 2.5.6
  Complete documentation for Confluence versions 2.5.4 to 2.5.8.

Confluence 2.6
Complete documentation for Confluence version 2.6

Confluence 2.7
Complete documentation for Confluence version 2.7

Confluence 2.8
Complete documentation for Confluence version 2.8

Confluence 2.9
Complete documentation for Confluence version 2.9

Confluence Development

Confluence Docs 2.10
Complete documentation for Confluence version 2.10

Confluence Docs 3.0
Complete documentation for Confluence version 3.0

Confluence Docs 3.1
Complete documentation for Confluence version 3.1

Confluence Docs 3.2
Documentation for the latest version of the Confluence wiki: Installation Guide, User Guide, Admin Guide and other techn...

Confluence Evaluator Resources

Confluence Extensions

Confluence Homepage Image

Confluence Hosted Evaluator Resources

Confluence Knowledge Base
Troubleshooting and support tips for Confluence

Confluence SharePoint Connector 1.0
Documentation for version 1.0.x of the Confluence SharePoint Connector.

Confluence SharePoint Connector 1.1
Documentation for the latest version of the Confluence SharePoint Connector. This product integrates Confluence with Mic...

Confluence User Community
This space is for discussing ideas, new features and suggestions for Confluence.

Crowd 1.0
Documentation for Crowd version 1.0.x

Crowd 1.1
Documentation for Crowd version 1.1.x

Crowd 1.2
Documentation for Crowd version 1.2.x

Crowd 1.3
Documentation for Crowd version 1.3.x

Crowd 1.4
Documentation for Crowd version 1.4.x

Crowd 1.5
Documentation for Crowd version 1.5.x

Crowd 1.6
Documentation for Crowd version 1.6.x

Crowd 2.0
Documentation for the latest version of Crowd single signon and identity management

Crowd Development

Crowd Extension

Crowd Knowledge Base
Troubleshooting and support tips for Crowd

Crucible 1.1

Documentation for Crucible 1.1

Crucible 1.2
Documentation for Crucible 1.2

Crucible 1.5
Documentation for Crucible 1.5

Crucible 1.6
Documentation for Crucible 1.6

Crucible 2.0
Documentation for Crucible 2.0

Crucible 2.1
Documentation for Crucible 2.1

Crucible 2.2
Latest documentation for Crucible: do your code reviews, your way

Crucible Knowledge Base
Troubleshooting and support tips for Crucible

Demonstration Space
A space which demonstrates Confluence functionality.

Doc Sprint

Enterprise Hosting
Documentation and resources for Enterprise Hosted versions of Confluence and JIRA

Evaluator Resources

FishEye 1.3
Documentation for FishEye 1.3

FishEye 1.4
Documentation for FishEye 1.4

FishEye 1.5
Documentation for FishEye 1.5

FishEye 1.6
Documentation for FishEye 1.6

FishEye 2.0
Documentation for FishEye 2.0

FishEye 2.1
Documentation for FishEye 2.1

FishEye 2.2
Latest documentation for FishEye, your view into your source code repository

FishEye and Crucible Development
Tutorials and reference for enhancing and integrating FishEye and Crucible

Fisheye Knowledge Base
Troubleshooting and support tips for Fisheye

Gadget Development

Gadgets and Dashboards 1.0
Documentation for version 1.0.x of Atlassian Gadgets and Dashboards

Gadgets and Dashboards 2.0
Documentation for the latest version of Atlassian Gadgets and Dashboards

GreenHopper 3.8
Documentation archive for GreenHopper 3.8

GreenHopper 4.0
Documentation for GreenHopper 4.0

GreenHopper 4.1
Documentation for GreenHopper 4.1

GreenHopper 4.2
Archive of the documentation for GreenHopper 4.2.

GreenHopper 4.3

Documentation for GreenHopper 4.3

- GreenHopper 4.3 documentation archive
  Documentation for GreenHopper 4.3
- Greenhopper Knowledge Base
- JIRA 3.x Developer Documentation Archive
  JIRA 3.x developer documentation archive
- JIRA 4.0
  Documentation for JIRA 4.0
- JIRA 4.0 Archive
  Archived documentation for JIRA 4.0
- JIRA Community Space
- JIRA Download Evaluator Resources
- JIRA Extensions
- JIRA Hosted Evaluator Resources
- JIRA Knowledge Base
  Troubleshooting and support tips for Jira
- JIRA Studio
- JIRA Studio Evaluator Resources
- Partner Program
- Plugin Framework
  Documentation for all versions of the Atlassian Plugin Framework.
- Plugin Framework 2.0
  Documentation for version 2.0 of the Atlassian Plugin Framework.
- Plugin Framework 2.1
  Documentation for version 2.1 of the Atlassian Plugin Framework.
- Plugin Framework 2.2
  Documentation for version 2.2 of the Atlassian Plugin Framework.
- Plugin Framework 2.3
  Documentation for version 2.3 of the Atlassian Plugin Framework.
- Shared Access Layer
  Documentation for the Atlassian Shared Access Layer (SAL)
- Summit 2010 Planning
- Summit 2010Speaker Guidelines
- Test space
- test space copy
  PLEASE DELETE
- The Confluence Test Space
  This is a space dedicated to testing and playing around with Confluence features. Everyone is invited to experiment.
- User Group Sponsorship

Known Bugs

There is a known bug that occurs whenever you enable or disable a plugin in Confluence 1.4-DR2. The error will appear to be 'null', but the following will appear in the logfiles:

```
[ERROR]  Tue Jan 04 18:01:20 CST 2005  [com.atlassian.core.util.DateUtils]
java.util.MissingResourceException: Can't find resource for bundle
java.util.PropertyResourceBundle, key core.dateutils.minutes
```

This is issue CONF-2513 and will be fixed in the next DR.

Release Notes 1.4-DR3

Confluence 1.4-DR3 is the third development release in the cycle leading up to Confluence 1.4. It is the first step in a process of making the Confluence User Interface simpler, and easier to navigate.

Who should upgrade?

Development releases are snapshots of the ongoing Confluence development process. We make them available for customers who are willing to risk an unpolished release in order to have early access to new features.

If you want to be running the most stable and most reliable version of Confluence, you should stick with the official, numbered releases.

Upgrade Procedure

Upgrading Confluence should be pretty easy. If you are upgrading from Confluence 1.3 or higher, you can find instructions here. *We strongly recommend that you backup your `confluence.home` directory and database before upgrading!*

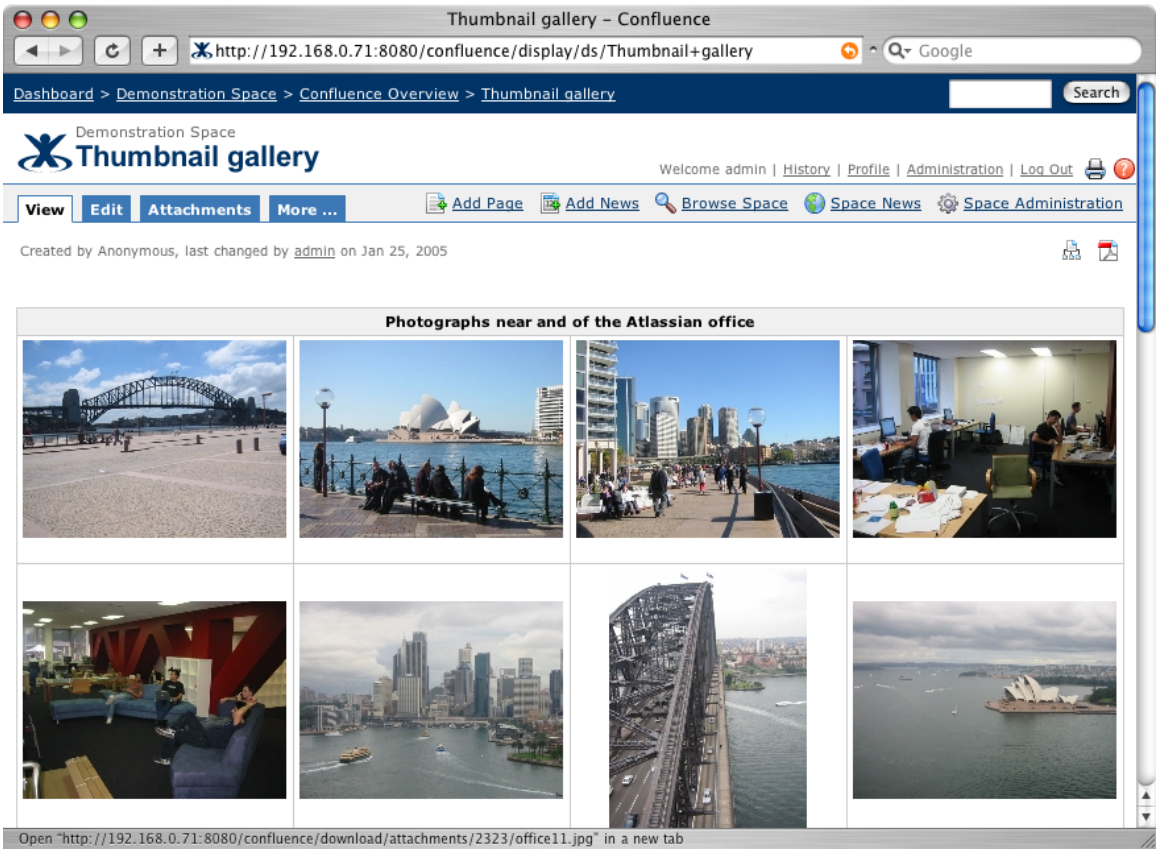
If you are upgrading from Confluence 1.2.3 or earlier, be sure to read the upgrade instructions in the Confluence 1.3 release notes.

Downloads

All development releases are available from Development Releases on the Atlassian website.

New Features in Confluence 1.4-DR3

User Interface Enhancements

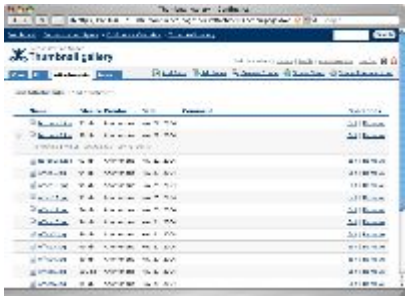


Some of the most frequent feedback we get about Confluence is that while users like the wealth of features on offer, the interface is often overwhelmed by too much information and too many options, making Confluence difficult to navigate, and difficult to introduce to non-technical users. Confluence 1.4-DR3 is the first step in an attempt to present a cleaner, more useable design for Confluence, while still providing a powerful interface for experienced users.

> 1.4-DR3 is a developer release. As such, some of the changes to the interface are incomplete, and some may not make it into the final version. Please offer your feedback on 1.4 UI Discussion.

While the obvious way to get a look at the new interface is just to explore this site, here are a few screenshots that might help:

Confluence 1.4 Interface Changes



New-look recently updated list for spaces
New attachment management interface
More Functional Edit Page Interface
Simpler Dashboard



Cleaner (and wider) page view

I'm also compiling a page to help people who are used to Confluence Classic, and might be a little lost in the new interface: 1.4 Interface - Where Did Everything Go?

Known Bugs

A page has been created in the Discussion space to gather comments, suggestions, compliments and complaints about the new interface: 1.4 UI Discussion

1.4 Interface - Where Did Everything Go?

If you're used to the Confluence Classic interface, the new 1.4 interface might be a little hard to navigate at first. We think that the new design is a vast improvement, but it's inevitable that people who have used the application since its release will be wondering where some of the buttons they were used to have gone...

Where Did They Go...

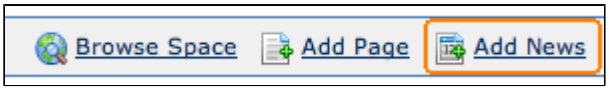
- Blogs?
- Rename Page?
- Move Page?
- Create Child Page?
- Page Locks?
- Page History? Short URL? Watch Page? Remove Page?

Like it? Don't? Join the discussion.

Blogs?

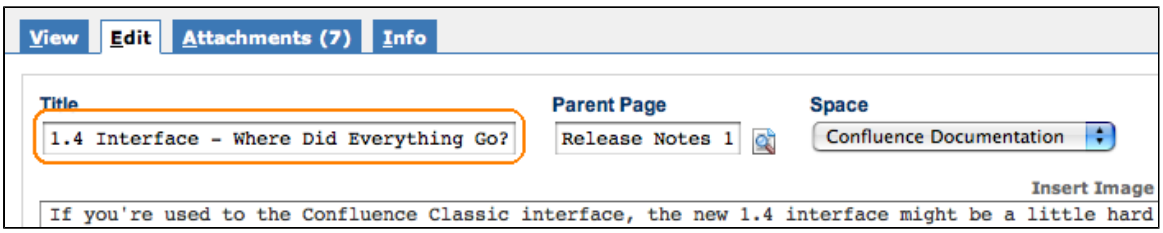
By incredibly popular demand, blogs have been renamed to "News" across the site. While blogs are all the rage at the moment (or, as my brother would say, very *zeitgeisty*), it's still far easier to explain "News" to someone familiar with blogs, than it is to explain blogs to someone who is familiar with news.





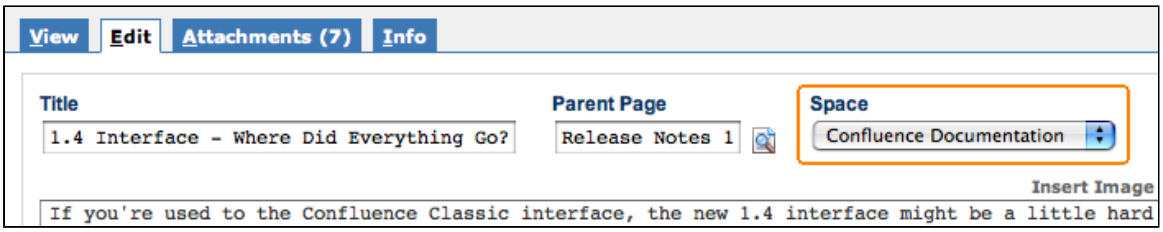
Rename Page?

Rename page is no longer a separate function. Just edit the page and change the page title, and Confluence will rename all the links to the page for you.



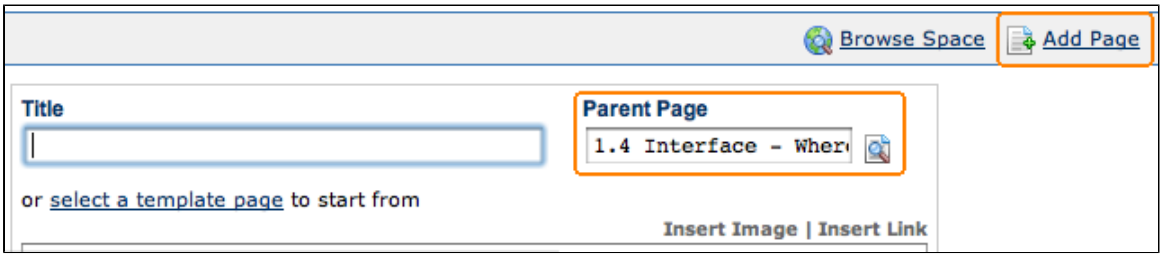
Move Page?

Move page is no longer a separate function. Just edit the page and change its space, and Confluence will move the page to the new space for you, renaming any links to the page. If the page has children, you will be given the option to move all the children as well.



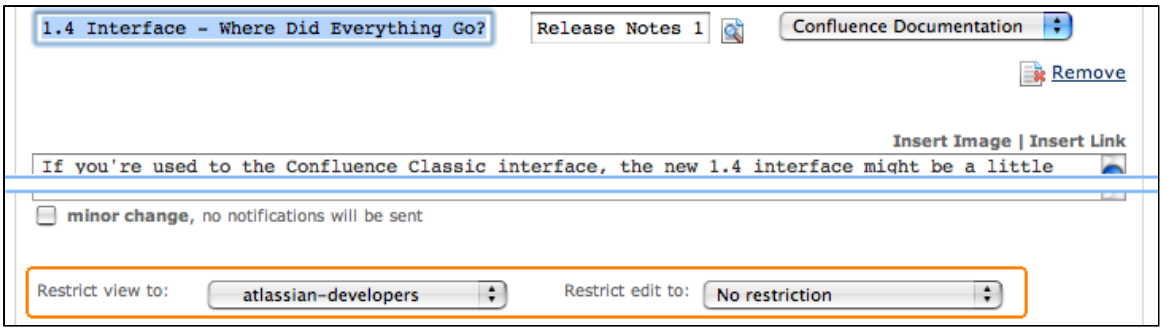
Create Child Page?

Whenever you select the "Add Page" link from the top of the page, your current page will be automatically inserted into the Parent Page box. We've found that this is the way most people use Confluence anyway.



Page Locks?

Page locks have been renamed "Permissions" and moved to the bottom of the edit page. We found that most people didn't understand (or couldn't find) locks, and thus didn't understand that Confluence has page-level edit security.



Page History? Short URL? Incoming Links? Hot Referrers?

All this information has been moved under the "Info" tab when you view a page.



Release Notes 1.4-DR4

Confluence 1.4-DR4

Confluence 1.4-DR4 is the fourth development release in the cycle leading up to Confluence 1.4. In 1.4-DR4 the user interface continues to improve, and our new wiki->HTML renderer lands, fixing a large number of bugs (but probably introducing a few more)

1.4-DR4a
One of the bugs that was introduced in 1.4-DR4 caused the edit option to disappear completely from the Confluence interface for anyone who was not a global administrator. Because that issue rendered Confluence mostly unuseable, 1.4-DR4a has been released including a fix to that problem. 1.4-DR4a also includes the patch for Confluence Security Advisory 2005-02-09.

We would like to remind everyonee who downloads the developer releases that they are *not* fully stable.

Who should upgrade?

Development releases are snapshots of the ongoing Confluence development process. We make them available for customers who are willing to risk an unpolished release in order to have early access to new features.

If you want to be running the most stable and most reliable version of Confluence, you should stick with the official, numbered releases.

Upgrade Procedure

Upgrading Confluence should be pretty easy. If you are upgrading from Confluence 1.3 or higher, you can find instructions here. *We strongly recommend that you backup your `confluence.home` directory and database before upgrading!*

If you are upgrading from Confluence 1.2.3 or earlier, be sure to read the upgrade instructions in the Confluence 1.3 release notes.

Downloads

All development releases are available from Development Releases on the Atlassian website.

New Features in Confluence 1.4-DR4

User Interface Enhancements

The user interface of Confluence continues to improve. A big thankyou to all the users who contributed feedback on DR3. While the changes

between DR3 and DR4 aren't as radical as the last release, we've been able to polish up a lot of the edges, and make Confluence even easier and more pleasant to use.

We're still gathering feedback on the UI changes, add your input here: 1.4 UI Discussion

Mail Archives

- You can now bulk-delete all mail within a space (mail deleted this way does not go to the trash, and **can not** be recovered later!)
- Pagination on the mail archives is a bit more sane: it doesn't try to tell you about every single page of mail in the archives.
- The mbox importer detects if you're trying to import something that isn't an mbox file, making you much less likely to break anything if you upload the wrong file.

Nofollow Support

As part of the effort to combat spam on wikis and blogs (Confluence being both), Google came up with some markup that will cause search-engines not to follow links. By removing the main benefit of wiki-spamming (increased search-engine PageRank), it's hoped that the noxious practice will stop being cost-effective, and eventually die out.

By default, all URLs inserted in a page (or in comments) will be given the `nofollow` tag. Inter-page links or shortcut links (i.e. CONF-2622) will not be tagged. If your wiki doesn't support public editing or commenting, or you just disapprove of nofollow on principle, the site administrator can turn the feature off in General Configuration.

V2Renderer

Over the last few releases of Confluence, it has become apparent that the engine we were using to convert wiki text to HTML was starting to cause us problems. There were a growing number of bugs that we *simply couldn't fix*, because the nature of the engine meant that fixing one bug would cause ten worse bugs to spring up in its place. After a few false starts, we can now introduce 'v2Renderer', which fixes quite a few of these problems already, and promises to make fixing the rest of them much easier:

Error formatting macro: jiraissues: java.lang.RuntimeException: A value with ID '11040' does not exist for the field 'fixVersion'.

Because this is the first public release of v2Renderer, it's likely that there are some situations in which the markup of pages may change. We've been pretty careful, but users who have been making heavy use of undocumented (or accidental) features of the old markup may encounter situations in which the page displays differently to before. Please, *please* report these problems to us, so that we can fix the renderer for the final release, or perhaps provide you with a script to painlessly convert your pages to a markup that works.

> **Renderer Bugs**
> If you file a bug about the v2Renderer, **PLEASE** make sure you put it in the Wiki/XHTML component: this makes finding all the relevant bugs so much easier for me. – Charles Miller

Macros

There is a new macro API that we will be rolling out for 1.4. However, there's a compatibility layer in place to make sure that most 1.3 macros will continue to function properly. Most macros written for Confluence 1.3 will continue to work in 1.4-DR4. Most of our macros are running without any changes at all, and many of the ones we did migrate were changed over in order to use the capabilities of the new engine to fix bugs.

We suggest that macro authors who are worried about compatibility test their macros with 1.4-DR4 and contact us as soon as possible to discuss how we can improve the compatibility layer before the final release, or help you migrate to the new API.

Known Bugs

The renaming of links when a page is moved or renamed is currently unreliable. This is a side-effect of the change in renderers, and will be fixed before the final version of 1.4 is released.

Release Notes 1.4-DR6

Confluence 1.4-DR6

Confluence 1.4-DR6 is the next development release in the cycle leading up to Confluence 1.4. In 1.4-DR6 we:

- introduce page-level view permissions,
- make Confluence radically more themeable,
- further improved the overall user interface and usability,
- add several interesting new plugin types,
- add file attachments to blog posts and user profile pages,
- let you customize Confluence logos at a space and global level, and
- help you keep track of all the email watches you might have set up in Confluence.

Phew! Got all that?

Who should upgrade?

Development releases are snapshots of the ongoing Confluence development process. We make them available for customers who are willing to risk an unpolished release in order to have early access to new features.

If you want to be running the most stable and most reliable version of Confluence, you should stick with the official, numbered releases.



Note for Oracle Users
Confluence 1.4-DR6 is currently not recommended for Oracle users, as it has not passed our test-suite against the Oracle database. This will (obviously) be fixed before we release Confluence 1.4.

Upgrade Procedure

Upgrading Confluence should be pretty easy. If you are upgrading from Confluence 1.3 or higher, you can find instructions here. *We strongly recommend that you backup your `confluence.home` directory and database before upgrading!*

If you are upgrading from Confluence 1.2.3 or earlier, be sure to read the upgrade instructions in the Confluence 1.3 release notes.



If you upgrade to Confluence 1.4-DR6 you will need to rebuild your search indexes before search will function correctly. Look for "Rebuild Search Index" on the global Administration page.

Downloads

All development releases are available from Development Releases on the Atlassian website.

New Features in Confluence 1.4-DR6

See also: Issues Resolved for 1.4-DR6

Page Level Permissioning

A highly-requested feature, and the reason this DR took so long to get out, you can now hide pages from other users. The option for hiding pages can be found at the bottom of the page's "Edit" screen.

- You can hide a page so it is only visible to yourself, or only visible to members of groups you belong to.
- Once a page is hidden, all its child pages (and their children, and so on) are also hidden.
- A new space-level permission has been introduced to determine which users are permitted to hide pages.

Please note that it may be possible for users to learn the names of hidden pages, even if they can't see their content. So don't make a page called "We Hate Bob Smith" if Bob Smith has access to your wiki. Just make a page called "We hate..." and put Bob's name in the body. 

Improved Themes

While the overwhelming reaction to the User Interface changes we have made in the 1.4-DR series has been positive, it has sparked continuing discussions on the user mailing list about various aspects of the presentation of a Confluence site. From these discussions we have determined:

- Everyone wants something slightly different.
- Some people want something completely different.

In response, we've moved to vastly improve the level of themeability that Confluence offers. Almost all of the user interface elements have been pushed into customisable decorators, so if you don't like the tabs, you will be able to install a theme that removes them entirely! We're hoping to bring out some example themes with the 1.4 release that do just this, but for now you can read the documentation for the new decorators in the Theme Module documentation.

Plugin Improvements

To make Confluence more flexible, we're introducing more ways for advanced users to extend the system:

Uploadable Plugins

Confluence administrators can now upload new plugins through the administrative interface, and have them installed without restarting the server.

Search Extractor Plugins

Extractor plugins allow you to add custom information to Confluence's full-text index when Confluence content is saved or updated. The most obvious application for this plugin type is to allow you to index attachment formats that Confluence does not yet support, but more creative plugin authors might make use of this as a sneaky way to store and retrieve metadata about any content in the Confluence system. Documentation for Extractor plugins is coming soon.

XWork Plugins

XWork plugins allow you to add new XWork (Webwork 2) actions to Confluence. Advanced programmers can add entirely new behaviours to the application: take a look at the examples in XWork-WebWork Module, in which we've implemented a rudimentary "Google Suggest"-like live search entry box for Confluence:

Servlet Plugins

Servlet plugins allow you to deploy servlets into Confluence dynamically – perfect for integrating Confluence with some legacy application that is only available as a servlet. See Servlet Module for an example.

Attachments for Blog Posts and User Profiles

You can now upload attachments to blog posts and user profiles, both frequently requested features. Users have the option to nominate one picture attached to their profile as their "profile picture", but we haven't quite decided what to do with this information yet.

There is a new global permission to determine if users can attach files to their profile. This permission is not assigned by default, so you will need to explicitly give it to users if you want to enable this feature.

Customizing Logos at a Space and Global Level

Space and global administrators may now easily configure the image displayed as the Confluence logo. This can be done for each space, or at a global level, without having to actually edit a single file.

Keeping track of your Email Watches

A Notifications link on your profile displays a list of all currently configured email watches you might have set up on various spaces or pages. You can now delete them in one, central, place or simply visit them.

Bug fixes

All bugs fixes made between Confluence 1.3.2 and Confluence 1.3.5 have been merged into the developer release branch, fixing many annoying bugs including the crash on selecting "Next" in search results. See also: Issues Resolved for 1.3.4, Issues Resolved for 1.3.5.

Issues Resolved for 1.4-DR6

Issues resolved for 1.4-DR5

Error formatting macro: jiraissues: java.lang.RuntimeException: A value with ID '11059' does not exist for the field 'fixVersion'.

Issues resolved for 1.4-DR6

Error formatting macro: jiraissues: java.lang.RuntimeException: A value with ID '11122' does not exist for the field 'fixVersion'.

Release Notes 1.4-DR7

Confluence 1.4-DR7

Confluence 1.4-DR7 is the first release-candidate build of Confluence 1.4. In 1.4-DR7 we have resolved over 120 issues: mostly related to fixing bugs and polishing up the application for a stable release. We anticipate that DR7 will be the last major build before the stable 1.4 release.

DR7 is not a public release. It has been built and deployed onto http://confluence.atlassian.com as part of our internal quality-control process, but we have decided that the developer time required to make this a public beta release would be better directed towards continuing the push towards a final, stable 1.4 build.

Who should upgrade?

Since Confluence 1.4-DR7 is not publicly available, nobody should upgrade. The upgrade note below is just for completeness, so when we come to gather all these release notes together into a single note for 1.4, we don't miss the warning.

Upgrade Procedure

Upgrading Confluence should be pretty easy. If you are upgrading from Confluence 1.3 or higher, you can find instructions here. *We strongly recommend that you backup your `confluence.home` directory and database before upgrading!*

If you are upgrading from Confluence 1.2.3 or earlier, be sure to read the upgrade instructions in the Confluence 1.3 release notes.

> When you first start Confluence 1.4-DR7 after upgrading, the server may take a few minutes to become fully functional. This is because Confluence is rearranging information in the database so that it can be looked up more efficiently. This will only be done the first time Confluence is started after upgrading.

New Features in Confluence 1.4-DR7

The changes made between DR6 and DR7 are mostly correctness and polish issues - major feature development on Confluence 1.4 has been completed. A few notable changes however are:

- The space content tree is now updated dynamically, leading to much faster page-load times and much less load on the server (for

large spaces, this page was turning into a reliable way to DOS Confluence). However, you must have a modern web browser with Javascript enabled for the tree to work. (A non-Javascript fallback will be included in 1.4-final - CONF-3098)
- You can now customise the site's name alongside the logo
- The display and navigation of page history diffs has been improved significantly
- Webdav configuration has been reinstated to the same functionality as 1.3.5
- An "insert image" popup has been added to the page editing screen, that can select images from the page's attachments. There is also an Attachments tab in the "insert link" popup. Allowing file upload from these popups has not yet been implemented (CONF-3099)
- Many improvements and fixes to email notifications
- Marking your change as a "minor edit" will prevent email notifications being sent
- The new Component plugin module adds new beans to the Spring context
- You can now uninstall plugins via the web interface
- The default theme has various keyboard shortcuts (see the tabs on this page)
- The Include Page Macro can now include pages from other spaces
- For developers, macros can now decide whether or not they have a body
- Report showing all attachments to pages in a space (check it out!)
- Generally improved notifications including a much more usable "File Attached" notification
- "View wiki source" link, for all those source junkies...

And much much more.

Release Notes 1.4-RC2

Confluence 1.4 Release Candidate 2



1.4-RC2 is now available for download here.

Confluence 1.4-RC2 is (barring some absolute disaster) the final public developer build for Confluence 1.4. It incorporates all the changes that were deployed in the private DR7 build, as well as around 60 new bug fixes and improvements.

Confluence 1.4-RC2 is being released to give Confluence customers the opportunity to test the upgrade path for their existing Confluence installation. By testing your migration with this pre-release, you can ensure that any problems you may experience will be fixed before the final 1.4 release, rather than having to rely on ad-hoc patches or wait for 1.4.1.

Who should upgrade?

While 1.4-RC2 is still pre-release software, and we do not recommend upgrading production Confluence sites until the release of 1.4-final, we would recommend any existing Confluence customer install this release on a test server, and try to import their existing Confluence data into it.

Upgrade Procedure

Upgrading Confluence should be pretty easy. If you are upgrading from Confluence 1.3 or higher, you can find instructions here. *We strongly recommend that you backup your `confluence.home` directory and database before upgrading!*

If you are upgrading from Confluence 1.2.3 or earlier, be sure to read the upgrade instructions in the Confluence 1.3 release notes.



When you first start Confluence 1.4-RC2 after upgrading, the server may take a few minutes to become fully functional. This is because Confluence is rearranging information in the database so that it can be looked up more efficiently. This will only be done the first time Confluence is started after upgrading.

New Features in Confluence 1.4-RC2

The changes made between DR7 and RC2 are mostly correctness and polish issues - major feature development on Confluence 1.4 has been completed. A few notable changes however are:

- You can upload attachments while editing a page – from the "Insert Link" popup
- The dynamic content-tree degrades gracefully on unsupported browsers
- Many bugs related to PDF export (and images included in PDF exports) have been fixed
- Ü characters in page titles no longer confuse the database
- Text properties set via the contentPropertyManager can now be longer than 255 characters, fixing related bugs in the {tasklist} and {excerpt} macros
- The {attachments} macro can display the list of attachments inline within a page
- New XML-RPC and SOAP plugin types allow you to deploy web services dynamically into Confluence
- A potential connection leak that has been causing problems for Oracle users has been fixed.

And much much more.

Release Notes 1.5-DR1

Confluence 1.5-DR1 is the first development release leading up to Confluence 1.5. Developer releases are a snapshot of our work in progress, allowing our customers to see what we're up to, and provide feedback

Who should upgrade?

Development releases are snapshots of the ongoing Confluence development process. We make them available for customers who are willing to risk an unpolished release in order to have early access to new features. Developer releases are not suitable for running on production systems.

If you want to be running the most stable and most reliable version of Confluence, you should stick with the official, numbered releases.



DB2 Compatibility
At the time of release, there were a number of issues performing database queries for labels or RSS feeds against DB2 databases. These issues will be fixed for the next release. For now, though, we do not recommend testing this release on DB2.



Macro Compatibility
Incompatibilities exist that may cause Confluence not to start if custom macros are deployed. When upgrading to Confluence 1.5-DR1, be sure to remove any custom Macro plugins from your `$confluencehome/plugins` and `WEB-INF/lib` directories. These incompatibilities will be resolved before the final, stable release.

Upgrade Procedure

Upgrading Confluence should be pretty easy. If you are upgrading from Confluence 1.4 or higher, you can find instructions here. *We strongly recommend that you backup your `confluence.home` directory and database before upgrading!*

If you are upgrading from Confluence 1.4.3 or earlier, be sure to read the upgrade instructions in the Confluence 1.4 release notes.

Downloads

All development releases are available from Development Releases on the Atlassian website.

New Features in Confluence 1.5-DR1

The three major features we've been working on are:

- WYSIWYG Editing
- Labels for content
- A dynamic RSS Builder

Consult JIRA for the full list of issues resolved for 1.5-DR1

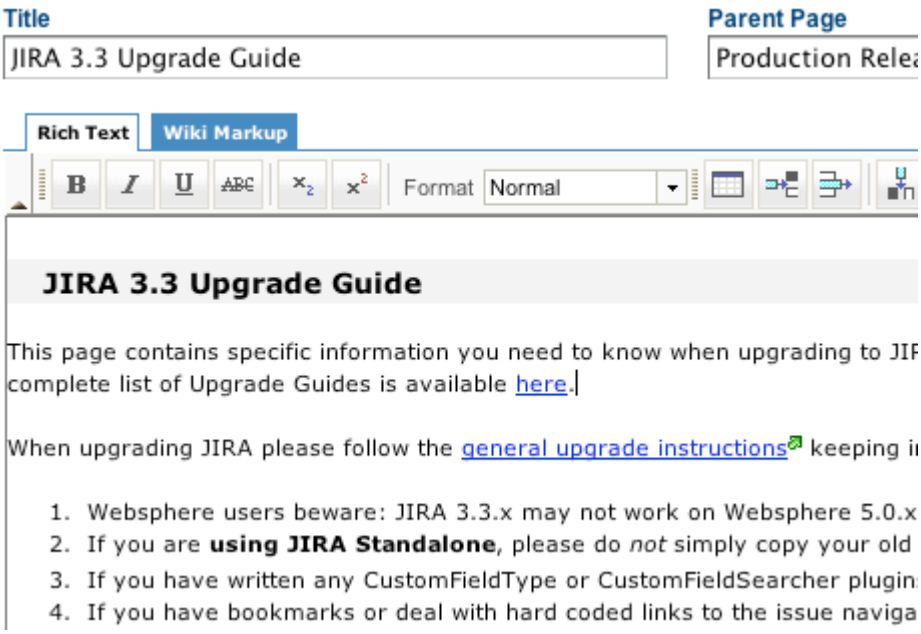
1.5-DR1 also incorporates all the bug-fixes that were made between Confluence 1.4 and Confluence 1.4.3.

WYSIWYG Editing



Browser Compatibility
The Confluence WYSIWYG editor is currently only compatible with Internet Explorer 6 on Windows, plus Mozilla and Firefox across platforms. Javascript must be enabled in the browser for the editor to function. Support for Safari under Mac OS X is currently not available. To track Safari compatibility, please follow this JIRA issue: CONF-3864

The WYSIWYG editor allows for Confluence pages to be edited directly through an editing GUI embedded in the web browser, without having to remember Confluence's wiki markup. It almost goes without saying that this has been our number one most requested features, and we're glad we can finally offer it!



Global administrators can enable WYSIWYG editing in the General Configuration screen of the site's preferences. They can also choose whether users are presented with the WYSIWYG editor by default, or whether users default to the old wiki markup text-field.

Unable to render embedded object: File (Wysiwyg Editor Global Prefs.png) not found.

If WYSIWYG markup is enabled, but an individual user does not like the default set by the administrator, they are free to override it in their preferences.

Unable to render embedded object: File (Wysiwyg Editor Personal Prefs.png) not found.

Labels for content

Another highly requested feature was the ability to categorise content within Confluence, beyond the rigid heirarchy allowed by spaces and parent-child relationships between pages. To this end we have introduced labels. Labels are simple one-word 'tags' that can be added to any page or blog-post that the user has permission to edit. Labels can be used to categorise content, bookmark it, flag it for attention, or anything else you can think of.

> **A Tag By Any Other Name**
> Picking a name for labels wasn't easy. Google's GMail service calls them labels, while other collaborative categorisation systems such as del.icio.us and Flickr call them tags. We decided that 'label' was a more natural description. For more information about the philosophy behind labels and tags, check out the Folksonomy page on Wikipedia

Labels can be added to any page from the edit screen, as a space-separated list of words. If you are browsing the site with a modern, Javascript-enabled browser, you'll also be able to use the dynamic web UI to add labels directly from the page itself.

Unable to render embedded object: File (view page add UI) not found.

Once a page is labeled, then clicking on the label will enable you to browse other pages that have been similarly tagged, or browse other labels that commonly occur on the same page.

Unable to render embedded object: File (label browser UI) not found.

You can also view the space's 200 most popular labels from the space browser, to get an idea of the most popular concerns within the space.

Unable to render embedded object: File (space level heatmap) not found.

Many other Confluence features interact with labels: they can be searched for through the search interface, the new RSS builder can filter pages by their label, and we're looking forward to building label support into Confluence's packaged macros, and even the dashboard.

Personal Labels

If you prepend `my:` to a label (for example, `my:todo` or `my:favourite`, then the label is a personal label – only visible to you. Personal labels allow you to discreetly tag content for your own purposes. You can browse your personal labels from your user profile. Any user can add their personal labels to a page, even if they don't have permission to edit it.

Unable to render embedded object: File (personal label page) not found.

Dynamic RSS Builder

A third highly-requested feature. Confluence has always provided a brace of useful RSS feeds, but the problem is that for every feed we provided, users wanted half a dozen more. The obvious answer is to let users build RSS feeds based on their own chosen criteria. You can access the RSS builder from the Confluence dashboard

Unable to render embedded object: File (dashboard link) not found.

Once in the builder, you can choose

- Which spaces to include in the feed
- Which types of content should be tracked
- Which labels, if any, you are interested in
- How many items to include in the feed
- Whether you want a single RSS entry per page or one for each time the page is edited
- Whether you want an RSS 2.0 or Atom 0.3 feed
- Whether Confluence should require authentication to view the feed

Unable to render embedded object: File (Builder form) not found.

Once you have decided what you want, Confluence will give you a URL to paste into your RSS reader. These URLs can be shared with other Confluence users, although they will only ever be allowed to see content that they have permission to view. If you have asked to authenticate, Confluence will require HTTP Basic Authentication, which is supported by most RSS readers.

We've also taken the opportunity to improve the presentation of our RSS feeds - including a lot more information in each feed so you can follow your Confluence site entirely from your newsreader.

Unable to render embedded object: File (Newsfire screenshot) not found.



The Atom 1.0 draft has just been accepted as an IETF standard. Future versions of Confluence will be phasing out use of Atom 0.3 in favour of Atom 1.0

Known Bugs

Confluence 1.5-DR1 is a preview, not a full Confluence release, and as such there are a number of known bugs included in the release (at no extra cost!). Important bugs include:

- Label and RSS database queries do not work on DB2
- WYSIWYG editor adds extra information to browser history on Firefox
- Going to a page when not logged in will present you with a 404 error page, instead of giving you the chance to log in
- Various WYSIWYG round-trip bugs

Release Notes 1.5-DR2

Confluence 1.5-DR2 is the first public development release leading up to Confluence 1.5. (Confluence 1.5-DR1 was an internal release only). Developer releases are a snapshot of our work in progress, allowing our customers to see what we're up to, and provide feedback

Who should upgrade?

Development releases are snapshots of the ongoing Confluence development process. We make them available for customers who are willing to risk an unpolished release in order to have early access to new features. Developer releases are not suitable for running on production systems.

If you want to be running the most stable and most reliable version of Confluence, you should stick with the official, numbered releases.



Macro Compatibility
Incompatibilities exist that may cause Confluence not to start if custom macros are deployed. When upgrading to Confluence 1.5-DR2, be sure to remove any custom Macro plugins from your `$confluencehome/plugins` and `WEB-INF/lib` directories. These incompatibilities should be resolved before the final, stable release.

Upgrade Procedure

Upgrading Confluence should be pretty easy. If you are upgrading from Confluence 1.4 or higher, you can find instructions here. *We strongly recommend that you backup your `confluence.home` directory and database before upgrading!*

If you are upgrading from Confluence 1.3.x or earlier, be sure to read the upgrade instructions in the Confluence 1.4 release notes.

Downloads

All development releases are available from Development Releases on the Atlassian website.

New Features in Confluence 1.5-DR2

We've been pretty busy, but the four major new features you can find in 1.5-DR2 are:

- WYSIWYG Editing
- Labels for content
- A dynamic RSS Builder
- Change summaries

Consult JIRA for the full list of issues resolved for 1.5-DR2. 1.5-DR2 also incorporates all the bug-fixes that were made between TestTest.

WYSIWYG Editing

> **Browser Compatibility**
> The Confluence WYSIWYG editor is currently only compatible with Internet Explorer 6 on Windows, plus Mozilla and Firefox across platforms. Javascript must be enabled in the browser for the editor to function. Support for Safari under Mac OS X is currently not available. To track Safari compatibility, please follow this JIRA issue: CONF-3864

The WYSIWYG editor allows for Confluence pages to be edited directly through an editing GUI embedded in the web browser, without having to remember Confluence's wiki markup. It almost goes without saying that this has been our number one most requested features, and we're glad we can finally offer it!

Global administrators can enable WYSIWYG editing in the General Configuration screen of the site's preferences. They can also choose whether users are presented with the WYSIWYG editor by default, or whether users default to the old wiki markup text-field.



If WYSIWYG markup is enabled, but an individual user does not like the default set by the administrator, they are free to override it via a "make this my default" link that will appear on whichever editor is currently not your default.
For the "feature mad" amongst us, here are some neat things you can do with the WYSIWYG editor:

- Full screen view - really useful for editing large pages. Click in the menu bar.
- Quickly switch between WYSIWYG and Wiki markup without a page refresh
- Change the size of your editing window to suit your browser. Your size preference is remembered across sessions. To change it, drag the handle in the bottom right hand corner of the editor.
- Undo and redo!

Labels for content

Another highly requested feature was the ability to categorise content within Confluence beyond the rigid heirarchy allowed by spaces and parent-child relationships between pages. To this end we have introduced labels.



Labels are simple one-word 'tags' that can be added to any page or blog-post the user has permission to edit. Labels can be used to categorise content, bookmark it, flag it for attention, or anything else you can think of.

> **A Tag By Any Other Name**
> Picking a name for labels wasn't easy. Google's GMail service calls them labels, while other collaborative categorisation systems such as del.icio.us and Flickr call them tags. We decided that 'label' was a more natural description. For more information about the philosophy behind labels and tags, check out the Folksonomy page on Wikipedia

Labels can be added to any page from the edit screen, as a space-separated list of words. If you are browsing the site with a modern, Javascript-enabled browser, you'll also be able to use the dynamic web UI to add labels while viewing the page through an interactive interface.



Once a page is labeled, then clicking on the label's name allows you to browse other pages with the same label, or view related labels that commonly occur on the same pages.
You can also view the space's most popular labels from the space browser, to get an idea of the most popular topics within the space.



Many other Confluence features interact with labels: they can be searched for through the search interface, the new RSS builder can filter pages by their label, and we're looking forward to building label support into Confluence's packaged macros, and even the dashboard.

Personal Labels

If you prepend `my:` to a label (for example, `my:todo` or `my:favourite`, then the label is a personal label – only visible to you. Personal labels allow you to tag content for your own purposes. You can browse your personal labels from your user profile. Any user can add their personal labels to any page, even when they don't have editing permission.



Personal Label Privacy
It is possible in Confluence 1.5-DR2 to see other people's personal labels in various views through the system. This is a known issue and in the final release, personal labels will be private to the user. (sharing personal labels at the user's discretion is also planned for the future)

Favourites

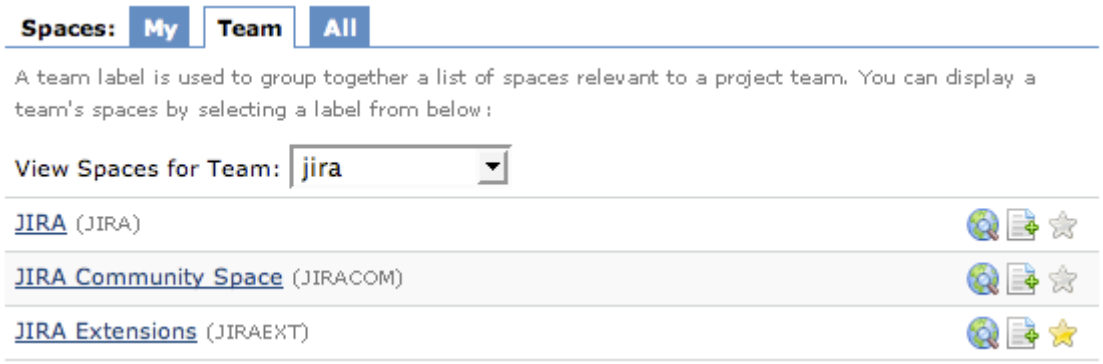
Favourites are a special personal label: `my:favourite` or `my:favorite`. Whenever you see the icon, it means you can label this content as being your favourite, and whenever you see the icon, it means that the content is currently in your list of favourites. You can view your favourites from the Labels tab of your user profile, but they come in handy for...

Labels and the Dashboard

Labels can be used on the Dashboard to create different views of the Confluence site. The list of spaces now offers the following tabs:

- **My** shows you spaces that you have labelled as a Favourite. You can label any space you see on the dashboard by clicking its icon.
- **Team** allows space admins to dynamically create lists of spaces on particular topics (see below. This tab only appears if there are teams that the user can see)
- **New** shows any spaces that have been created in the last week (This tab only appears if there are new spaces)
- **All** shows all the spaces in the Confluence site

A "team" is a group of spaces that share a common team label. Spaces can be labeled from the Advanced tab of Browse Space.



As of Confluence 1.5-DR2, anyone can add space-level labels. This will be fixed in the final release.

The Recent Changes list on the dashboard will show only content that has been changed in the spaces that are currently listed in the space view. So if you are currently viewing the 'My' tab, only updates in your favourite spaces will be shown on the dashboard.

If you are logged in, Confluence will remember your most recently selected tab and team on the dashboard.

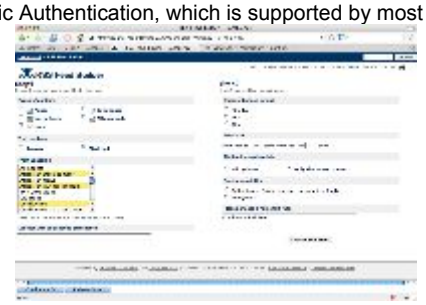
Dynamic RSS Builder

A third highly-requested feature. Confluence has always provided a brace of useful RSS feeds, but the problem is that for every feed we provided, users wanted half a dozen more. The obvious answer is to let users build RSS feeds based on their own chosen criteria. You can access the RSS builder from the Confluence dashboard

Once in the builder, you can choose

- Which spaces to include in the feed
- Which types of content should be tracked
- Which labels, if any, you are interested in
- How many items to include in the feed
- Whether you want a single RSS entry per page or one for each time the page is edited
- Whether you want an RSS 2.0 or Atom 0.3 feed
- Whether Confluence should require authentication to view the feed

Once you have decided what you want, Confluence will give you a URL to paste into your RSS reader. These URLs can be shared with other Confluence users, although they will only ever be allowed to see content that they have permission to view. If you have asked to authenticate, Confluence will require HTTP Basic Authentication, which is supported by most RSS readers.

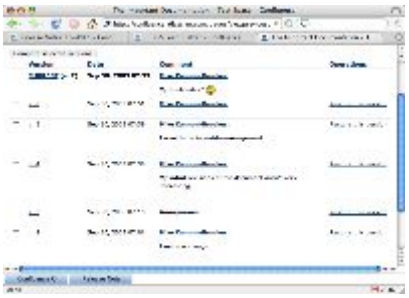


We've also taken the opportunity to improve the presentation of our RSS feeds - including a lot more information in each feed so you can follow your Confluence site entirely from your newsreader.

The Atom 1.0 draft has just been accepted as an IETF standard. Future versions of Confluence will be phasing out use of Atom 0.3 in favour of Atom 1.0

Change summaries

In a very late addition (added during our recent ShipIt Day 2), Confluence now has change summaries. These allow you to add a comment to each edit that will appear in change histories, allowing you to keep a more complete record of how and why a particular page has been modified.



There is also a {`change-history`} macro to enable you to display a page's history within its body if you so desire.

Other Things to Check Out

Embed Flash and Movies

You can now embed Flash content or movies (Quicktime or Windows Media) into a page as easily as you can an image: just attach the Flash or movie file to the page, then include it as you would include an image (!filename.mov!).

Export Pages as Word Documents

You can now export pages straight into Word from the Info tab. This is extremely useful for emailing around content to non-Confluence users, printing a document or just creating a backup in Word.



Improved Search Interface
Results returned from Confluence's search engine now have:

- Improved contextual results, showing the most important text around where your query was matched in the page
- Contextual results for any attachment: see where a search was matched even **inside** PDF, Word, PowerPoint or Excel documents!
- Search results for attachments give you more (and clearer) information about what the attachment is, and where it's from!search-result-fragments.png|thumbnail, align=center!
  Improved Gallery Macro

The gallery macro has been spruced up, and now has a slideshow view:

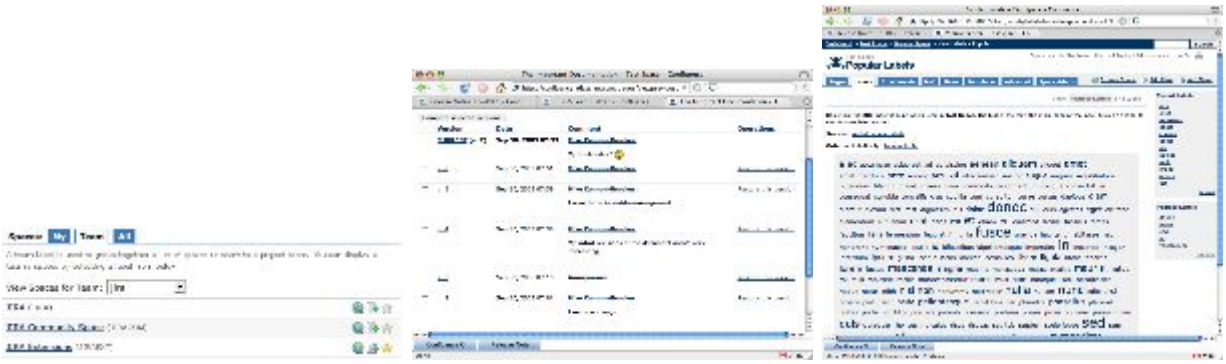
Confluence 1.5DR2 Screenshots



These release notes exported to Word (even on a Mac!) - great for printing too.
The new RSS builder lets you construct RSS feeds of exactly the content you want.
Screenshot of the new 1.5 WYSIWYG editor.



View page label interface showing autocomplete.
Improved search showing fragments from attachments, file sizes and types.
My personal labels! Only for me!



The new Dashboard space tabs, here showing the "jira" team tab and one favourite space.
A very basic (quite boring - sorry, it's late) example of change summaries.
Popular labels for a given space.



Alphabetically organised labels for a given space.

Also:

- Additions to the Confluence Remote API including:
  - Comment manipulation
  - Label manipulation
  - Attachment uploading and editing
  - Improved user- and permissions management
- Import and restore now have progress indicators
- Backup and restore use significantly less memory
- The embedded database has been upgraded to HSQL 1.8, which should be significantly more reliable
- Collapsed breadcrumbs now expand with a single mouse click

Known Bugs

Confluence 1.5-DR2 is a preview, not a full Confluence release, and as such there are a number of known bugs included in the release (at no extra cost!). Important bugs include:

- The left-navigation theme is currently broken
- The "make this my default editor" link does not always appear (try switching back and forth between views)
- The Info page may cause a Hibernate exception when its parent has page level permissions
- The space export may be unreliable, and does not properly back up labels.
- Incoming Trackback pings are not recognised
- Various WYSIWYG round-trip inconsistencies

table test

Release Notes 2.0-RC1

Confluence 2.0-RC1 is the first release candidate for Confluence 2.0 (previously known as Confluence 1.5). It resolves almost 150 issues since the 1.5-RC2 release.

Who should upgrade?

Development releases are snapshots of the ongoing Confluence development process. We make them available for customers who are willing to risk an unpolished release in order to have early access to new features. Developer releases are not suitable for running on production systems.

If you want to be running the most stable and most reliable version of Confluence, you should stick with the official, numbered releases.

Upgrade Procedure

Upgrading Confluence should be pretty easy. If you are upgrading from Confluence 1.4 or higher, you can find instructions here. *We strongly recommend that you backup your `confluence.home` directory and database before upgrading!*

If you are upgrading from Confluence 1.3.x or earlier, be sure to read the upgrade instructions in the Confluence 1.4 release notes.



Important Notice
If, after you have upgraded Confluence, the "Recent Changes" list on the dashboard is empty, this may be because you need to re-build your site's search index. You can do this under `Administration  Content Indexing  Rebuild Search Index`

Downloads

All development releases are available from Development Releases on the Atlassian website.

New Features in Confluence 2.0-RC1

The 2.0-RC1 has largely concentrated on fixing bugs and polishing existing features. If you are upgrading from Confluence 1.4 or earlier, you should read the 1.5-DR2 Release Notes for a description of the major new features in Confluence 2.0.

Consult JIRA for the full list of issues resolved for 2.0-RC1.

Notable Features and Improvements

- You can now place markup such as bold or italics within a word if you are using the WYSIWYG editor
- The following characters are now permitted in page titles: `! & ( ) * ~ $ _` Pages can not, however, start with `~` or `$`.
- The most recent change comment for a page is now included in the page's searchable text
- `rel="nofollow"` added to links to printable versions of pages, to prevent them from being indexed by search engines
- The Recent Changes list on the dashboard now correctly reflects whether you are viewing all spaces, a team, or your favourite spaces
- Further improvements to the edit page UI including:
    - Improved label editing
    - Inline page previewing
    - Hiding of lesser-used functions such as moving a page or editing page permissions
- Attachment downloads now support HTTP conditional get
- Space exports now include that space's labels and page-level properties
- Support for labels in the SOAP and XML-RPC remote APIs has been completed
- Determining if a user exists via the remote API is now supported

Release Notes 2.0-RC2

Confluence 2.0-RC2 is the second release candidate for Confluence 2.0 (previously known as Confluence 1.5). It resolves almost 80 issues since the 1.5-RC1 release.

Who should upgrade?

Development releases are snapshots of the ongoing Confluence development process. We make them available for customers who are willing to risk an unpolished release in order to have early access to new features. Developer releases are not suitable for running on production systems.

If you want to be running the most stable and most reliable version of Confluence, you should stick with the official, numbered releases.

Upgrade Procedure

Upgrading Confluence should be pretty easy. If you are upgrading from Confluence 1.4 or higher, you can find instructions here. *We strongly recommend that you backup your `confluence.home` directory and database before upgrading!*

If you are upgrading from Confluence 1.3.x or earlier, be sure to read the upgrade instructions in the Confluence 1.4 release notes.



Important Notice
If, after you have upgraded Confluence, the "Recent Changes" list on the dashboard is empty, this may be because you need to re-build your site's search index. You can do this under `Administration  Content Indexing  Rebuild Search Index`



Important Notice 2
If, after you have upgraded Confluence and you find the page breadcrumbs do not show the page's parents correctly, log into your site as administrator, then visit http://yoursite.example.com/admin/permissions/rebuild_ancestor_table.action (substituting your own site's URL, obviously) to rebuild the ancestor table (this will be fixed for the final release.

Downloads

All development releases are available from Development Releases on the Atlassian website.

New Features in Confluence 2.0-RC2

The 2.0-RC2 has largely concentrated on fixing bugs and polishing existing features. If you are upgrading from Confluence 1.4 or earlier, you

should read the 1.5-DR2 Release Notes for a description of the major new features in Confluence 2.0.

Consult JIRA for the full list of issues resolved for 2.0-RC2.

Error formatting macro: jiraissues: java.lang.RuntimeException: A value with ID '11500' does not exist for the field 'fixVersion'.

Release Notes 2.3-DR1

Confluence 2.3-DR1 is the first public development release leading up to Confluence 2.3. Development releases are a snapshot of our work in progress, allowing our customers to see what we're up to.

Please note the following

- **Development releases are not safe** — Development releases are snapshots of the ongoing Confluence development process. As such:
  - While we try to keep these releases stable, they have not undergone the same degree of testing as a full release.
  - Features in development releases may be incomplete, or may change or be removed before the next full release.
- **No upgrade path** — Because development releases represent work in progress, we **can not** provide a supported upgrade path between development releases, or from any development release to the eventual final release. Thus, it is possible that you will not be able to migrate any data you store in a Confluence development release to a future Confluence release.

Who should upgrade?

This development release is being made available specifically for Confluence plugin developers. The changes to the Confluence API discussed below mean that many plugins will need to be updated to work with Confluence 2.3.

For all production use and testing of Confluence, please use the latest official release.

The bundled Tangosol library with this development release has a license which expires on August 31, 2006. This release will not operate after that date.

This release should not be used for testing a clustered deployment; the clustering user interface is not yet complete.

Upgrade Procedure

Follow the normal upgrade instructions to upgrade from Confluence 2.2.x to this release. *We strongly recommend that you backup your confluence-home directory and database before upgrading!*

If you are upgrading from Confluence 2.1.x or earlier, be sure to read the upgrade instructions in the Confluence 2.2 release notes.

Downloads

All development releases are available from Development Releases on the Atlassian website.

New features in Confluence 2.3

In short, the new features of Confluence 2.3 are:

- Clustering support (not available in 2.3-DR1)
- People directory
- Bundled Plugin Repository Client

Architecture changes between Confluence 2.2 and 2.3-DR1

Summarised below are architectural changes that might be relevant to plugin developers. Please add a comment below if there is anything else that should be listed here.

We've kicked the Bucket!

Most of the functionality of Confluence's biggest internal library, bucket, has been split into three new components: atlassian-config, atlassian-spring and atlassian-event.

Where practical, the old interface is still available as deprecated classes and methods. **However, some plugins will not compile unless modified to use the new package names.**

Please refer to the source code or Javadoc for details on the new interface.

Tangosol Coherence replaces EhCache

Tangosol Coherence is now the caching library used by Confluence, both in clustered and non-clustered mode. To facilitate this, some of the caching APIs have been updated.

As mentioned above, the Tangosol license included with this release is specifically for testing purposes and will expire on August 31, 2006.

Bandana configuration stored in database

The Bandana Confluence configuration files previously stored in confluence-home are now stored in the database. An upgrade from a previous version will automatically move existing configuration settings to the database.

Release Notes 2.3-DR2

Confluence 2.3-DR2 is the second public development release leading up to Confluence 2.3. Development releases are a snapshot of our work in progress, allowing our customers to see what we're up to.

Who should upgrade?



Please note the following

- **Development releases are not safe** — Development releases are snapshots of the ongoing Confluence development process. As such:
  - While we try to keep these releases stable, they have not undergone the same degree of testing as a full release.
  - Features in development releases may be incomplete, or may change or be removed before the next full release.
- **No upgrade path** — Because development releases represent work in progress, we **can not** provide a supported upgrade path between development releases, or from any development release to the eventual final release. Thus, it is possible that you will not be able to migrate any data you store in a Confluence development release to a future Confluence release.

This development release is being made available specifically for Confluence plugin developers. The changes to the Confluence API discussed below mean that many plugins will need to be updated to work with Confluence 2.3.

For all production use and testing of Confluence, please use the latest official release.

This release should not be used for testing a clustered deployment; the clustering user interface is not yet complete.

Upgrade Procedure

Follow the normal upgrade instructions to upgrade from Confluence 2.2.x to this release. *We strongly recommend that you backup your confluence-home directory and database before upgrading!*

If you are upgrading from Confluence 2.1.x or earlier, be sure to read the upgrade instructions in the Confluence 2.2 release notes.

Downloads

All development releases are available from Development Releases on the Atlassian website.

Things to look out for

- The Tangosol Coherence jars included in this DR are now fully licensed.
- The API for retrieving historical versions of content has changed.

Architecture changes between Confluence 2.3-DR1 and 2.3-DR2

Content history API changed

The `getPreviousVersions` method of `ContentEntityObject` has been removed, because it was insanely inefficient as soon as a piece of content started having lots of versions. It has been replaced with the following methods on the `ContentEntityManager`:

```
01.ContentEntityObject getPreviousVersion(ContentEntityObject ceo);
02.
03.ContentEntityObject getNextVersion(ContentEntityObject ceo);
04.
05.ContentEntityObject getOtherVersion(ContentEntityObject ceo, int version);
06.
07./**
08.   * Get a VersionHistorySummary for all previous versions of a ContentEntityObject, starting
with the current
09.   * content.
10.   *
11.   * @param ceo the entity object to return the version history of
12.   * @return the full version history of that object, as VersionHistorySummary objects.
13.   */
14.List getVersionHistorySummaries(ContentEntityObject ceo);
```

The `VersionHistorySummary` class defines a limited set of Content data that is relevant to viewing version histories.

Release Notes 2.6-dr1

Confluence 2.6-dr1 is a public development release leading up to Confluence 2.6. Development releases are a snapshot of our work in progress, allowing our customers to see what we're up to.

Who should upgrade?



Please note the following

- **Development releases are not safe** — Development releases are snapshots of the ongoing Confluence development process. As such:
  - While we try to keep these releases stable, they have not undergone the same degree of testing as a full release.
  - Features in development releases may be incomplete, or may change or be removed before the next full release.
- **No upgrade path** — Because development releases represent work in progress, we **can not** provide a supported upgrade path between development releases, or from any development release to the eventual final release. Thus, it is possible that you will not be able to migrate any data you store in a Confluence development release to a future Confluence release.

This development release is being made available specifically for Confluence plugin developers to test their existing plugins against the significant style changes that have been made.

For all production use and testing of Confluence, please use the latest official release.

Upgrade Procedure

Follow the normal upgrade instructions to upgrade from Confluence 2.5.x to this release. *We strongly recommend that you backup your confluence-home directory and database before upgrading!*

Downloads

All development releases are available from Development Releases on the Atlassian website.

Note about themes

Confluence 2.5.x themes are expected to be compatible with 2.6 without authors needing to make any change to their existing themes. This is possible because Confluence will, by default, include all Confluence 2.5.x specific styles automatically.

However, if you would like to upgrade your theme to use the latest style and typography changes in Confluence 2.6, you will need to update the way you include stylesheets in your theme. Full instructions on how to do so can be found here.

Including Cascading Stylesheets in Themes for Confluence 2.6

Firstly, you will need to disable the inclusion of Confluence 2.5.x styles. These were included by default to allow 2.5.x themes to remain compatible in Confluence version 2.6 and after.

To disable 2.5.x styles, add the following to your theme's `atlassian-plugin.xml`:

```
1....
2.<theme key="aqua" i18n-name-key="com.atlassian.confluence.themes.aqua.name" name="Aqua Theme" class
="com.atlassian.confluence.themes.BasicTheme">
3.          <description key="com.atlassian.confluence.themes.aqua.desc"/>
4.          ...
5.          <param name="includeClassicStyles" value="false"/>
6.          ...
7.</theme>
8....
```

Note the flag `includeClassicStyles` is set to `false`. You will have to set this explicitly as the default is `true`.

Secondly, you will need to declare your custom theme style sheet in `atlassian-plugin.xml` like this:

```
1....
2.<theme key="aqua" i18n-name-key="com.atlassian.confluence.themes.aqua.name" name="Aqua Theme" class
="com.atlassian.confluence.themes.BasicTheme">
3.          <description key="com.atlassian.confluence.themes.aqua.desc"/>
4.          ...
5.          <resource type="stylesheet" name="my.css" location="styles/my-css.vm"/>
6.          ...
7.</theme>
8....
```

Differences from Confluence 2.5.x

- You no longer have to include your theme stylesheet in the main decorator using `#pluginStylesheet` anymore. Confluence will load your theme's stylesheet automatically provided that its the active theme
- The resource is declared in the `theme` module instead of the `layout` module
- You need to start your custom style sheet (say `my-css.vm`), by copying over the latest styles from `http://yourhost/contextPath/styles/main-action.css`. This step is necessary as Confluence now includes either your theme stylesheet or the default stylesheet, **not both**. This implies:
    - you can no longer rely on the default styles being there to style parts of the Confluence you are not directly theming
    - you are no longer overridding styles with your plugin style sheet. It is now the primary stylesheet
    - you will need to merge any new styles in later versions of Confluence into your theme's style sheet

Multiple style sheets

It is possible to configure your theme to use multiple style sheets. This feature may useful if you want to break up your main style sheet into a few smaller style sheets with more defined purposes. You can declare these like so:

```
01....
02.<theme key="aqua" i18n-name-key="com.atlassian.confluence.themes.aqua.name" name="Aqua Theme"
class="com.atlassian.confluence.themes.BasicTheme">
03.          <description key="com.atlassian.confluence.themes.aqua.desc"/>
04.          ...
05.          <resource type="stylesheet" name="my1.css" location="styles/my-css1.vm"/>
06.          <resource type="stylesheet" name="my2.css" location="styles/my-css2.vm"/>
07.          <resource type="stylesheet" name="my3.css" location="styles/my-css3.vm"/>
08.          ...
09.</theme>
10....
```

These style sheets will be included in the order in which they are declared.

Release Notes 2.7-m2 ("Milestone 2")

Do not use this release to upgrade your production systems.

For all production use and testing of Confluence, please use the latest official release.

Confluence 2.7-m2 is the first milestone development release for 2.7. This is a public **development release** leading up to Confluence 2.7. Development releases are a snapshot of our work in progress, allowing our customers to see what we're up to.

Who should upgrade?

Please note the following

- **Development releases are not safe** — Development releases are snapshots of the ongoing Confluence development process. As such:
  - While we try to keep these releases stable, they have not undergone the same degree of testing as a full release.
  - Features in development releases may be incomplete, or may change or be removed before the next full release.
- **No upgrade path** — Because development releases represent work in progress, we **can not** provide a supported upgrade path between development releases, or from any development release to the eventual final release. Thus, it is possible that you will not be able to migrate any data you store in a Confluence development release to a future Confluence release.

Our milestone releases aim to provide plugin developers with an opportunity to see the latest changes in the code.

Each milestone release has passed all our automatic tests, and has been used for one week on our official internal Confluence server. Most of the issues solved have been reviewed too.

However, since our milestones releases are timeboxed (i.e. they get released every two weeks, no matter how far we have come implementing features and bugfixes), there is always a chance that we have new known bugs, which are scheduled to be fixed in the next milestone.

Milestone releases have not been load- or stress-tested. So, for example, they might behave well on a small installation but show severe problems when subjected to many users.

Upgrade Procedure

Follow the normal upgrade instructions to upgrade from Confluence 2.6.x to this release. *We strongly recommend that you backup your confluence-home directory and database before upgrading!*

Downloads

All development releases are available from Development Releases on the Atlassian website.

Bugs fixed in this release

View the list of fixes in JIRA.

Release Notes 2.7-m4 ("Milestone 4")

Do not use this release to upgrade your production systems.

For all production use and testing of Confluence, please use the latest official release.

Confluence 2.7-m4 is the second milestone development release for 2.7. (Milestone 3 has been skipped because of stability concerns.) This is a public **development release** (DR) leading up to Confluence 2.7. Development releases are a snapshot of our work in progress, allowing our customers to see what we're up to.

Who should upgrade?

Please note the following

- **Development releases are not safe** — Development releases are snapshots of the ongoing Confluence development process. As such:
  - While we try to keep these releases stable, they have not undergone the same degree of testing as a full release.
  - Features in development releases may be incomplete, or may change or be removed before the next full release.
- **No upgrade path** — Because development releases represent work in progress, we **can not** provide a supported upgrade path between development releases, or from any development release to the eventual final release. Thus, it is possible that you will not be able to migrate any data you store in a Confluence development release to a future Confluence release.

Our milestone releases aim to provide plugin developers with an opportunity to see the latest changes in the code.

Each milestone release has passed all our automatic tests, and has been used for one week on our official internal Confluence server. Most

of the issues solved have been reviewed too.

However, since our milestones releases are timeboxed (i.e. they get released every two weeks, no matter how far we have come implementing features and bugfixes), there is always a chance that we have new known bugs, which are scheduled to be fixed in the next milestone, or completely new bugs unknown even to us.

Milestone releases have not been load- or stress-tested. So, for example, they might behave well on a small installation but show severe problems when subjected to many users.

Upgrade Procedure

Follow the normal upgrade instructions to upgrade from Confluence 2.6.x to this release. *We strongly recommend that you backup your confluence-home directory and database before upgrading!*

Downloads

All development releases are available from Development Releases on the Atlassian website.

Issues resolved in this release

View the list of fixes and newly implemented features in JIRA. They contain **all** the issues resolved since 2.6.x, not just the ones fixed since the previous milestone.

Release Notes 2.7-m5 ("Milestone 5")



Do not use this release to upgrade your production systems.

For all production use and testing of Confluence, please use the latest official release.

Confluence 2.7-m5 is a milestone development release for 2.7. This is a public **development release** (DR) leading up to Confluence 2.7. Development releases are a snapshot of our work in progress, allowing our customers to see what we're up to.

Who should upgrade?



Please note the following

- **Development releases are not safe** — Development releases are snapshots of the ongoing Confluence development process. As such:
  - While we try to keep these releases stable, they have not undergone the same degree of testing as a full release.
  - Features in development releases may be incomplete, or may change or be removed before the next full release.
- **No upgrade path** — Because development releases represent work in progress, we **can not** provide a supported upgrade path between development releases, or from any development release to the eventual final release. Thus, it is possible that you will not be able to migrate any data you store in a Confluence development release to a future Confluence release.

Our milestone releases aim to provide plugin developers with an opportunity to see the latest changes in the code.

Each milestone release has passed all our automatic tests, and has been used for one week on our official internal Confluence server. Most of the issues solved have been reviewed too.

However, since our milestones releases are timeboxed (i.e. they get released every two weeks, no matter how far we have come implementing features and bugfixes), there is always a chance that we have new known bugs, which are scheduled to be fixed in the next milestone, or completely new bugs unknown even to us.

Milestone releases have not been load- or stress-tested. So, for example, they might behave well on a small installation but show severe problems when subjected to many users.

Upgrade Procedure

Follow the normal upgrade instructions to upgrade from Confluence 2.6.x to this release. *We strongly recommend that you backup your confluence-home directory and database before upgrading!*

Downloads

All development releases are available from Development Releases on the Atlassian website.

Issues resolved in this release

View the list of fixes and newly implemented features in JIRA. They contain **all** the issues resolved during development of 2.7, not just the ones fixed since the previous milestone.

Known problems

For some reason, an outdated version of the new Dynamic Tasklist 2 was included in this milestone. Since Milestones are not real customer releases, we do ship them with bugs like this. If you want to see the latest and greatest Dynamic Tasklist 2 in action, please simply upgrade to the latest version using the plugin repository from the administration console.

Also, there is an annoying (2.7-specific) bug in this release that causes warnings for all colours used by the colour-macro: http://jira.atlassian.com/browse/CONF-10001/. It will be resolved by the time the next milestone gets released.

Apart from that, M5 is a nice Milestone release and definitely worth a look for all plugin developers - especially now that we are getting closer and closer to the end of the release cycle, and fewer and fewer major architectural changes will be made.

Release Notes 2.7-rc1 ("Release Candidate 1")



Do not use this release to upgrade your production systems.

For all production use and testing of Confluence, please use the latest official release.

Confluence 2.7-rc1 is a release candidate for 2.7. This is a public **development release** (DR) leading up to Confluence 2.7. Development releases are a snapshot of our work in progress, allowing our customers to see what we're up to.

Who should upgrade?



Please note the following

- **Development releases are not safe** — Development releases are snapshots of the ongoing Confluence development process. As such:
  - While we try to keep these releases stable, they have not undergone the same degree of testing as a full release.
  - Features in development releases may be incomplete, or may change or be removed before the next full release.
- **No upgrade path** — Because development releases represent work in progress, we **can not** provide a supported upgrade path between development releases, or from any development release to the eventual final release. Thus, it is possible that you will not be able to migrate any data you store in a Confluence development release to a future Confluence release.

Our milestones and release candidates aim to provide plugin developers with an opportunity to see the latest changes in the code.

Each milestone/release candidate has passed all our automatic tests, and has been used for one week on our official internal Confluence server. Most of the issues solved have been reviewed too.

However, since our milestones and release candidates re timeboxed (i.e. they get released every two weeks, no matter how far we have come implementing features and bugfixes), there is always a chance that we have new known bugs, which are scheduled to be fixed in the next milestone, or completely new bugs unknown even to us.

Milestone releases and release candidates have not been load- or stress-tested. So, for example, they might behave well on a small installation but show severe problems when subjected to many users.

Upgrade Procedure

Follow the normal upgrade instructions to upgrade from Confluence 2.6.x to this release. *We strongly recommend that you backup your confluence-home directory and database before upgrading!*

Downloads

All development releases are available from Development Releases on the Atlassian website.

Issues resolved in this release

View the list of fixes and newly implemented features in JIRA. They contain **all** the issues resolved during development of this release, not just the ones fixed since the previous milestone.

Known problems

None.

You should really download this Release Candidate and check if your plugin works with it. If not, use the last few days before the official 2.7.0

release to fix it

Release Notes 2.8-m2 ("Milestone 2")

Do not use this release to upgrade your production systems.

For all production use and testing of Confluence, please use the latest official release.

Confluence 2.8-m2 is a milestone development release for 2.8. This is a public **development release** (DR) leading up to Confluence 2.8. Development releases are a snapshot of our work in progress, allowing our customers and especially plugin-developers to see what we're up to.

Who should upgrade?

Please note the following

- **Development releases are not safe** — Development releases are snapshots of the ongoing Confluence development process. As such:
  - While we try to keep these releases stable, they have not undergone the same degree of testing as a full release.
  - Features in development releases may be incomplete, or may change or be removed before the next full release.
- **No upgrade path** — Because development releases represent work in progress, we **can not** provide a supported upgrade path between development releases, or from any development release to the eventual final release. Thus, it is possible that you will not be able to migrate any data you store in a Confluence development release to a future Confluence release.

Our milestone releases aim to provide plugin developers with an opportunity to see the latest changes in the code.

Each milestone release has passed all our automatic tests, and has been used for one week on our official internal Confluence server. Most of the issues solved have been reviewed too.

However, since our milestones releases are timeboxed (i.e. they get released every two weeks, no matter how far we have come implementing features and bugfixes), there is always a chance that we have new known bugs, which are scheduled to be fixed in the next milestone, or completely new bugs unknown even to us.

Milestone releases have not been load- or stress-tested. So, for example, they might behave well on a small installation but show severe problems when subjected to many users.

Upgrade Procedure

Follow the normal upgrade instructions to upgrade from Confluence 2.7.x to this release. *We strongly recommend that you backup your confluence-home directory and database before upgrading!*

Downloads

All development releases are available from Development Releases on the Atlassian website.

Issues resolved in this release

- Velocity template engine upgrade

Confluence's velocity template engine has been upgrade from 1.3 to 1.5. This shouldn't mean too much to end users however there could be some compatibility problems with existing themes and plugins. Check this upgrade guide that we use internally as well: http://confluence.atlassian.com/display/DOC/Migrating+to+Velocity+1.5

- File upload component upgrade

Pell multipart has been replaced with the Jakarta Commons Fileupload component to handle web browser file uploads. Again this won't have much impact on the Confluence experience but could cause some odd problems when attaching files to Confluence.

- PDF space export

There has been some optimizations made the space PDF export which should result in less server memory usage during export

- New GZIP compression filter

Confluence team have integrated a new GZIP compression filter to achieve more efficient downloads and page views. This is not currently enabled by default, so to test this with your plugin you need to turn it on in the management console. We are using it internally at Atlassian, and the performance improvement while browsing pages is quite substantial.

You can view the complete list of fixes and newly implemented features in JIRA. They contain **all** the issues resolved during development of 2.8, not just the ones fixed since the previous milestone.

Known problems

There are a few cosmetic UI problems related to space logos and to the login screen, none of them impact productivity. Notification mails are empty because of a Velocity macro bug.

Release Notes 2.8-m3 ("Milestone 3")



Do not use this release to upgrade your production systems.

For all production use and testing of Confluence, please use the latest official release.

Confluence 2.8-m3 is a milestone development release for 2.8. This is a public **development release** (DR) leading up to Confluence 2.8. Development releases are a snapshot of our work in progress, allowing our customers and especially plugin-developers to see what we're up to.

Who should upgrade?



Please note the following

- **Development releases are not safe** — Development releases are snapshots of the ongoing Confluence development process. As such:
  - While we try to keep these releases stable, they have not undergone the same degree of testing as a full release.
  - Features in development releases may be incomplete, or may change or be removed before the next full release.
- **No upgrade path** — Because development releases represent work in progress, we **can not** provide a supported upgrade path between development releases, or from any development release to the eventual final release. Thus, it is possible that you will not be able to migrate any data you store in a Confluence development release to a future Confluence release.

Our milestone releases aim to provide plugin developers with an opportunity to see the latest changes in the code.

Each milestone release has passed all our automatic tests, and has been used for one week on our official internal Confluence server. Most of the issues solved have been reviewed too.

However, since our milestones releases are timeboxed (i.e. they get released every two weeks, no matter how far we have come implementing features and bugfixes), there is always a chance that we have new known bugs, which are scheduled to be fixed in the next milestone, or completely new bugs unknown even to us.

Milestone releases have not been load- or stress-tested. So, for example, they might behave well on a small installation but show severe problems when subjected to many users.

Upgrade Procedure

Follow the normal upgrade instructions to upgrade from Confluence 2.7.x to this release. *We strongly recommend that you backup your confluence-home directory and database before upgrading!*

Downloads

All development releases are available from Development Releases on the Atlassian website.

Issues resolved or improved in this release

- First draft of UI improvements: We are working hard on moving the HTML and CSS delivered by Confluence towards standards-compliance and accessibility. We are at the same time reworking our UI structure. We feel that the UI has been growing a bit too organically over the past years, resulting in somewhat confusing navigation like the "remove page"-button only available when you are editing a page, for example. We have implemented an initial dropdown-menu for adding content, thus already cleaning up the UI a bit. All the navigation-changes we make are targetted at the default theme, however some HTML&CSS changes may also affect the other themes.
  In the future our changes will enable plugin-developers to hook into Confluence easier than before, and skinning Confluence will be considerably easier. However, plugin-developers should be aware of the changes we are making, and check (by examining this Milestone release) whether they have to schedule some maintenance work for their plugins. More changes will be made, so don't rush it - but plan for it.
- First draft of Page Reordering. We are working on implementing a highly anticipated feature that will dramatically improve the quality of content generated from a Confluence space. So far the PDF files you generate from Confluence have been alphabetically ordered - which for example makes even the Confluence User Manual quite awkward to read. With the new page ordering feature it will be able to structure your Wiki exactly as you like - or of course leave it in alphabetic mode if you prefer that. Our goal is to deliver a

really useful solution, and this Milestone only shows you the first 50% of the feature. We are currently busy implementing all the feedback and bug-reports we get from Sarah (our documentation queen) and the next Milestone releases will deliver further improvements and bugfixes.

You can view the complete list of fixes and newly implemented features in JIRA. They contain **all** the issues resolved during development of 2.8, not just the ones fixed since the previous milestone.

Known problems

No major known problems, but plenty of smaller to medium sized bugs, mainly related to page reordering. They will get tackled in M4.

Release Notes 2.8-m4 ("Milestone 4")



Do not use this release to upgrade your production systems.

For all production use and testing of Confluence, please use the latest official release.

Confluence 2.8-m4 is a milestone development release for 2.8. This is a public **development release** (DR) leading up to Confluence 2.8. Development releases are a snapshot of our work in progress, allowing our customers and especially plugin-developers to see what we're up to.

Who should upgrade?



Please note the following

- **Development releases are not safe** — Development releases are snapshots of the ongoing Confluence development process. As such:
  - While we try to keep these releases stable, they have not undergone the same degree of testing as a full release.
  - Features in development releases may be incomplete, or may change or be removed before the next full release.
- **No upgrade path** — Because development releases represent work in progress, we **can not** provide a supported upgrade path between development releases, or from any development release to the eventual final release. Thus, it is possible that you will not be able to migrate any data you store in a Confluence development release to a future Confluence release.

Our milestone releases aim to provide plugin developers with an opportunity to see the latest changes in the code.

Each milestone release has passed all our automatic tests, and has been used for one week on our official internal Confluence server. Most of the issues solved have been reviewed too.

However, since our milestones releases are timeboxed (i.e. they get released every two weeks, no matter how far we have come implementing features and bugfixes), there is always a chance that we have new known bugs, which are scheduled to be fixed in the next milestone, or completely new bugs unknown even to us.

Milestone releases have not been load- or stress-tested. So, for example, they might behave well on a small installation but show severe problems when subjected to many users.

Upgrade Procedure

Follow the normal upgrade instructions to upgrade from Confluence 2.7.x to this release. *We strongly recommend that you backup your confluence-home directory and database before upgrading!*

Downloads

All development releases are available from Development Releases on the Atlassian website.

Issues resolved or improved in this release

This release mainly fixes bugs and glitches of the previous milestone, and extends the functionality slightly. You will notice that a new dropdown-menu has been added which allows you to access the browse-space actions easier, and that most macros now honor the new page ordering that can be set in recently added page-reorder-component.

We had forgotten to release the source code of the previous milestones, this has been taken care of now, thanks Alain from Adaptavist for pointing this out! 

The coolest improvement however is that we are able to deliver the first ShipIt 7 project in this milestone - less than two weeks since the actual event! (see http://blogs.atlassian.com/developer/2007/09/atlassian_shipit_day_vi.html) for an overview of ShipIt 6 last year). Use the new drop-down menu to "Browse Labels", select one of the most popular ones, and you will now be able to filter by multiple labels - just increase or decrease the number of labels by using the plus- and minus-signs next to the related labels. Admittedly the UI still needs some

finishing touches, but it works fine already and it will make 21 voters happy once released officially (http://jira.atlassian.com/browse/CONF-4577). More ShipIt projects can be expected to make it into M5 and M6.

You can view the complete list of fixes and newly implemented features in JIRA. They contain **all** the issues resolved during development of 2.8, not just the ones fixed since the previous milestone.

Known problems

There are still a few known bugs in this release, most noticeably the broken printable view (CONF-10583) which will be fixed soon. Please continue reporting problems through JIRA, your feedback has been very helpful and a lot of it has been incorporated already.

The CacheManager API has changed in milestone 4, and all plugins that use the old CacheManager interface will be broken (CONF-10602). In milestone 5, Confluence has a workaround for plugins that get the cacheManager injected. That is, the following code will work correctly in 2.8-m5 (but is currently broken in 2.8-m4):

```
01.import com.atlassian.user.impl.cache.CacheManager;
02.
03.public class MyClass {
04.           private CacheManager cacheManager;
05.
06.           public void setCacheManager(CacheManager cacheManager) {
07.                     this.cacheManager = cacheManager;
08.           }
09.}
```

Plugins that access the CacheManager not via dependency injection, but statically through the ContainerManager will fail at runtime with a ClassCastException. These plugins should be rewritten to use dependency injection, or use the new `com.atlassian.cache.CacheManager` interface. The following code is broken in 2.8-m4, and will continue to be broken for the final release of 2.8:

```
1.CacheManager cacheManager = (CacheManager) ContainerManager.getComponent("cacheManager");
```

Plugins that wish to avoid using the deprecated CacheManager interface should change to use `com.atlassian.cache.CacheManager`.

Release Notes 2.8-m6 ("Milestone 6")



Do not use this release to upgrade your production systems.

For all production use and testing of Confluence, please use the latest official release.

Confluence 2.8-m6 is a milestone development release for 2.8. This is a public **development release** (DR) leading up to Confluence 2.8. Development releases are a snapshot of our work in progress, allowing our customers and especially plugin-developers to see what we're up to.

Who should upgrade?



Please note the following

- **Development releases are not safe** — Development releases are snapshots of the ongoing Confluence development process. As such:
  - While we try to keep these releases stable, they have not undergone the same degree of testing as a full release.
  - Features in development releases may be incomplete, or may change or be removed before the next full release.
- **No upgrade path** — Because development releases represent work in progress, we **can not** provide a supported upgrade path between development releases, or from any development release to the eventual final release. Thus, it is possible that you will not be able to migrate any data you store in a Confluence development release to a future Confluence release.

Our milestone releases aim to provide plugin developers with an opportunity to see the latest changes in the code.

Each milestone release has passed all our automatic tests, and has been used for one week on our official internal Confluence server. Most of the issues solved have been reviewed too.

However, since our milestones releases are timeboxed (i.e. they get released every two weeks, no matter how far we have come implementing features and bugfixes), there is always a chance that we have new known bugs, which are scheduled to be fixed in the next milestone, or completely new bugs unknown even to us.

Milestone releases have not been load- or stress-tested. So, for example, they might behave well on a small installation but show severe problems when subjected to many users.

Upgrade Procedure

Follow the normal upgrade instructions to upgrade from Confluence 2.7.x to this release. *We strongly recommend that you backup your confluence-home directory and database before upgrading!*

Downloads

All development releases are available from Development Releases on the Atlassian website.

Issues resolved or improved in this release

Milestone 6 incorporates Milestone 5 changes (which were released but never publicly documented, sorry)

Milestone 5 (the missed one)

The user interface improvements include the most drastic changes to be done for Confluence 2.8. After adding a new Actions menu to the page, we have removed the unnecessary tab navigation. This completes our migration to a menu-based user interface. The content of the page is now the most important aspect of the page, as all the editing and modification features are now out of the way when reading a page. The UI has also been improved in other small ways in this milestone, and we have fixed many bugs with the printable view and the default themes.

Plugin components that expect to be injected with the old, deprecated CacheManager interface will now work again. This means plugins such as the JIRA issues macro will be useable once more. Plugins that use the ContainerManager.getComponent method of obtaining the cache manager will continue to malfunction; there is no plan to support this usage. More commentary about this at CONF-10602.

The Page Ordering feature mostly contains bug fixes. The majority of changes are done in the back end as we prepare to remove extJS library and substitute the page tree widget with a custom component.

Milestone 6 (the current one)

User Interface improvements:

- performance improvements through better clientside-caching on the edit-page
- user menu has been included. The only thing missing to wrap up the UI changes is to move the remaining icon-buttons (PDF, watch pages, etc) to menu items as well. This will be delivered in two weeks.
- New and improved Tree component, that lets you drag more comfortably than before
- Added the ability to cancel a page move by pressing 'esc' after the user starts dragging a page
- Added automatic scrolling functionality. If a user drags a page to the top/bottom of the screen in an attempt to reach a page that is currently outside of the screen, it will automatically scroll up/down.
- Known problem: No indication of subnodes being loaded
- bugfixes

Technical improvements and API changes

- JMX has been added to the list of Confluence capabilities for improved monitoring
- Raising a support case is now possible directly from the admin console. This will attach logs and other relevent system information automatically.
- JQuery now default JS library. We have removed ExtJS as a Javascript-dependency, and will standardize on JQuery in all of Confluence
- Coherence cache lease duration has been set at 10 minutes to support system recovery when threads do not release their leases properly. A warning message will be logged detailing that a lease has timed out to help with tracking down errant lease management.

Known problems

- We found a severe performance problem during our loadtests. Confluence works exactly like before under low load, but once you go above a request per second you will notice a much higher load than acceptable. We will fix that for the final release.
- Upgrade issue for clustered deployments. Don't use this Milestone on a Cluster.
- a few glitches when viewing with IE6&7

Release Notes 2.8-m7 ("Milestone 7")



Do not use this release to upgrade your production systems.

For all production use and testing of Confluence, please use the latest official release.

Confluence 2.8-m7 is a milestone development release for 2.8. This is a public **development release** (DR) leading up to Confluence 2.8. Development releases are a snapshot of our work in progress, allowing our customers and especially plugin-developers to see what we're up to.

Who should upgrade?



Please note the following

- **Development releases are not safe** — Development releases are snapshots of the ongoing Confluence development process. As such:
  - While we try to keep these releases stable, they have not undergone the same degree of testing as a full release.
  - Features in development releases may be incomplete, or may change or be removed before the next full release.
- **No upgrade path** — Because development releases represent work in progress, we **can not** provide a supported upgrade path between development releases, or from any development release to the eventual final release. Thus, it is possible that you will not be able to migrate any data you store in a Confluence development release to a future Confluence release.

Our milestone releases aim to provide plugin developers with an opportunity to see the latest changes in the code.

Each milestone release has passed all our automatic tests, and has been used for one week on our official internal Confluence server. Most of the issues solved have been reviewed too.

However, since our milestones releases are timeboxed (i.e. they get released every two weeks, no matter how far we have come implementing features and bugfixes), there is always a chance that we have new known bugs, which are scheduled to be fixed in the next milestone, or completely new bugs unknown even to us.

Milestone releases may not have not been load- or stress-tested, or maybe they have but some performance problems still persist. So, for example, they might behave well on a small installation but show severe problems when subjected to many users.

Upgrade Procedure

Follow the normal upgrade instructions to upgrade from Confluence 2.7.x to this release. *We strongly recommend that you backup your confluence-home directory and database before upgrading!*

Downloads

All development releases are available from Development Releases on the Atlassian website.

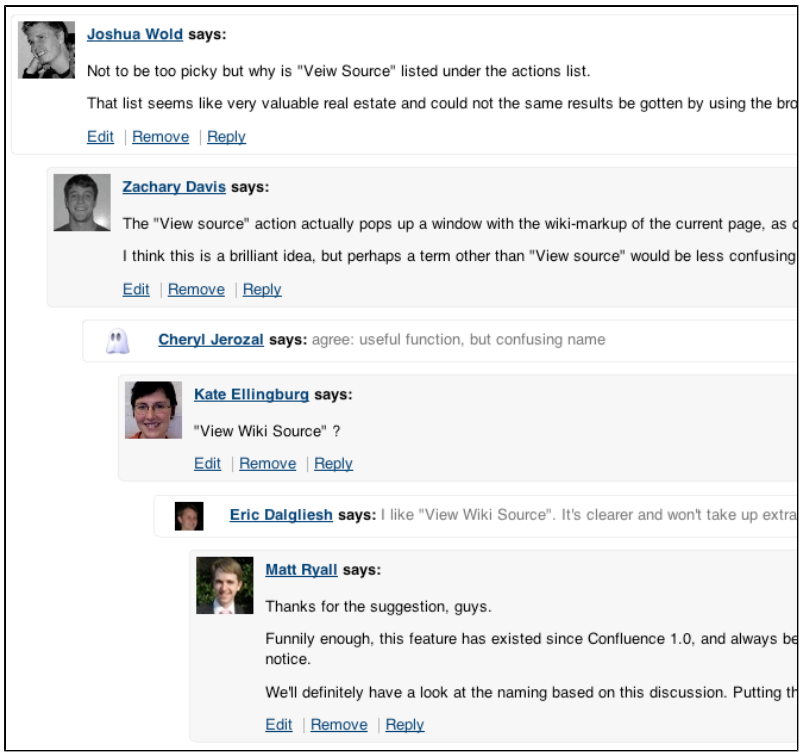
Issues resolved or improved in this release

This page describes the changes you can expect from Confluence 2.8 Milestone 7. M7 is a "small" milestone that delivers some features that just didn't make it into M6, but which can't wait for M8 nor RCs either because we need **your** feedback as soon as possible!

End-user improvements

User Interface improvements:

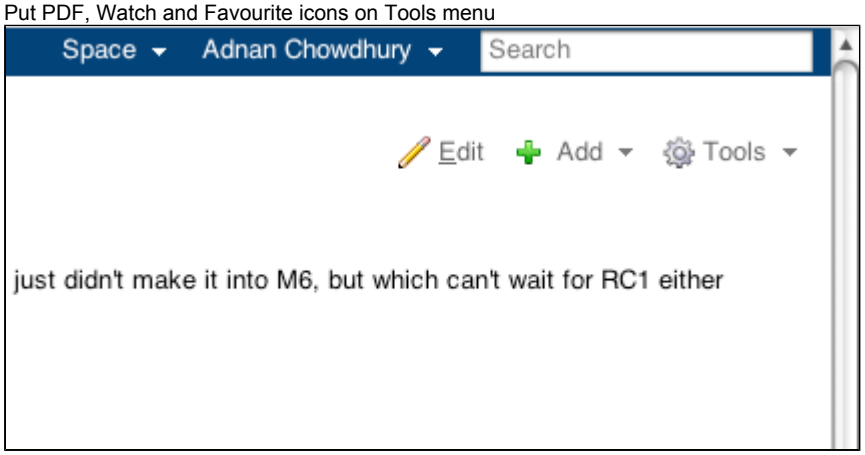
- Matt Ryall's ShipIt project: improved comment threading, with dynamic collapsing



- Edit page title change

The title is not duplicated anymore, you simply edit the title directly

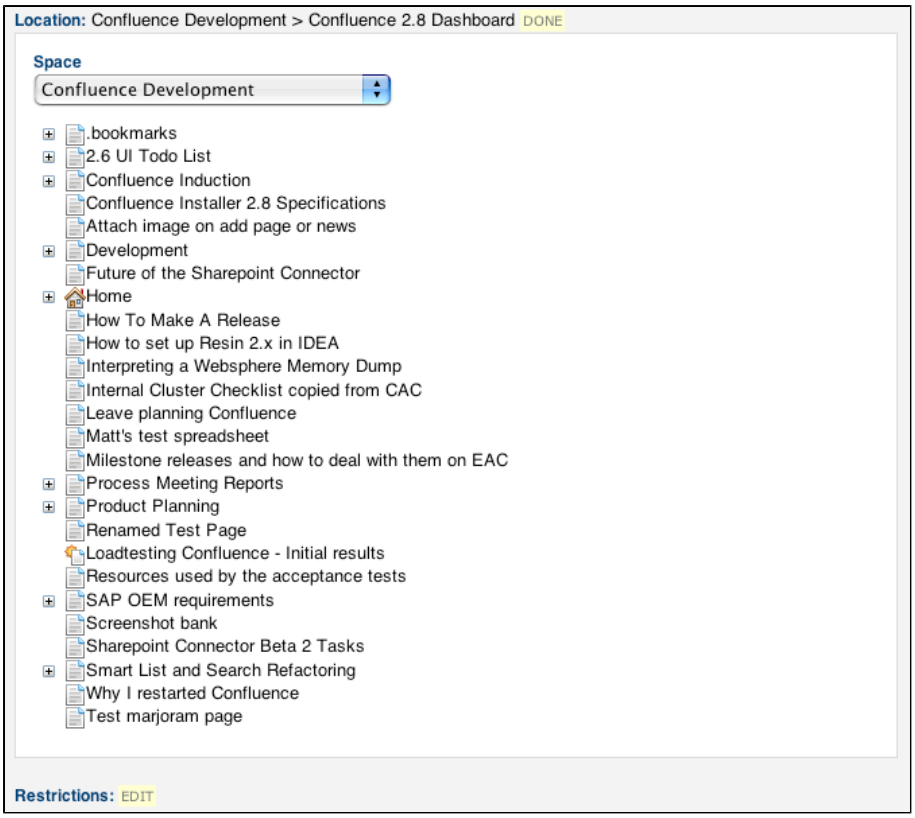
- Put PDF, Watch and Favourite icons on Tools menu



- Replace printable view with an improved print stylesheet for better printing directly from the page
- Improved editor caching
- Fixed editor layout bug in Safari
- Menus appear on hover
- Move command in Tools menu

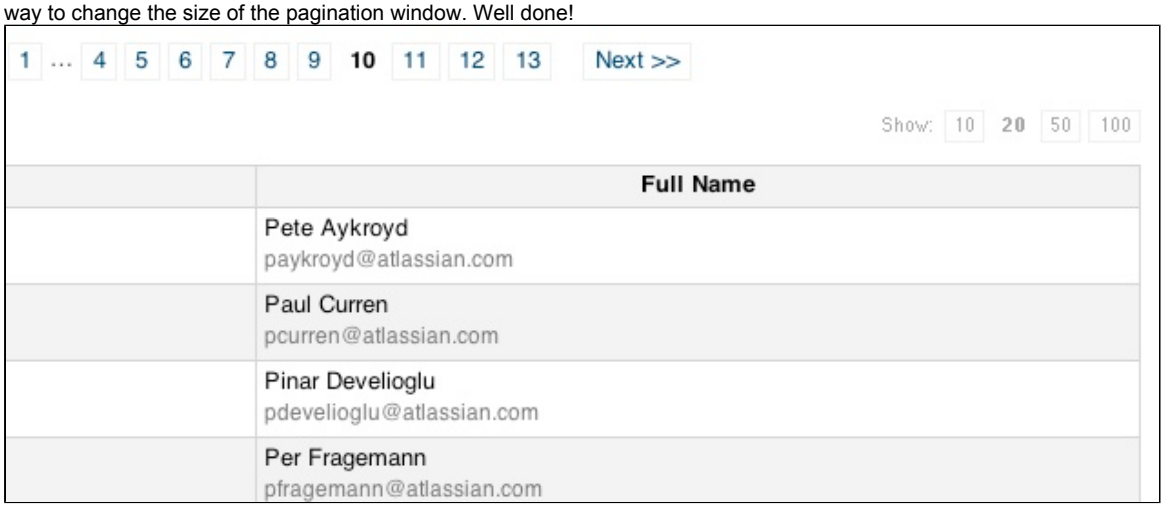
Page Ordering:

- Included the page ordering tree into the edit-page, ironed out some bugs



Misc:

- Our student developer Chris Broadfoot added a few missing links into our user-management section, and incorporated a pretty slick way to change the size of the pagination window. Well done!



API changes

For plugins which configure Confluence's web interface, Condition implementations which depend on GlobalHelper are now deprecated, and should be converted to use the new WebInterfaceContext interface. The easiest way for most plugins to do this is to extend BaseConfluenceCondition rather than the now-deprecated AbstractConfluenceCondition. (Javadoc links to come once this is published.)

Known problems

- We found a severe performance problem during our loadtests. Confluence works exactly like before under low load, but once you go above a request per second you will notice a much higher load than acceptable. We will fix that for the final release.
- Upgrade issue for clustered deployments. Don't use this Milestone on a Cluster.
- a few glitches when viewing with IE6&7, and when using the new page tree in a browser without FireBug installed.

Outlook

We are in bugfix-mode now. While a few icons and CSS-styles will still change, the main priority is now to iron out all our bugs during the next two weeks.

Release Notes 2.8-m9 ("Milestone 9")

Do not use this release to upgrade your production systems.

For all production use and testing of Confluence, please use the latest official release.

Confluence 2.8-m9 is a milestone development release for 2.8. This is a public **development release** (DR) leading up to Confluence 2.8. Development releases are a snapshot of our work in progress, allowing our customers and especially plugin-developers to see what we're up to.

Who should upgrade?

Please note the following

- **Development releases are not safe** — Development releases are snapshots of the ongoing Confluence development process. As such:
  - While we try to keep these releases stable, they have not undergone the same degree of testing as a full release.
  - Features in development releases may be incomplete, or may change or be removed before the next full release.
- **No upgrade path** — Because development releases represent work in progress, we **can not** provide a supported upgrade path between development releases, or from any development release to the eventual final release. Thus, it is possible that you will not be able to migrate any data you store in a Confluence development release to a future Confluence release.

Our milestone releases aim to provide plugin developers with an opportunity to see the latest changes in the code.

Each milestone release has passed all our automatic tests, and has been used for one week on our official internal Confluence server. Most of the issues solved have been reviewed too.

However, since our milestones releases are timeboxed (i.e. they get released every two weeks, no matter how far we have come implementing features and bugfixes), there is always a chance that we have new known bugs, which are scheduled to be fixed in the next milestone, or completely new bugs unknown even to us.

Milestone releases may not have not been load- or stress-tested, or maybe they have but some performance problems still persist. So, for example, they might behave well on a small installation but show severe problems when subjected to many users.

Upgrade Procedure

Follow the normal upgrade instructions to upgrade from Confluence 2.7.x to this release. *We strongly recommend that you backup your confluence-home directory and database before upgrading!*

Downloads

All development releases are available from Development Releases on the Atlassian website.

Issues resolved or improved in this release

All functional changes are now complete. We are only fixing some remaining bugs now.

Known problems

- We found a major performance problem during our loadtests. Confluence works exactly like before under low load, but once you go above a request per second you will notice a much load than acceptable. We will fix that for the final release.
- a few glitches when viewing with IE6&7, and when using the new page tree. Will be fixed for the final version

Outlook

We are in bugfix-mode now. While a few icons and CSS-styles will still change, the main priority is now to iron out all our bugs during the next one or two weeks.

Release Notes 2.9-m2 ("Milestone 2")

Do not use this release to upgrade your production systems.

For all production use and testing of Confluence, please use the latest official release.
This release is a milestone development release for 2.9. This is a public **development release** (DR) leading up to Confluence 2.9. Development releases are a snapshot of our work in progress, allowing our customers and especially plugin-developers to see what we're up

to.

Who should upgrade?



Please note the following

- **Development releases are not safe** — Development releases are snapshots of the ongoing Confluence development process. As such:
  - While we try to keep these releases stable, they have not undergone the same degree of testing as a full release.
  - Features in development releases may be incomplete, or may change or be removed before the next full release.
- **No upgrade path** — Because development releases represent work in progress, we **can not** provide a supported upgrade path between development releases, or from any development release to the eventual final release. Thus, it is possible that you will not be able to migrate any data you store in a Confluence development release to a future Confluence release.

Our milestone releases aim to provide plugin developers with an opportunity to see the latest changes in the code.

Each milestone release has passed all our automatic tests, and has been used for one week on our official internal Confluence server. Most of the issues solved have been reviewed too, and usually milestone releases even have been load- and performance-tested for a while.

However, since our milestones releases are timeboxed (i.e. they get released every two weeks, no matter how far we have come implementing features and bugfixes), there is always a chance that we have new known bugs, which are scheduled to be fixed in the next milestone, or completely new bugs unknown even to us.

Additionally, our performance-testing and compatibility testing for databases and application servers is not done to the full extent. So, for example, a milestone release might behave well on a small installation but show severe problems when subjected to many users.

Upgrade Procedure

Follow the normal upgrade instructions to upgrade from Confluence 2.8.x to this release. *We strongly recommend that you backup your confluence-home directory and database before upgrading!*

Downloads

All development releases are available from Development Releases on the Atlassian website.

Issues resolved or improved in this release

Page Ordering: We have been improving the page-tree a bit, fitting in some work that was intended for Confluence 2.8 but didn't make it. The tree now has a "revert to alphabetical mode" button which will revert manually ordered nodes back to the original alphanumeric ordering. Additionally we added one small improvement which will help to battle spam: administrators will see two additional buttons on the page tree that enable easy deletion and preview of pages. We are not entirely sure yet whether (and how) these two buttons will ship, so don't raise improvement issues for those two buttons yet. (IF they ship with 2.9, then we will put in some more polish of course)

Unfortunately there are a few known (uncritical) bugs related to the revert-to-alphabetic-mode in edge-cases, which will be addressed by M3.

Search: The Search has improved internally (the ranking should be a bit more relevant than before), and also what is being searched has changed: mail and personal spaces are searched by default, and using the space key in the search boosts hits within that space.

Known problems

Some small issues, mainly with the UI, but nothing serious.

Release Notes 2.9-m3 ("Milestone 3")



Do not use this release to upgrade your production systems.

For all production use and testing of Confluence, please use the latest official release.
This release is a milestone development release for 2.9. This is a public **development release** (DR) leading up to Confluence 2.9.
Development releases are a snapshot of our work in progress, allowing our customers and especially plugin-developers to see what we're up to.

Who should upgrade?

Please note the following

- **Development releases are not safe** — Development releases are snapshots of the ongoing Confluence development process. As such:
  - While we try to keep these releases stable, they have not undergone the same degree of testing as a full release.
  - Features in development releases may be incomplete, or may change or be removed before the next full release.
- **No upgrade path** — Because development releases represent work in progress, we **can not** provide a supported upgrade path between development releases, or from any development release to the eventual final release. Thus, it is possible that you will not be able to migrate any data you store in a Confluence development release to a future Confluence release.

Our milestone releases aim to provide plugin developers with an opportunity to see the latest changes in the code.

Each milestone release has passed all our automatic tests, and has been used for one week on our official internal Confluence server. Most of the issues solved have been reviewed too, and usually milestone releases even have been load- and performance-tested for a while.

However, since our milestones releases are timeboxed (i.e. they get released every two weeks, no matter how far we have come implementing features and bugfixes), there is always a chance that we have new known bugs, which are scheduled to be fixed in the next milestone, or completely new bugs unknown even to us.

Additionally, our performance-testing and compatibility testing for databases and application servers is not done to the full extent. So, for example, a milestone release might behave well on a small installation but show severe problems when subjected to many users.

Upgrade Procedure

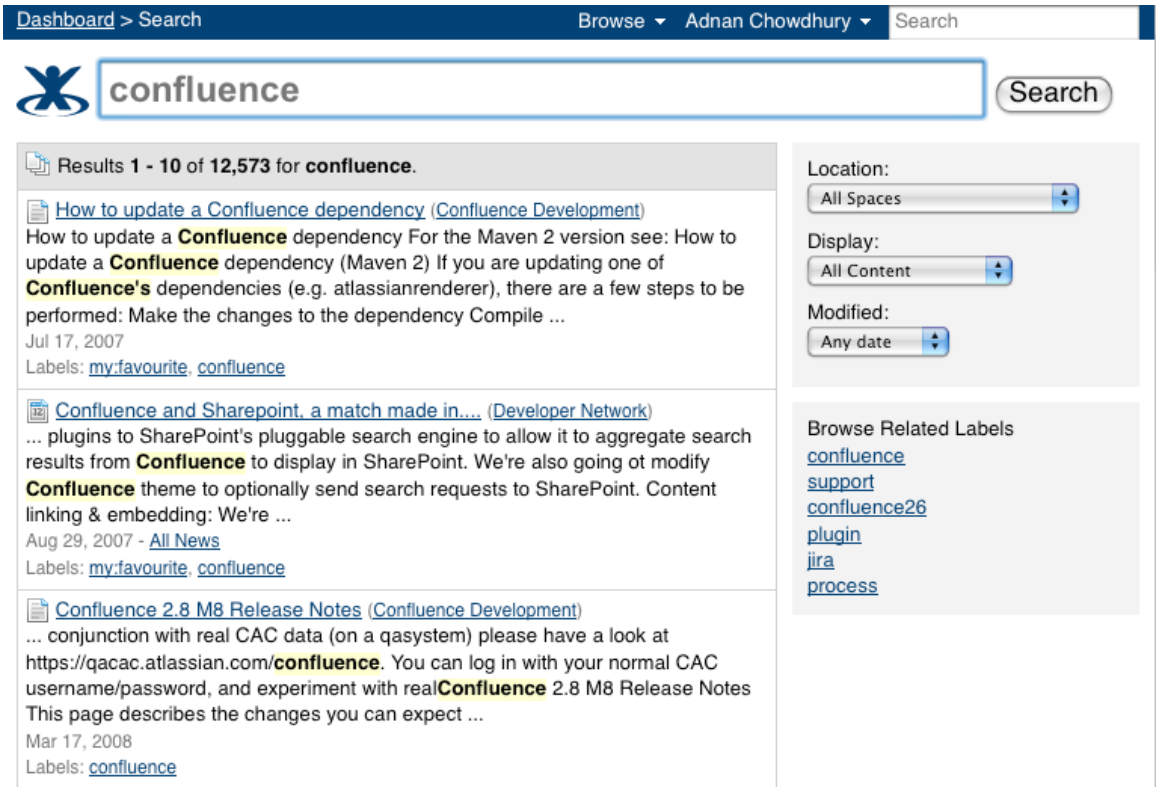
Follow the normal upgrade instructions to upgrade from Confluence 2.8.x to this release. *We strongly recommend that you backup your confluence-home directory and database before upgrading!*

Downloads

All development releases are available from Development Releases on the Atlassian website.

Issues resolved or improved in this release

Milestone 3 is all about the UI: It features major parts of the new search UI, and we have added two more menu items to the "Browse" menu.



Some effort was spent on making the page tree more stable, and we have spent more time on internal code improvements (which did however not ship in this milestone)

Known problems

Several small bugs. And a slowdown of search performance by 50%, which will be addressed in a later milestone release.

Release Notes 2.9-m5 ("Milestone 5")



Do not use this release to upgrade your production systems.

For all production use and testing of Confluence, please use the latest official release.
This release is a milestone development release for 2.9. This is a public **development release** (DR) leading up to Confluence 2.9. Development releases are a snapshot of our work in progress, allowing our customers and especially plugin-developers to see what we're up to.

Who should upgrade?



Please note the following

- **Development releases are not safe** — Development releases are snapshots of the ongoing Confluence development process. As such:
  - While we try to keep these releases stable, they have not undergone the same degree of testing as a full release.
  - Features in development releases may be incomplete, or may change or be removed before the next full release.
- **No upgrade path** — Because development releases represent work in progress, we **can not** provide a supported upgrade path between development releases, or from any development release to the eventual final release. Thus, it is possible that you will not be able to migrate any data you store in a Confluence development release to a future Confluence release.

Our milestone releases aim to provide plugin developers with an opportunity to see the latest changes in the code.

Each milestone release has passed all our automatic tests, and has been used for one week on our official internal Confluence server. Most of the issues solved have been reviewed too, and usually milestone releases even have been load- and performance-tested for a while.

However, since our milestones releases are timeboxed (i.e. they get released every two weeks, no matter how far we have come implementing features and bugfixes), there is always a chance that we have new known bugs, which are scheduled to be fixed in the next milestone, or completely new bugs unknown even to us.

Additionally, our performance-testing and compatibility testing for databases and application servers is not done to the full extent. So, for example, a milestone release might behave well on a small installation but show severe problems when subjected to many users.

Upgrade Procedure

Follow the normal upgrade instructions to upgrade from Confluence 2.8.x to this release. *We strongly recommend that you backup your confluence-home directory and database before upgrading!*

Downloads

All development releases are available from Development Releases on the Atlassian website.

Issues resolved or improved in this release

We are rapidly approaching the end of the 2.9 release cycle. There will be a really small M6 later this week, and then that's it, we are going into bugfix and compatibility-testing mode, so we can release on the 29th of July.

Plugins

We have done a significant amount of work in the area of plugin loading (and the plugin classloaders) to improve performance. There are no API changes, but it is very possible that these changes may expose bugs. While we tried to remain backwards compatible, it is **crucial** that you test for compatibility of your plugins.

Editor

The **preview** in the editor is now capable of rendering content exactly as it appears on the page. However, to do this, we have had to disable interaction with the preview. You will no longer be able to follow links or interact with the preview in any other way.

When **drafts are saved** automatically to the server, this will be shown to you at the top of the editor.

Engine room

Whisper the words "Action class hierarchy" anywhere close to a Confluence developer and he or she will cringe. Not anymore as of this Milestone. Plenty of inheritance trees have been cut, and miraculously Confluence still works. This will enable us to work faster and more efficiently in the future.

Page tree and UI

We had a few nasty UI bugs recently where plugin JavaScript code would interfere with our menus and other JavaScript code, and we had a few Pagetree bugs. The most important ones have been fixed in this release.

Discovery Team

Author auto-complete field added to the search result screen to allow filtering searches by contributor (author or editor).

The 'new search' (or reset search) link on the search screen filter left the building somewhere around milestone 3. This has now been re-instated.

A couple of 2.9 specific bug fixes are also included.

Known problems

Several small bugs, specifically around the new author-search.

Release Notes 2.9-rc1 ("Release Candidate 1")



Do not use this release to upgrade your production systems.

For all production use and testing of Confluence, please use the latest official release.
This release is a milestone development release for 2.9. This is a public **development release** (DR) leading up to Confluence 2.9. Development releases are a snapshot of our work in progress, allowing our customers and especially plugin-developers to see what we're up to.

Who should upgrade?



Please note the following

- **Development releases are not safe** — Development releases are snapshots of the ongoing Confluence development process. As such:
  - While we try to keep these releases stable, they have not undergone the same degree of testing as a full release.
  - Features in development releases may be incomplete, or may change or be removed before the next full release.
- **No upgrade path** — Because development releases represent work in progress, we **can not** provide a supported upgrade path between development releases, or from any development release to the eventual final release. Thus, it is possible that you will not be able to migrate any data you store in a Confluence development release to a future Confluence release.

Our milestone releases aim to provide plugin developers with an opportunity to see the latest changes in the code.

Each milestone release has passed all our automatic tests, and has been used for one week on our official internal Confluence server. Most of the issues solved have been reviewed too, and usually milestone releases even have been load- and performance-tested for a while.

However, since our milestones releases are timeboxed (i.e. they get released every two weeks, no matter how far we have come implementing features and bugfixes), there is always a chance that we have new known bugs, which are scheduled to be fixed in the next milestone, or completely new bugs unknown even to us.

Additionally, our performance-testing and compatibility testing for databases and application servers is not done to the full extent. So, for example, a milestone release might behave well on a small installation but show severe problems when subjected to many users.

Upgrade Procedure

Follow the normal upgrade instructions to upgrade from Confluence 2.8.x to this release. *We strongly recommend that you backup your confluence-home directory and database before upgrading!*

Downloads

All development releases are available from Development Releases on the Atlassian website.

Issues resolved or improved in this release

We are now at the end of the 2.9 release cycle.

We have released our first release candidate RC1 today, you can see it on our public Confluence installation at http://confluence.atlassian.com. We are aware of one remaining major bug that we want to fix (JIRA-issues macro not working in preview mode), but apart from that and some minor fixes we will not change the code anymore.

While this is not the final release, and the official release notes will be unveiled next week, you should definitly have a look now if everything still works fine with your plugins, if you haven't done so yet. We made some changes to the plugin subsystem and to the action class hierarchy (please refer to the other milestone release notes), and we are aware of some plugins (such as the Gliffy plugin) which had to get changed a bit to be compatible with 2.9.

The RC1 is functionally almost equivalent to M5/M6 which was announced in early July, but has a big amount of additional bugfixes, check out our roadmap in the JIRA project at http://jira.atlassian.com/browse/CONF?report=com.atlassian.jira.plugin.system.project:roadmap-panel for more details.

Known problems

Jira-issues macro does not work in preview mode.

Release Notes 2.10-m1 ("Milestone 1")

Do not use this release to upgrade your production systems.

For all production use and testing of Confluence, please use the latest official release.
This release is a milestone development release for 2.10 ("Two Ten"). This is a public **development release** (DR) leading up to Confluence 2.10. Development releases are a snapshot of our work in progress, allowing our customers and especially plugin-developers to see what we're up to.

Who should upgrade?

Please note the following

- **Development releases are not safe** — Development releases are snapshots of the ongoing Confluence development process. As such:
  - While we try to keep these releases stable, they have not undergone the same degree of testing as a full release.
  - Features in development releases may be incomplete, or may change or be removed before the next full release.
- **No upgrade path** — Because development releases represent work in progress, we **can not** provide a supported upgrade path between development releases, or from any development release to the eventual final release. Thus, it is possible that you will not be able to migrate any data you store in a Confluence development release to a future Confluence release.

Our milestone releases aim to provide plugin developers with an opportunity to see the latest changes in the code.

Each milestone release has passed all our automatic tests, and has been used for one week on our official internal Confluence server. Most of the issues solved have been reviewed too, and usually milestone releases even have been load- and performance-tested for a while.

However, since our milestones releases are timeboxed (i.e. they get released every two weeks, no matter how far we have come implementing features and bugfixes), there is always a chance that we have new known bugs, which are scheduled to be fixed in the next milestone, or completely new bugs unknown even to us.

Additionally, our performance-testing and compatibility testing for databases and application servers is not done to the full extent. So, for example, a milestone release might behave well on a small installation but show severe problems when subjected to many users.

Upgrade Procedure

Follow the normal upgrade instructions to upgrade from Confluence 2.9.x to this release. *We strongly recommend that you backup your confluence-home directory and database before upgrading!*

Downloads

All development releases are available from Development Releases on the Atlassian website.

Issues resolved or improved in this release

It's all about small improvements this time.

During thhe last weeks we have fixed a whole bunch of bugs, worked on major functionality, and done a lot of planning. Therefore the M1 is a rather small release - it contains the bugfixes, but no major work yet. We have developed a new feature and two great smaller improvements which are almost done, but didn't make it into this Milestone, so expect a lot more in M2 in two weeks.

But wait, there is one small new feature that might have a big impact on EAC. Jens Schumacher, developer on Confluence Hosted, has been working on a top secret mission, and his new feature means you now can store CSS per space and for the whole installation.

Inside Atlassian we now have a small competition going with developers competing for the nicest/coolest/weirdest CSS-based design to be delivered within the next two weeks.

The competition only started yesterday, but there are some funky screenshots available already Obviously Confluence still has areas that are completely hardwired and not stylable yet, but we are working on this as we go.









Release Notes 2.10-m2 ("Milestone 2")

Do not use this release to upgrade your production systems.

For all production use and testing of Confluence, please use the latest official release.
This release is a milestone development release for 2.10 ("Two Ten"). This is a public **development release** (DR) leading up to Confluence 2.10. Development releases are a snapshot of our work in progress, allowing our customers and especially plugin-developers to see what we're up to.

Who should upgrade?

Please note the following

- **Development releases are not safe** — Development releases are snapshots of the ongoing Confluence development process. As such:
  - While we try to keep these releases stable, they have not undergone the same degree of testing as a full release.
  - Features in development releases may be incomplete, or may change or be removed before the next full release.
- **No upgrade path** — Because development releases represent work in progress, we **can not** provide a supported upgrade path between development releases, or from any development release to the eventual final release. Thus, it is possible that you will not be able to migrate any data you store in a Confluence development release to a future Confluence release.

Our milestone releases aim to provide plugin developers with an opportunity to see the latest changes in the code.

Each milestone release has passed all our automatic tests, and has been used for one week on our official internal Confluence server. Most of the issues solved have been reviewed too, and usually milestone releases even have been load- and performance-tested for a while.

However, since our milestones releases are timeboxed (i.e. they get released every two weeks, no matter how far we have come implementing features and bugfixes), there is always a chance that we have new known bugs, which are scheduled to be fixed in the next milestone, or completely new bugs unknown even to us.

Additionally, our performance-testing and compatibility testing for databases and application servers is not done to the full extent. So, for example, a milestone release might behave well on a small installation but show severe problems when subjected to many users.

Upgrade Procedure

Follow the normal upgrade instructions to upgrade from Confluence 2.9.x to this release. *We strongly recommend that you backup your confluence-home directory and database before upgrading!*

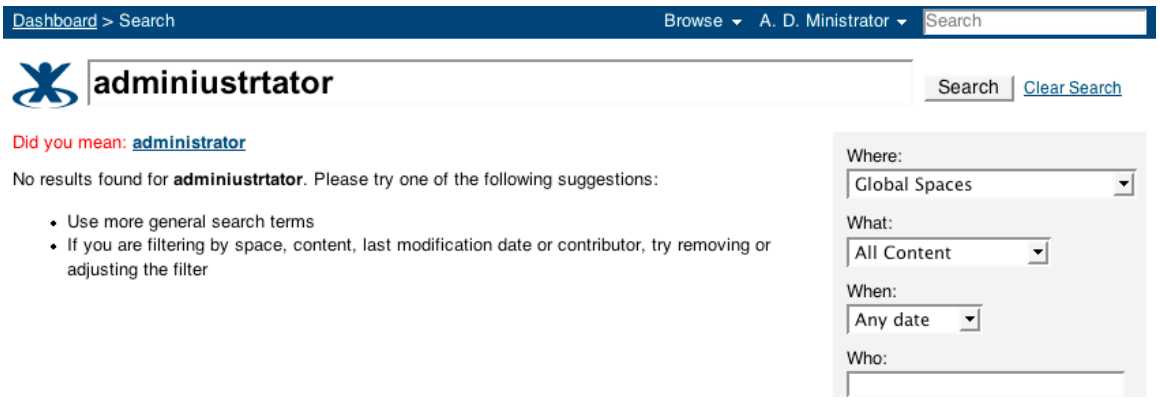
Downloads

All development releases are available from Development Releases on the Atlassian website.

Issues resolved or improved in this release

"Did you mean?"

Dave Loeng's auto-suggest feature is the first major 20%-project to make it into Confluence. Try it out by mistyping a search, and you will find a link suggesting more relevant searches. **This is an actual example of me typing way too fast:**

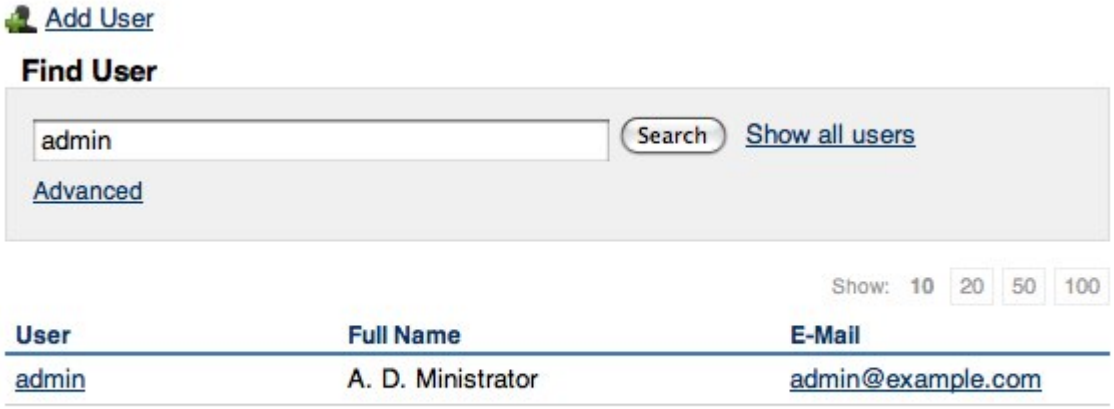


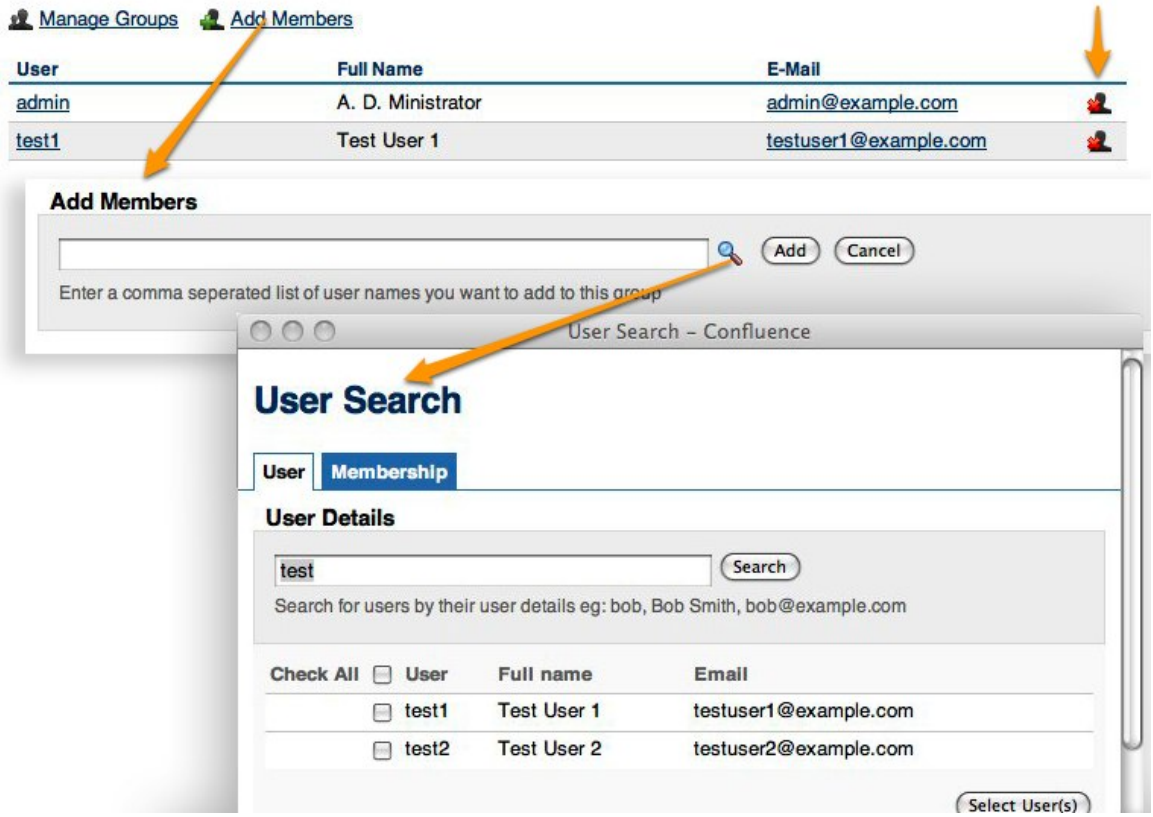
Try misspelling someone's difficult surname, like say, 'fraggemnan'.

User management

Confluence Hosted developer Jens Schumacher has been very busy hacking on Confluence again, this time he improved the user-management a lot - so much more convenient than before.

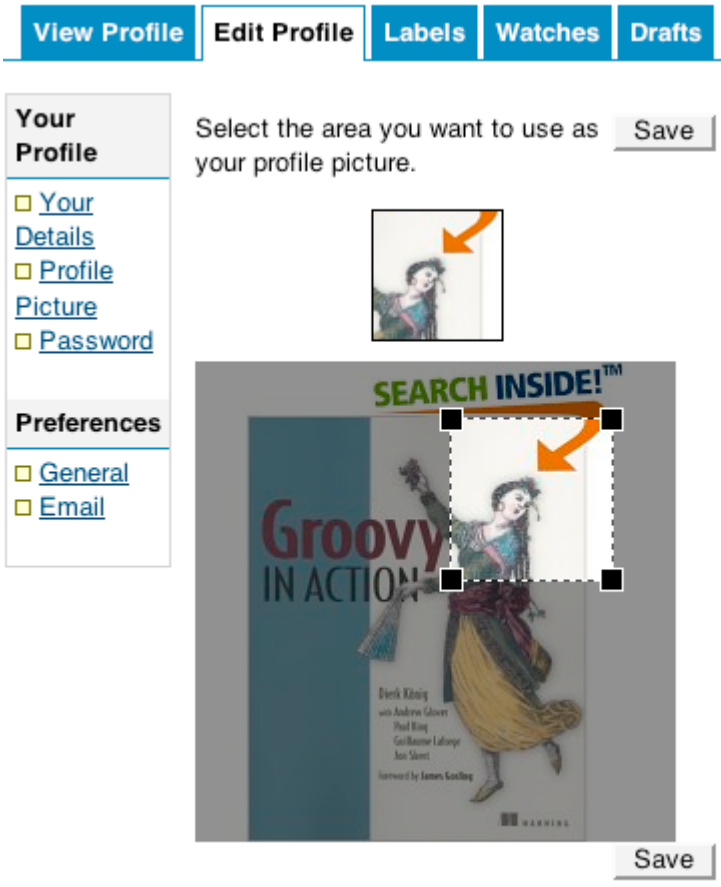
- Improved Search
- Adding and removing users when viewing a group
- New table styles
- Improved User Picker





Avatar cropping and deletion

Another, smaller 20%-project also made it into this release: Charles Miller and Dmitry baranovskiy enabled Confluence to allow cropping of uploaded avatar-images, and while they were at it the highly desired "delete images you never intended to upload in the first place" issue got solved too in order to make another 23 voters happy.



Backend changes

The Engine Room team has implemented several important backend changes in this milestone:

- Integrated the first version of Plugins 2.0, including an initial migration the Confluence plugin repository to Plugins 2. An updated version of this plugin and a converted dynamictasklist are planned for the next milestone.
- Trusted authentication and other Seraph-based authentication methods are now available for calls to the Confluence RPC methods (CONF-8680). This makes it practical to write Confluence gadgets which use the remote API to retrieve data. Tom Davies is using this to implement his Crucible Confluence review plugin.

Release Notes 2.10-m3 ("Milestone 3")

Do not use this release to upgrade your production systems.

For all production use and testing of Confluence, please use the latest official release.
This release is a milestone development release for 2.10 ("Two Ten"). This is a public **development release** (DR) leading up to Confluence 2.10. Development releases are a snapshot of our work in progress, allowing our customers and especially plugin-developers to see what we're up to.

Who should upgrade?

Please note the following

- **Development releases are not safe** — Development releases are snapshots of the ongoing Confluence development process. As such:
  - While we try to keep these releases stable, they have not undergone the same degree of testing as a full release.
  - Features in development releases may be incomplete, or may change or be removed before the next full release.
- **No upgrade path** — Because development releases represent work in progress, we **can not** provide a supported upgrade path between development releases, or from any development release to the eventual final release. Thus, it is possible that you will not be able to migrate any data you store in a Confluence development release to a future Confluence release.

Our milestone releases aim to provide plugin developers with an opportunity to see the latest changes in the code.

Each milestone release has passed all our automatic tests, and has been used for one week on our official internal Confluence server. Most of the issues solved have been reviewed too, and usually milestone releases even have been load- and performance-tested for a while.

However, since our milestones releases are timeboxed (i.e. they get released every two weeks, no matter how far we have come implementing features and bugfixes), there is always a chance that we have new known bugs, which are scheduled to be fixed in the next milestone, or completely new bugs unknown even to us.

Additionally, our performance-testing and compatibility testing for databases and application servers is not done to the full extent. So, for example, a milestone release might behave well on a small installation but show severe problems when subjected to many users.

Upgrade Procedure

Follow the normal upgrade instructions to upgrade from Confluence 2.9.x to this release. *We strongly recommend that you backup your confluence-home directory and database before upgrading!*

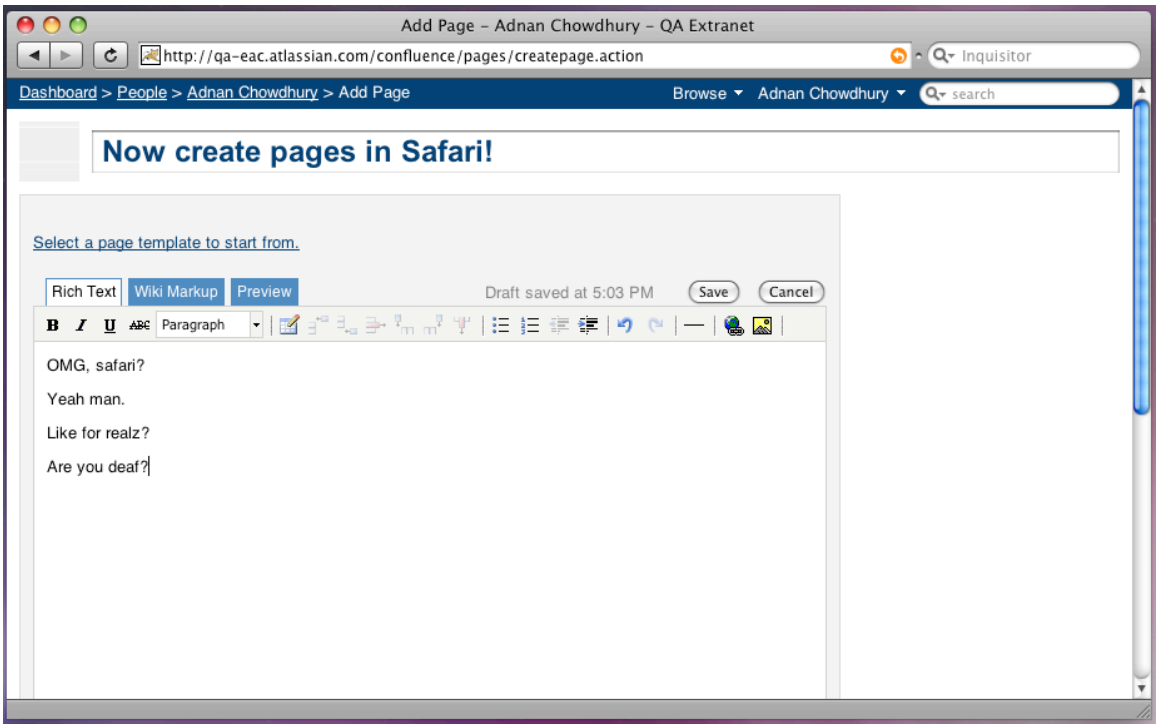
Downloads

All development releases are available from Development Releases on the Atlassian website.

Issues resolved or improved in this release

TinyMCE 3 is in

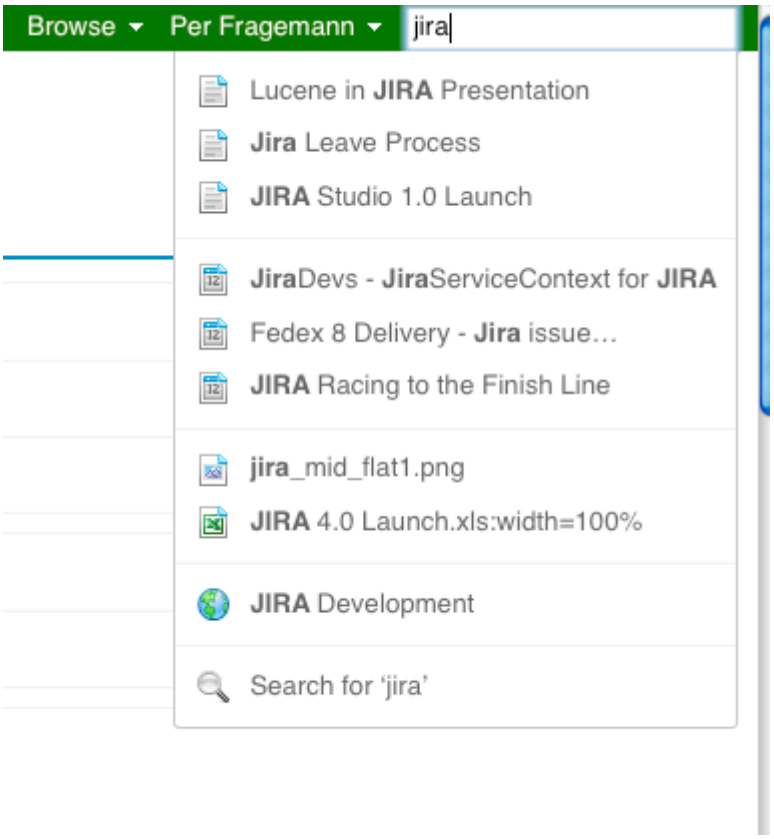
This milestone contains the long-awaited TinyMCE 3 upgrade. It comes at a cost (we had to disable the emoticons, special characters, colour-pickers and fullscreen mode) to make it into M3, but they will be back in M4, and this gives us a much more stable Rich Text Editor two weeks earlier than expected. Please note: we have **not** yet started working on roundtripping bugs (e.g. stuff that breaks when switching from RTE to WikiMarkup and back). We will start attacking bugs those towards the end of the release cycle and during the 3.0 cycle. But the improvements in the RTE itself are extremely valuable already, including **support for safari**. If you encounter problems please don't hesitate to ask or raise issues. We want to make this a really solid experience.



Oh, and the new RTE starts up much faster too! Kudos to the relentless work of the Writer-Team (especially David T and Agnes), who have been working for many hours over the past months to get us here.

Quick Navigation

- Matches titles as you type
- Shows 3 pages, 3 news items, 2 attachments, 3 people, 2 spaces. If no titles are matched the category isn't shown.
- People's profile photos are shown
- Use your keyboard to navigate
- Now with % more rounded corners



Clear Search

- Improved search result page and results
  - Lighter 'Showing ...' line
  - Better fit with a grid, softening of filter panel corners and better spacing
  - Emphasis of the titles
  - Clearer, simpler result format
  - Real profile pictures of real people
- Clear filter instead of clear search
- Search field at the bottom of the page



Small improvements

Per's first 20% project displays a lock-icon next to documents that have restricted access.



Enjoy!

Notes on deploying to your test-instance

With the TinyMCE upgrade, we have renamed the TinyMCE plugin key to match the Confluence version. Hence, you may find that there are two TinyMCE plugins installed which can cause errors when editing (e.g. the pop-ups not working). If so, you will need to disable both and then re-enable the 2.10-m3 version of TinyMCE. This will be fixed for the final 2.10 release with an upgrade task to uninstall the old plugin.

Release Notes 2.10-m4 ("Milestone 4")



Do not use this release to upgrade your production systems.

For all production use and testing of Confluence, please use the latest official release.
This release is a milestone development release for 2.10 ("Two Ten"). This is a public **development release** (DR) leading up to Confluence 2.10. Development releases are a snapshot of our work in progress, allowing our customers and especially plugin-developers to see what we're up to.

Who should upgrade?



Please note the following

- **Development releases are not safe** — Development releases are snapshots of the ongoing Confluence development process. As such:
  - While we try to keep these releases stable, they have not undergone the same degree of testing as a full release.
  - Features in development releases may be incomplete, or may change or be removed before the next full release.
- **No upgrade path** — Because development releases represent work in progress, we **can not** provide a supported upgrade path between development releases, or from any development release to the eventual final release. Thus, it is possible that you will not be able to migrate any data you store in a Confluence development release to a future Confluence release.

Our milestone releases aim to provide plugin developers with an opportunity to see the latest changes in the code.

Each milestone release has passed all our automatic tests, and has been used for one week on our official internal Confluence server. Most of the issues solved have been reviewed too, and usually milestone releases even have been load- and performance-tested for a while.

However, since our milestones releases are timeboxed (i.e. they get released every two weeks, no matter how far we have come implementing features and bugfixes), there is always a chance that we have new known bugs, which are scheduled to be fixed in the next milestone, or completely new bugs unknown even to us.

Additionally, our performance-testing and compatibility testing for databases and application servers is not done to the full extent. So, for example, a milestone release might behave well on a small installation but show severe problems when subjected to many users.

Upgrade Procedure

Follow the normal upgrade instructions to upgrade from Confluence 2.9.x to this release. *We strongly recommend that you backup your confluence-home directory and database before upgrading!*

Downloads

All development releases are available from Development Releases on the Atlassian website.

Issues resolved or improved in this release

TinyMCE upgrade complete

- Colour Text, Insert Emoticons and Custom Character are now back in the editor!
- Fullscreen mode is also back and even better. It no longer pops up in a new window, but enlargens the editor in the current page you are editing.

Quick Navigation polish and performance improvements

- performance significantly better than m3
- can handle concurrent requests up to 30 users without breaking a sweat (on medium datasets) - this should address the timeout issues people have been having

Plugins

- New build of JIRA Issues macros. Bug-fixes, but new feature: all RSS and Email version of the macro will be static, and not require javascript or Ajax.
- Content Filtering Macros: contentbylabel and blogposts have a new, SearchAPI v2 backend.
  - All existing parameters should continue to work
  - New, standard set of parameters for various kinds of filtering
- Content Filtering macros: new {recently-changed} macro, which is a reimplementation of {recently-updated}. All existing recently-updated params should work, with the addition of the new standard params listed above. Will eventually replace {recently-updted}.

- Added an optional whitelist for the RSS macro and the HTML-include macro.

Release Notes 2.10-m5 ("Milestone 5")



Do not use this release to upgrade your production systems.

For all production use and testing of Confluence, please use the latest official release.
This release is a milestone development release for 2.10 ("Two Ten"). This is a public **development release** (DR) leading up to Confluence 2.10. Development releases are a snapshot of our work in progress, allowing our customers and especially plugin-developers to see what we're up to.

Who should upgrade?



Please note the following

- **Development releases are not safe** — Development releases are snapshots of the ongoing Confluence development process. As such:
  - While we try to keep these releases stable, they have not undergone the same degree of testing as a full release.
  - Features in development releases may be incomplete, or may change or be removed before the next full release.
- **No upgrade path** — Because development releases represent work in progress, we **can not** provide a supported upgrade path between development releases, or from any development release to the eventual final release. Thus, it is possible that you will not be able to migrate any data you store in a Confluence development release to a future Confluence release.

Our milestone releases aim to provide plugin developers with an opportunity to see the latest changes in the code.

Each milestone release has passed all our automatic tests, and has been used for one week on our official internal Confluence server. Most of the issues solved have been reviewed too, and usually milestone releases even have been load- and performance-tested for a while.

However, since our milestones releases are timeboxed (i.e. they get released every two weeks, no matter how far we have come implementing features and bugfixes), there is always a chance that we have new known bugs, which are scheduled to be fixed in the next milestone, or completely new bugs unknown even to us.

Additionally, our performance-testing and compatibility testing for databases and application servers is not done to the full extent. So, for example, a milestone release might behave well on a small installation but show severe problems when subjected to many users.

Upgrade Procedure

Follow the normal upgrade instructions to upgrade from Confluence 2.9.x to this release. *We strongly recommend that you backup your confluence-home directory and database before upgrading!*

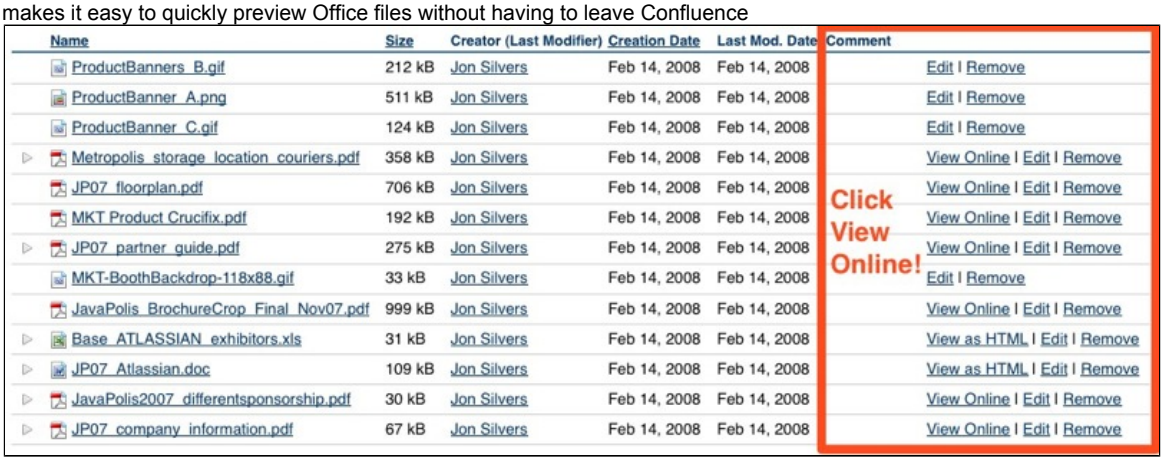
Downloads

All development releases are available from Development Releases on the Atlassian website.

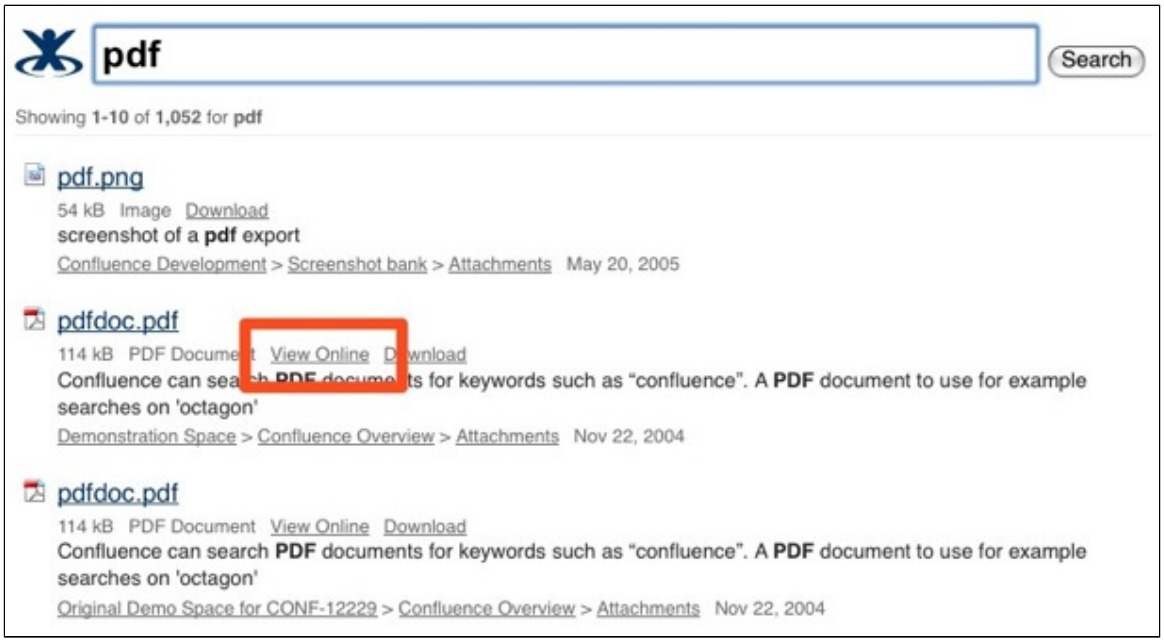
Issues resolved or improved in this release

Office Connector

- View Online" or "View as HTML" link for supported file types on the Attachments screen, Attachments macro, and Search Results. This makes it easy to quickly preview Office files without having to leave Confluence



Name	Size	Creator (Last Modifier)	Creation Date	Last Mod. Date	Comment	
ProductBanners_B.gif	212 kB	Jon Silvers	Feb 14, 2008	Feb 14, 2008		Edit | Remove
ProductBanner_A.png	511 kB	Jon Silvers	Feb 14, 2008	Feb 14, 2008		Edit | Remove
ProductBanner_C.gif	124 kB	Jon Silvers	Feb 14, 2008	Feb 14, 2008		Edit | Remove
Metropolis_storage_location_couriers.pdf	358 kB	Jon Silvers	Feb 14, 2008	Feb 14, 2008		View Online | Edit | Remove
JP07_floorplan.pdf	706 kB	Jon Silvers	Feb 14, 2008	Feb 14, 2008		View Online | Edit | Remove
MKT Product Crucifix.pdf	192 kB	Jon Silvers	Feb 14, 2008	Feb 14, 2008		View Online | Edit | Remove
JP07_partner_guide.pdf	275 kB	Jon Silvers	Feb 14, 2008	Feb 14, 2008		View Online | Edit | Remove
MKT-BoothBackdrop-118x88.gif	33 kB	Jon Silvers	Feb 14, 2008	Feb 14, 2008		Edit | Remove
JavaPolis_BrochureCrop_Final_Nov07.pdf	999 kB	Jon Silvers	Feb 14, 2008	Feb 14, 2008		View Online | Edit | Remove
Base_ATLASSIAN_exhibitors.xls	31 kB	Jon Silvers	Feb 14, 2008	Feb 14, 2008		View as HTML | Edit | Remove
JP07_Atlassian.doc	109 kB	Jon Silvers	Feb 14, 2008	Feb 14, 2008		View as HTML | Edit | Remove
JavaPolis2007_differentsponsorship.pdf	30 kB	Jon Silvers	Feb 14, 2008	Feb 14, 2008		View Online | Edit | Remove
JP07_company_information.pdf	67 kB	Jon Silvers	Feb 14, 2008	Feb 14, 2008		View Online | Edit | Remove



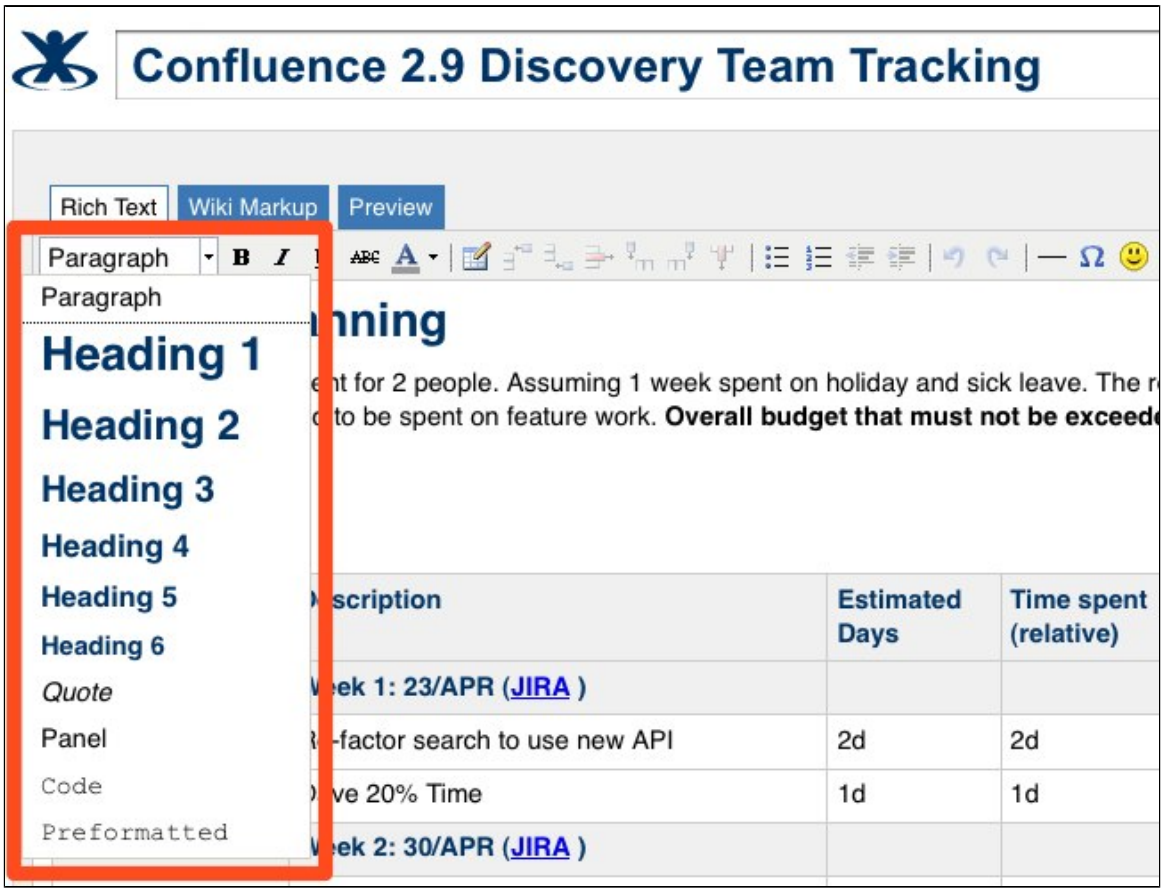
- New look for the PPT and PDF flash viewer with a download button, and an improved fullscreen.



- Conversion Queue for managing performance when you may have a lot of users using the Office Connector.

Rich Text Editor

- Tab & Shift+Tab in lists are now shortcuts for indent/outdent
- Lots of bug fixes (especially around table editing)
- Round tripping bugs around links were fixed (shortcut links, using colors & images in links)
- The format dropdown is now styled properly

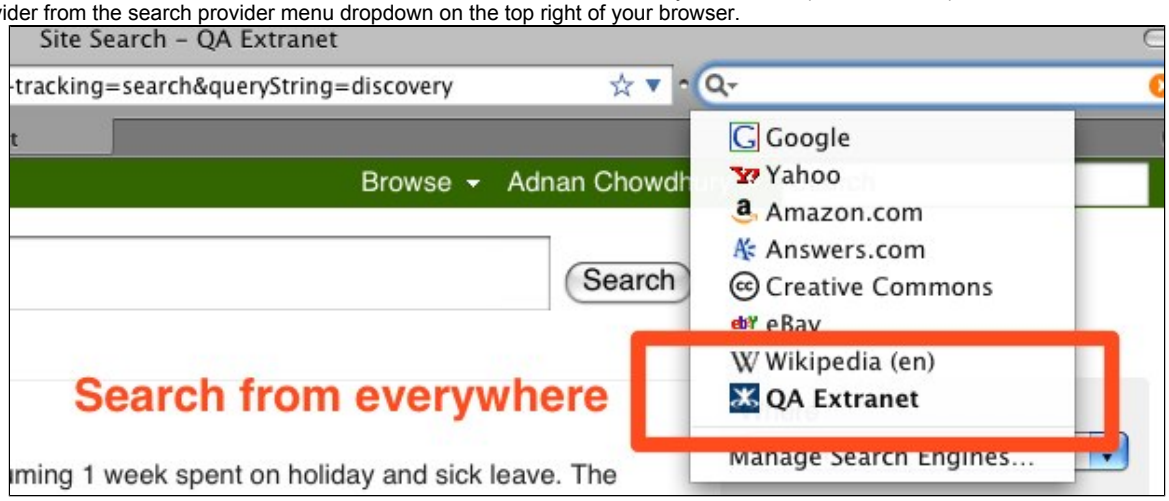


Quick nav improvements

- Is now at least twice as fast since we have reduced the pause by half.
- More page results! 6 in fact.
- See more of each search result. You can see the whole title (just hover!), and more letters with better truncation.
- People should not be getting weird timeout messages anymore
- New administrative feature to limit the number of simultaneous quick nav searches, or to turn it off.

Open Search

- You can now search Confluence from the convenience of the search box in your browser (for FF and IE7). Just add the search provider from the search provider menu dropdown on the top right of your browser.



Site Search improvements

- Normal search is at least 5 times faster than it use to be.
- Better relevancy algorithm relevant

Miscellaneous Small Improvements

- Cleaner, sweeter breadcrumbs and 10 more rounded corners.

Plugins

- New version of Dynamic Task List (3.1.1)
  - Fixes TASK-108
- New version of Chart Plugin (1.12)
  - Fixes CONF-12792
- SNAPSHOT version of JIRA Issues Macro
  - Fixes lots of stuff
- 1.5-SNAPSHOT version of Advanced Macros
  - removed RSS macro
  - upgrade content-by-label
  - upgraded blogposts
  - upgraded recently-updated
  - Fixed Gallery Macro: ADVMACROS-39, ADVMACROS-48, ADVMACROS-49
- 1.5-SNAPSHOT version of Dashboard Macros
  - renamed recently-updated to recently-updated-dashboard to make way for new code in Advanced Macros
- New version of HTML Macros (1.5)
  - Includes RSS Macros now to take advantage of shared whitelist

atlassian-plugins upgrade to v 2.1.0.rc1

- added filter plugin type
- added standard decorators that are available to plugins
- added support for xml plugin artifacts

Release Notes 2.10-m7 ("Milestone 7")



Do not use this release to upgrade your production systems.

For all production use and testing of Confluence, please use the latest official release.
This release is a milestone development release for 2.10 ("Two Ten"). This is a public **development release** (DR) leading up to Confluence 2.10. Development releases are a snapshot of our work in progress, allowing our customers and especially plugin-developers to see what we're up to.

Who should upgrade?



Please note the following

- **Development releases are not safe** — Development releases are snapshots of the ongoing Confluence development process. As such:
  - While we try to keep these releases stable, they have not undergone the same degree of testing as a full release.
  - Features in development releases may be incomplete, or may change or be removed before the next full release.
- **No upgrade path** — Because development releases represent work in progress, we **can not** provide a supported upgrade path between development releases, or from any development release to the eventual final release. Thus, it is possible that you will not be able to migrate any data you store in a Confluence development release to a future Confluence release.

Our milestone releases aim to provide plugin developers with an opportunity to see the latest changes in the code.

Each milestone release has passed all our automatic tests, and has been used for one week on our official internal Confluence server. Most of the issues solved have been reviewed too, and usually milestone releases even have been load- and performance-tested for a while.

However, since our milestones releases are timeboxed (i.e. they get released every two weeks, no matter how far we have come implementing features and bugfixes), there is always a chance that we have new known bugs, which are scheduled to be fixed in the next milestone, or completely new bugs unknown even to us.

Additionally, our performance-testing and compatibility testing for databases and application servers is not done to the full extent. So, for example, a milestone release might behave well on a small installation but show severe problems when subjected to many users.

Upgrade Procedure

Follow the normal upgrade instructions to upgrade from Confluence 2.9.x to this release. *We strongly recommend that you backup your confluence-home directory and database before upgrading!*

Downloads

All development releases are available from Development Releases on the Atlassian website.

Issues resolved or improved in this release

Office Connector

- Changed the preview link to just "View". Previously it was "View Online" or "View as HTML". It will always be on the right of the actions links. This makes the alignment of the other links nicer.
- The slide viewer will always be downloading a few slides ahead of the current slide. This should make it appear to be loading slides faster when you advance the slides.
- Other minor refinements of the slide viewer UI.

RTE

- Keyboard shortcuts added for insert link/image (ctrl+k and ctrl+m)
- Keyboard shortcuts fixed for headings (ctrl+1 for heading one and so on)
- Links pasted into the rich text editor will work (for most browsers and to some extent for all browsers)
- Initial indent within a code macro is now displayed properly.
- Extra newlines are no longer added to bodies of most macros Eg {note}inline body{note} will be left alone by round-trip.
- Less extra whitespace will be added to wiki markup by going to rich text.
- Draft saving should now work in all browsers again.
- Editor tabs are now rounded in safari and firefox.

QuickNav

- A faulty optimisation that was causing problems with multi-word searches has been fixed.

Search

- Increased range of data boosting to cover a full year (instead of one month)

Did-you-mean

The administration screens for the did-you-mean configuration have been substantially improved.

Default Space Content

The default content used for the home page in a newly created space or personal space has been improved to show recently updated content, and a pagetree of the space.

Plugins

Included most recent bug-fix versions of:

- LeftNav Theme
- Userlister plugin
- ToC Plugin
- Social Bookmarking Plugin
- Email Page Plugin
- Live Search Plugin
- IM Presence Plugin
- Dashboard Macros Plugin
- Contributors Plugin
- Attachments Plugin
- JIRA Plugin
- Advanced Macros

Engine Room

Multiple bug fixes.

Release Notes 2.10-m8 ("Milestone 8")



Do not use this release to upgrade your production systems.

For all production use and testing of Confluence, please use the latest official release.
This release is a milestone development release for 2.10 ("Two Ten"). This is a public **development release** (DR) leading up to Confluence 2.10. Development releases are a snapshot of our work in progress, allowing our customers and especially plugin-developers to see what

we're up to.

Who should upgrade?



Please note the following

- **Development releases are not safe** — Development releases are snapshots of the ongoing Confluence development process. As such:
  - While we try to keep these releases stable, they have not undergone the same degree of testing as a full release.
  - Features in development releases may be incomplete, or may change or be removed before the next full release.
- **No upgrade path** — Because development releases represent work in progress, we **can not** provide a supported upgrade path between development releases, or from any development release to the eventual final release. Thus, it is possible that you will not be able to migrate any data you store in a Confluence development release to a future Confluence release.

Our milestone releases aim to provide plugin developers with an opportunity to see the latest changes in the code.

Each milestone release has passed all our automatic tests, and has been used for one week on our official internal Confluence server. Most of the issues solved have been reviewed too, and usually milestone releases even have been load- and performance-tested for a while.

However, since our milestones releases are timeboxed (i.e. they get released every two weeks, no matter how far we have come implementing features and bugfixes), there is always a chance that we have new known bugs, which are scheduled to be fixed in the next milestone, or completely new bugs unknown even to us.

Additionally, our performance-testing and compatibility testing for databases and application servers is not done to the full extent. So, for example, a milestone release might behave well on a small installation but show severe problems when subjected to many users.

Upgrade Procedure

Follow the normal upgrade instructions to upgrade from Confluence 2.9.x to this release. *We strongly recommend that you backup your confluence-home directory and database before upgrading!*

Downloads

All development releases are available from Development Releases on the Atlassian website.

Issues resolved or improved in this release

Writer

Bug fixes:

- CONF-13670 : fixed bug when Html Export used to fail on recently updated macro displaying an attachment.

Office Connector

- CONF-13613 : fixed several bugs in the pdf preview
- CONF-13643 : fixed previewing viewfile macro gives error on a new page with attachments
- Numerous other bugs. See the Office Connector release notes

Discovery

- Added recentness boosting to QuickNav.
- Added the space name to the tooltips for QuickNav results.
- Fixed bug preventing Did-You-Mean feature being enabled via the admin console.

Plugins

- Upgraded to Jira Macros Plugin v2.8.7, which contains several bugfixes (m7 contained v2.8.6)

Bug Fixing

Significant bug fixes since M7 include:

- CONF-13580: a blocker bug when editing pages with certain types of macros.
- CONF-13543: a blocker bug causing drafts to not be saved in certain situations
- CONF-13521: a serious issue with the RTE forcing the use of the base url.
- Several bugs (CONF-13338) fixed relating to the escaping of attachment filenames.
- Fixed the HTML around the quick search box that was causing a large right hand margin on some browsers.

Release Notes 2.10-rc1 ("Release Candidate 1")

Do not use this release to upgrade your production systems.

For all production use and testing of Confluence, please use the latest official release.
This release is a milestone development release for 2.10. This is a public **development release** (DR) leading up to Confluence 2.10. Development releases are a snapshot of our work in progress, allowing our customers and especially plugin-developers to see what we're up to.

Who should upgrade?

Please note the following

- **Development releases are not safe** — Development releases are snapshots of the ongoing Confluence development process. As such:
  - While we try to keep these releases stable, they have not undergone the same degree of testing as a full release.
  - Features in development releases may be incomplete, or may change or be removed before the next full release.
- **No upgrade path** — Because development releases represent work in progress, we **can not** provide a supported upgrade path between development releases, or from any development release to the eventual final release. Thus, it is possible that you will not be able to migrate any data you store in a Confluence development release to a future Confluence release.

Upgrade Procedure

Follow the normal upgrade instructions to upgrade from Confluence 2.9.x to this release. *We strongly recommend that you backup your confluence-home directory and database before upgrading!*

Downloads

All development releases are available from Development Releases on the Atlassian website.

Issues resolved or improved in this release

We are now reaching the end of the 2.10 release cycle.

We have released our first release candidate RC1 today, you can see it on our public Confluence installation at http://confluence.atlassian.com. There should be no more code changes made to Confluence from this point on unless any important issues are found in this release candidate.

While this is not the final release, and the official release notes will be unveiled next week, you should definitely have a look now to confirm everything still works fine with your plugins.

The RC1 is functionally almost equivalent to M8 which was announced in early November, but has a number of additional bugfixes. Check out our roadmap in the JIRA project

Known problems

The Edit Grid plugin doesn't work with this version of Confluence. See the Jira issue for more details.

Release Notes 3.0-beta2 ("Beta 2")

Welcome to our Beta Phase

The Confluence team is proud to present the first public beta release of our upcoming Confluence 3 release. It contains all the features we intend to ship. We will be publishing Beta 3 next week, and then a rapid succession of Release Candidates in two weeks. We aim at shipping Confluence 3.0 in (very) early June.

So, calling all plugin developers: If you want to make sure your plugin plays nicely with 3.0 on the launch date, and if you maybe even want to use our new features from within your plugin, this is really the last chance to get started. Download the Beta right now

Note: This is still not a stable release, it is not meant for production use. So our normal Milestone disclaimer still applies:

Milestone release advisory

Do not use this release to upgrade your production systems.

For all production use and testing of Confluence, please use the latest official release.
This release is a public **development release** ('milestone') leading up to Confluence 3.2, which will probably ship in mid March 2010. Development releases are a snapshot of our work in progress, allowing our customers and especially plugin developers to see what we're up to.

Who should upgrade?

Please note the following

- **Development releases are not safe** — Development releases are snapshots of the ongoing Confluence development process. As such:
  - While we try to keep these releases stable, they have not undergone the same degree of testing as a full release.
  - Features in development releases may be incomplete, or may change or be removed before the next full release.
- **No upgrade path** — Because development releases represent work in progress, we **can not** provide a supported upgrade path between development releases, or from any development release to the eventual final release. Thus, it is possible that you will not be able to migrate any data you store in a Confluence development release to a future Confluence release.

In supplying milestone releases, our aim is to provide plugin developers with an opportunity to see the latest changes in the code.

Each milestone release has passed all our automatic tests and has been used for one week on our official internal Confluence server. Most of the issues solved have been reviewed too, and all of our milestone releases even have been performance tested for a while.

However, since our milestones releases are timeboxed (i.e. they get released every two weeks, no matter how far we have come implementing features and bugfixes), there is always a chance that we have new known bugs that are scheduled to be fixed in the next milestone, or completely new bugs unknown even to us.

Additionally, we have not completed our performance testing and compatibility testing for databases and application servers. So, for example, a milestone release may behave well on a small installation but show severe problems when subjected to many users.

Upgrade Procedure

Follow the normal upgrade instructions to upgrade from Confluence 3.1.x to this release. *We strongly recommend that you backup your Confluence home directory and database before upgrading!*

Downloads

All development releases are available from the development releases page on the Atlassian website.

What is in it?

All the features we intend to ship are available in this beta. Unfortunately. the official documentation is not ready yet, so until we get a more official summary of the new features, please refer to our recent Milestone release notes, and try the beta out yourself.

Improvements

Since publishing Milestone 9, we have focussed on polishing our existing features. Here is a rundown on the changes since then:

Community features

Some important bug fixes:

Some IE issues: CONF-15497, CONF-15137, CONF-15474, CONF-15421, CONF-14593
Improvements to User Follow: CONF-15335, CONF-15279, CONF-15427, CONF-15286, CONF-15016
Removed the annoying page reload when following a user, CONF-15290

Macro Browser

Preview button removed and replaced by a refresh icon:



Some important bug fixes:

CONF-14720Inserting a macro via the macro browser would always scroll to the top, which is very annoying if you are editing the end of the page.

CONF-15487 It is now possible to insert more than one macro in Safari via the macro browser!

Office Connector

CONF-14798 Performance improvements on the backend for the viewfile macro.

Also, the initial loading of the Flash front end should be a little snappier. There was a default "pre-load" screen that added 1-2 seconds to the load time. We also changed the background color to white so it also adds the impression of not appearing until it's loaded.

old	new
JIRA Video.ppt 	JIRA Video.ppt

Engine Room

Confluence has been upgraded to the latest releases of Plugins 2 and SAL.

We did a small performance improvement by caching PersonalInformation objects by user name (CONF-15484), which will in particular speed up the rendering of profile pictures.

As part of this release cycle we fixed a lot of smaller bugs, but the most notable ones were related to content indexing:

CONF-15352 Fixed a race condition during content indexing which would result in leaked file handles, which had been introduced in M9
CONF-15483 Fixed a problem with a reindex being started before the plugin system was loaded. This would result in empty documents in the index.

Random

CONF-14322 was fixed, so change comments containing ampersands and angle brackets will no longer affect the system.

Known Issues

The Flash preview is broken in Office Connector that is bundled with beta 2. See CONF-15612. It's already fixed and it will be bundled in Beta 3 next week.

Release Notes 3.0-m3 ("Milestone 3")

Milestone release advisory

Do not use this release to upgrade your production systems.

For all production use and testing of Confluence, please use the latest official release.
This release is a public **development release** ('milestone') leading up to Confluence 3.2, which will probably ship in mid March 2010. Development releases are a snapshot of our work in progress, allowing our customers and especially plugin developers to see what we're up to.

Who should upgrade?

Please note the following

- **Development releases are not safe** — Development releases are snapshots of the ongoing Confluence development process. As such:
  - While we try to keep these releases stable, they have not undergone the same degree of testing as a full release.
  - Features in development releases may be incomplete, or may change or be removed before the next full release.
- **No upgrade path** — Because development releases represent work in progress, we **can not** provide a supported upgrade path between development releases, or from any development release to the eventual final release. Thus, it is possible that you will not be able to migrate any data you store in a Confluence development release to a future Confluence release.

In supplying milestone releases, our aim is to provide plugin developers with an opportunity to see the latest changes in the code.

Each milestone release has passed all our automatic tests and has been used for one week on our official internal Confluence server. Most of the issues solved have been reviewed too, and all of our milestone releases even have been performance tested for a while.

However, since our milestones releases are timeboxed (i.e. they get released every two weeks, no matter how far we have come implementing features and bugfixes), there is always a chance that we have new known bugs that are scheduled to be fixed in the next milestone, or completely new bugs unknown even to us.

Additionally, we have not completed our performance testing and compatibility testing for databases and application servers. So, for example, a milestone release may behave well on a small installation but show severe problems when subjected to many users.

Upgrade Procedure

Follow the normal upgrade instructions to upgrade from Confluence 3.1.x to this release. *We strongly recommend that you backup your Confluence home directory and database before upgrading!*

Downloads

All development releases are available from the development releases page on the Atlassian website.

Issues resolved or improved in this milestone since Confluence 2.10

User Hover

The first cut of User Hover has made it into m3. This mainly provides Confluence with the stepping block for features like Status. At present, it's only really useful to e-mail someone, or visit their user profile (both are notoriously difficult pre-m3).

Just hover over any user's name or profile picture, you should see the following information:



Performance

The Engine Room Team has started attacking general performance and also cluster performance. Our loadtests already show a significant improvement for high loads, and we hope it will show on EAC too.

In particular we made the following changes:

- Remove blocking in OGNL (CONF-14121)
- Rely on Coherence to do locking instead of using sychronized methods (CONF-14093)
- Remove unncessary sychronization on Hibernates UpdateTimestampsCache (CONF-14098)

More backend improvements related to caching and retrieval of macros have been done, which should make rendering long pages faster.

Overall we have seen a significant performance improvement under light to mid-level load, even more so on the clustered version of Confluence. This is mostly due to CONF-14093.

Rich Text Editor

Context menus are now available in the RTE (Agnes' ShipIt 9 project). Just right click in the editor! They are especially useful for table editing.



Notifications which show just relevant changes

When watching a page in Confluence, you normally get all the content whenever something changes. This isn't so useful, so we've added the option to get just the changes to the content emailed to you.

Just check the "Show changed content" box in Email Preferences. If enough people find this useful we are considering enabling it by default.



Viewing changes for content has also been similarly improved. By default, unchanged content is hidden and you can click the ellipsis to expand it.

Both improvements rely on 20% work done by David Taylor to dramatically improve the diffing algorithm used for Confluence content.

20% projects

Batching web resources

Confluence is now using Atlassian Plugins 2.2. beta4, which supports batching of web resources. This should result in a performance improvement, with fewer requests to the server for jss/css files.

See [here] for more details.

Roundtrip bugs

~~Yes, it's lame, but~~ DonW fixed a couple of particularly annoying round-trip RTE bugs in his 20% time:

- Images no longer get attached to the previous paragraph
- Emoticons can all be escaped in wiki markup, and will automatically be escaped when going from Rich Text to wiki markup.

Known issues

Plenty of known bugs. Check out JIRA.

Release Notes 3.0-m4 ("Milestone 4")

Milestone release advisory

Do not use this release to upgrade your production systems.

For all production use and testing of Confluence, please use the latest official release.
This release is a public **development release** ('milestone') leading up to Confluence 3.2, which will probably ship in mid March 2010. Development releases are a snapshot of our work in progress, allowing our customers and especially plugin developers to see what we're up to.

Who should upgrade?

Please note the following

- **Development releases are not safe** — Development releases are snapshots of the ongoing Confluence development process. As such:
  - While we try to keep these releases stable, they have not undergone the same degree of testing as a full release.
  - Features in development releases may be incomplete, or may change or be removed before the next full release.
- **No upgrade path** — Because development releases represent work in progress, we **can not** provide a supported upgrade path between development releases, or from any development release to the eventual final release. Thus, it is possible that you will not be able to migrate any data you store in a Confluence development release to a future Confluence release.

In supplying milestone releases, our aim is to provide plugin developers with an opportunity to see the latest changes in the code.

Each milestone release has passed all our automatic tests and has been used for one week on our official internal Confluence server. Most of the issues solved have been reviewed too, and all of our milestone releases even have been performance tested for a while.

However, since our milestones releases are timeboxed (i.e. they get released every two weeks, no matter how far we have come implementing features and bugfixes), there is always a chance that we have new known bugs that are scheduled to be fixed in the next milestone, or completely new bugs unknown even to us.

Additionally, we have not completed our performance testing and compatibility testing for databases and application servers. So, for example, a milestone release may behave well on a small installation but show severe problems when subjected to many users.

Upgrade Procedure

Follow the normal upgrade instructions to upgrade from Confluence 3.1.x to this release. *We strongly recommend that you backup your Confluence home directory and database before upgrading!*

Downloads

All development releases are available from the development releases page on the Atlassian website.

Issues resolved or improved in this milestone since Confluence 3.0 Milestone 3

Experimental Macro Browser

The Macro Browser is now available in the wiki markup editor, however it is currently **not in a state for proper use**. It is quite buggy and has been included for QA and the dev team, so please do not raise any issues against it yet.

Community

User Hover Fixes

A lot of outstanding issues with User Hover were fixed for m4.

First Cut of User Profiles

The profile tabs have been rearranged into a more intuitive order, with general and email preferences moved to a separate "Settings" tab:



Additional fields have now been added to your user profile.



Performance Improvements

We did the following small performance improvements:

- Caching/Queries optimized for page lookups (CONF-14273)
- Adjusted cache sizes (CONF-14294)
- Optimized access to attachments (CONF-14342)
- Less cache replication in a cluster (CONF-14339)

Although the performance improvement hasn't been as dramatically as in M2, but we still see a solid trend towards shorter and less varying response times. The overall performance improvement is around 10%-15% with more to come in the next milestones.

Known issues

- Performance issues with frequently updated attachments. One of the performance tweaks in M4 can cause issues if you have attachments with thousands of versions. (CONF-14422)
- The Macro Browser is very fresh and contains many bugs.

Check out JIRA for the full list of known bugs

Release Notes 3.0-m5 ("Milestone 5")

Milestone release advisory

Do not use this release to upgrade your production systems.

For all production use and testing of Confluence, please use the latest official release.
This release is a public **development release** ('milestone') leading up to Confluence 3.2, which will probably ship in mid March 2010. Development releases are a snapshot of our work in progress, allowing our customers and especially plugin developers to see what we're up to.

Who should upgrade?

Please note the following

- **Development releases are not safe** — Development releases are snapshots of the ongoing Confluence development process. As such:
  - While we try to keep these releases stable, they have not undergone the same degree of testing as a full release.
  - Features in development releases may be incomplete, or may change or be removed before the next full release.
- **No upgrade path** — Because development releases represent work in progress, we **can not** provide a supported upgrade path between development releases, or from any development release to the eventual final release. Thus, it is possible that you will not be able to migrate any data you store in a Confluence development release to a future Confluence release.

In supplying milestone releases, our aim is to provide plugin developers with an opportunity to see the latest changes in the code.

Each milestone release has passed all our automatic tests and has been used for one week on our official internal Confluence server. Most of the issues solved have been reviewed too, and all of our milestone releases even have been performance tested for a while.

However, since our milestones releases are timeboxed (i.e. they get released every two weeks, no matter how far we have come implementing features and bugfixes), there is always a chance that we have new known bugs that are scheduled to be fixed in the next milestone, or completely new bugs unknown even to us.

Additionally, we have not completed our performance testing and compatibility testing for databases and application servers. So, for example, a milestone release may behave well on a small installation but show severe problems when subjected to many users.

Upgrade Procedure

Follow the normal upgrade instructions to upgrade from Confluence 3.1.x to this release. *We strongly recommend that you backup your Confluence home directory and database before upgrading!*

Downloads

All development releases are available from the development releases page on the Atlassian website.

Issues resolved or improved in this milestone since Milestone 3

Macro Browser (in Markup Editor)

The Macro Browser is now available in the wiki markup editor! Features are listed below.









Browse Macros

You can now browse through a list of categorised macros and select a macro to insert. Macros that do not have categories define can be found in the 'All' categories tab. To insert a macro, you can either double click on the macro or click it once and then click the 'Next' button.

Insert Macro

For macros that have parameter metadata defined, you will see parameter labels and descriptions with the appropriate input fields. So far we are handling boolean and select-one-from-list type parameters. Everything else is considered a string. Please note that there is no field validation.

For macros that don't have the parameter metadata defined, you will see a single input field for the macro parameters and its notation guide help (if any) below.

Preview Macro

A default preview will be loaded only if the macro doesn't have required parameters. The preview can be updated by clicking on the 'Preview' button. This button location will change in the future.

Edit Macro

You can also edit a macro by placing your cursor in the **start tag** of a macro and then clicking the Macro Browser icon. It should display the 'Insert X Macro' page with it's parameters populated.

Known Issues

- Macro summary and parameter descriptions are currently under review by the tech writing team, so please ignore grammar/spelling/crappy English for this milestone.
- The macro icons (displayed in browse macros) are temporary. Hopefully we will be able to get much sexier icons.
- Some macros falsely declare they have bodies (e.g. gallery) so you may see some unnecessary body text areas

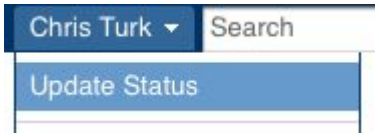
Community

Several major community features have their debut in this milestone.

Status

This milestone, we are releasing the first version of the Status feature. There is still a lot of work to be done, it's not particularly polished but we want your feedback.





First of all, set your status from the user menu, where you'll get a dialog box allowing you to enter your new status.

You can see your status and others' statuses from user hover, under the "Status List" option in user profiles, the "Favourite People" tab (see below), personal space sidebar (see below).

Follow

Marking another user as "favourite" is now meaningful. In milestone 5, you'll be able to get a list of statuses from the people you follow.



Check out the "Favourite People" tab in your user profile – you can get there via your user menu.

It will show you who you're following, and who's following you. You can also quickly follow someone with the user search, and see a list of status updates from the people you follow.

Personal Space Sidebar

With this release, we now have some structured form of a user profile for every user. On the right of any personal space*, there'll be a sidebar with various details of the person's space you're viewing. Don't worry, you can easily collapse the sidebar and it'll shift out of sight. Once you collapse the sidebar, it will stay collapsed for all personal spaces for the duration of your login session.

Here's what it looks like:



We plan to put more information on the sidebar, at the moment it looks a little silly if the user doesn't have a status set and hasn't filled out any of the extra profile fields we added last milestone.

This only works in the Default Theme.

Known Issues

- There are a few known issues with the Community work, the biggest thing is what you see at the moment is a first cut of the functionality, we plan to "sexy" it up a bit and make the status and profile pages look a lot nicer.
- You will also notice that when you favourite another user you will need to wait for the index to be flushed for your changes to be applied. This will be improved in a future milestone.

Bugfixes

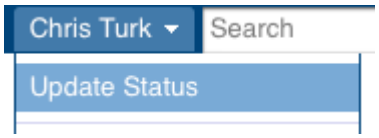
There are two relevant bug fixes in this milestone:

- Configuration to prevent anonymous users accessing user profile information (CONF-13276)
- Cannot insert links with Unicode characters in the URL (CONF-12707).

Engine Room

The Engine Room team has delivered one more performance improvement and a significant security improvement this milestone:

- Attachment retrieval should cache attachment IDs (CONF-14422)
- Anti-XSS mode is now enabled by default (CONF-12573).

We also fixed the regression in the previous milestone (M4) 

where pages with thousands of attachment versions brought all attachment retrievals to a grinding halt.

Release Notes 3.0-m6 ("Milestone 6")

Milestone release advisory

Do not use this release to upgrade your production systems.

For all production use and testing of Confluence, please use the latest official release.
This release is a public **development release** ('milestone') leading up to Confluence 3.2, which will probably ship in mid March 2010. Development releases are a snapshot of our work in progress, allowing our customers and especially plugin developers to see what we're up to.

Who should upgrade?

Please note the following

- **Development releases are not safe** — Development releases are snapshots of the ongoing Confluence development process. As such:
  - While we try to keep these releases stable, they have not undergone the same degree of testing as a full release.
  - Features in development releases may be incomplete, or may change or be removed before the next full release.
- **No upgrade path** — Because development releases represent work in progress, we **can not** provide a supported upgrade path between development releases, or from any development release to the eventual final release. Thus, it is possible that you will not be able to migrate any data you store in a Confluence development release to a future Confluence release.

In supplying milestone releases, our aim is to provide plugin developers with an opportunity to see the latest changes in the code.

Each milestone release has passed all our automatic tests and has been used for one week on our official internal Confluence server. Most of the issues solved have been reviewed too, and all of our milestone releases even have been performance tested for a while.

However, since our milestones releases are timeboxed (i.e. they get released every two weeks, no matter how far we have come implementing features and bugfixes), there is always a chance that we have new known bugs that are scheduled to be fixed in the next milestone, or completely new bugs unknown even to us.

Additionally, we have not completed our performance testing and compatibility testing for databases and application servers. So, for example, a milestone release may behave well on a small installation but show severe problems when subjected to many users.

Upgrade Procedure

Follow the normal upgrade instructions to upgrade from Confluence 3.1.x to this release. *We strongly recommend that you backup your Confluence home directory and database before upgrading!*

Downloads

All development releases are available from the development releases page on the Atlassian website.

Summary

This milestone is mainly about polishing features that were introduced previously, and a bit of back-end work for security and performance. We have fixed plenty of bugs of M5 too, and will keep fixing for M7.

Known issues

Given the recent work that has been done to improve "recent updates" (CONF-14434), updates shown on the dashboard will be affected shortly after the upgrade. You may notice that only changes made _after_ the upgrade will show. This will be rectified as soon as the index is rebuilt.

Improvements and new Features

New PDF Export

- This milestone contains the first cut of the Improved PDF export, which also give you more control over the conversion process by allowing you to use CSS. Please find the detailed documentaton on CAC. CONF-2079

There are still plenty of known bugs with the new PDF, we included it a bit early to be able to get feedback on the CSS-styling process, not

for beta-testing purposes...

Engine Room

The ER team has implemented several more important fixes and improvements in this milestone:

- *supportability* - integrated thread-dump tool in the administration screen (CONF-12395)
- *security* - editing comments in wiki markup is no longer double-HTML-escaped in anti-XSS mode (CONF-14601)
- *cluster performance* - avoid unnecessary updates to distributed cache to improve cluster performance (CONF-14657)
- *general performance* - improve performance of retrieving many attachments, like on the attachments page (CONF-14422)
- *front-end performance* - add caching headers to attachments rendered with a "?version=x" parameter (CONF-8034).

Office Connector

- When launching an external editor from Confluence, you don't have to login again. CONF-14705
- You can now monitor what is in the conversion queue and what is actually being converted CONF-14707

Macro Browser

The following bugs have been fixed:

Errors were reported by the JIRA trusted connection.

- APP_UNKNOWN; Unknown Application: {0}; ["confluence:4557196"]

JIRA Issues (0 issues)										
Type	Key	Summary	Assignee	Reporter	Priority	Status	Resolution	Created	Updated	Due

Status and Follow

The following bugs have been fixed:

Errors were reported by the JIRA trusted connection.

- APP_UNKNOWN; Unknown Application: {0}; ["confluence:4557196"]

JIRA Issues (0 issues)										
Type	Key	Summary	Assignee	Reporter	Priority	Status	Resolution	Created	Updated	Due

Bugfix-team

- It's safe to use Sections and Columns in the RTE again. (CONF-13823, CONF-14282)

Release Notes 3.0-m7 ("Milestone 7")

Milestone release advisory

Do not use this release to upgrade your production systems.

For all production use and testing of Confluence, please use the latest official release.
This release is a public **development release** ('milestone') leading up to Confluence 3.2, which will probably ship in mid March 2010. Development releases are a snapshot of our work in progress, allowing our customers and especially plugin developers to see what we're up to.

Who should upgrade?

Please note the following

- **Development releases are not safe** — Development releases are snapshots of the ongoing Confluence development process. As such:
  - While we try to keep these releases stable, they have not undergone the same degree of testing as a full release.
  - Features in development releases may be incomplete, or may change or be removed before the next full release.
- **No upgrade path** — Because development releases represent work in progress, we **can not** provide a supported upgrade path between development releases, or from any development release to the eventual final release. Thus, it is possible that you will not be able to migrate any data you store in a Confluence development release to a future Confluence release.

In supplying milestone releases, our aim is to provide plugin developers with an opportunity to see the latest changes in the code.

Each milestone release has passed all our automatic tests and has been used for one week on our official internal Confluence server. Most of the issues solved have been reviewed too, and all of our milestone releases even have been performance tested for a while.

However, since our milestones releases are timeboxed (i.e. they get released every two weeks, no matter how far we have come implementing features and bugfixes), there is always a chance that we have new known bugs that are scheduled to be fixed in the next milestone, or completely new bugs unknown even to us.

Additionally, we have not completed our performance testing and compatibility testing for databases and application servers. So, for example, a milestone release may behave well on a small installation but show severe problems when subjected to many users.

Upgrade Procedure

Follow the normal upgrade instructions to upgrade from Confluence 3.1.x to this release. *We strongly recommend that you backup your Confluence home directory and database before upgrading!*

Downloads

All development releases are available from the development releases page on the Atlassian website.

Summary

The Confluence team is proud to present another incremental milestone, called M7. We are currently aiming at shipping M8 in two weeks and M9 in four weeks. M9 will be almost feature-complete and therefore similar to a beta release. After a couple of betas and release candidates we intend to ship Confluence 3 in late May. As always, things can still change a little, so there is no fixed date. But if you are a plugin developer, you want to have your plugin tested well before the big marketing buzz kicks in. We try to remain as backward-compatible as we can for Confluence 3, but if something slipped through the cracks we need your feedback, and we need it now. So do have a look at this milestone release and tell us what you think.

Macro Browser and Rich Text Editor

You can now choose the kind of right-click context menu you want in the Rich Text Editor. A new button on the toolbar, handily called "Toggle context menu", toggles between the browser's default context menu and our TinyMCE one.

You can also tell the TinyMCE context menu that it isn't wanted via the "Disable context menu" option.



Macro Browser search has been improved and now accepts multiple words (order not important) and camel-casing (matches either case in consecutive words).
The filtering logic has been added to AUI as a new *commented* method of AJS, "filterBySearch" - feel free to use it.



Community Improvements

This milestone shows a number of improvements from the Community Team.

Profile Layout

In the spirit of making the profile page actually page you want to visit, it's been given a facelift and has now come closer in line with out stardard UI.

It also includes your current status is clearly visible at the top of the profile page.

For now you have a Status Updates section on the right hand side of the profile page, the next milestone will see this will be replaced with the new Recently Updated Macro (see below) so it will include other kinds of activity as well. It's here now because otherwise you wouldn't have anywhere to delete or clear your status, eventually will be moved to another tab and/or page.



User Status Updates

- Delete And Clear Status
  You can delete status entries from the list on the right. Someone hacks in your status, you can now always delete it. You can also clear your current status, so any location showing your current status will show nothing. This won't remove it from your history but your current status will now be considered cleared.



- Status Icon

  A new status Icon was introduced by the Design team. The lightbulb has been replaced with some sexy quotes (“).
- Immediate Updates
  Updating your status will use Javascript-Magic to update your status on the current page, giving you immediate feedback, with no need to refresh the page.



Recently Updated Macro

This is the first milestone that shows the hard work that David Loeng and Chris Broadfoot have been doing with our Recently Updated Macro.

To build up activity in the old recently updated macro the process would simply query the lastModified date of all the relevant content, then sort it accordingly. The side effect of this was, if two people edited the same page, you would only see the one edit in your list. Not only that, if you wanted to follow a certain user's activity and someone edited a page after them, you wouldn't see any update!

We've worked very hard on this and changed the way Confluence stores changes to documents and now tracks **all changes** every time a document, page, profile, status, anything, is updated.

This screen shot shows two edits from two different people on the same document, something that wasn't previously possible. These changes were critical for implementing a proper follow feature for 3.0.



Bug Fixes

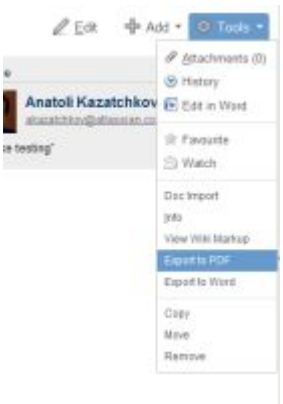
A number of bugs have been fixed by the community team, here are a list of some of the main ones:

- CONF-14773 Expanding and collapsing the personal sidebar will no longer affect the comments on that page/blogpost.
- CONF-14870 Delete links for user status items are now properly integrated with the new XSRF protection and actually work.
- CONF-14778 Printing pages from a personal space no longer includes the sidebar
  Other bug fixes:
- CONF-14689 Comment edits now send notifications

Plus a number of other fixes.

PDF export team

Removed the PDF Export++ option. The new PDF export is what you get when you select just **PDF Export**. The old PDF export was removed.



Added javascript to the space export tree that will select/deselect all descendants of the clicked node. This change applies to all export formats not just PDF. This should save hours for the Atlassian Techwriting team.



Bugs fixed

CONF-14906 - Fonts are too large in the page index macro when exported to PDF
CONF-14905 - Note macro missing colored background in PDF export
CONF-14902 - content by label macro looks ugly when exported to PDF
CONF-14900 - Bookmarks macro is producing really large font when exported to PDF
CONF-14899 - Attachments macro is showing velocity junk when exported to PDF
CONF-14898 - Panel macros (warning, note, info) could use some extra padding in PDF export
CONF-14897 - Spaces macro renders as a single bullet item list when exported to PDF
CONF-14895 - Recently-used-labels macro with a table style doesn't look like a table at all
CONF-14894 - Blockquote text too large in pdf export

Release Notes 3.0-m8 ("Milestone 8")

Milestone release advisory

Do not use this release to upgrade your production systems.

For all production use and testing of Confluence, please use the latest official release.
This release is a public **development release** ('milestone') leading up to Confluence 3.2, which will probably ship in mid March 2010. Development releases are a snapshot of our work in progress, allowing our customers and especially plugin developers to see what we're up to.

Who should upgrade?

Please note the following

- **Development releases are not safe** — Development releases are snapshots of the ongoing Confluence development process. As such:
  - While we try to keep these releases stable, they have not undergone the same degree of testing as a full release.
  - Features in development releases may be incomplete, or may change or be removed before the next full release.
- **No upgrade path** — Because development releases represent work in progress, we **can not** provide a supported upgrade path between development releases, or from any development release to the eventual final release. Thus, it is possible that you will not be able to migrate any data you store in a Confluence development release to a future Confluence release.

In supplying milestone releases, our aim is to provide plugin developers with an opportunity to see the latest changes in the code.

Each milestone release has passed all our automatic tests and has been used for one week on our official internal Confluence server. Most of the issues solved have been reviewed too, and all of our milestone releases even have been performance tested for a while.

However, since our milestones releases are timeboxed (i.e. they get released every two weeks, no matter how far we have come implementing features and bugfixes), there is always a chance that we have new known bugs that are scheduled to be fixed in the next milestone, or completely new bugs unknown even to us.

Additionally, we have not completed our performance testing and compatibility testing for databases and application servers. So, for example, a milestone release may behave well on a small installation but show severe problems when subjected to many users.

Upgrade Procedure

Follow the normal upgrade instructions to upgrade from Confluence 3.1.x to this release. *We strongly recommend that you backup your Confluence home directory and database before upgrading!*

Downloads

All development releases are available from the development releases page on the Atlassian website.

Summary

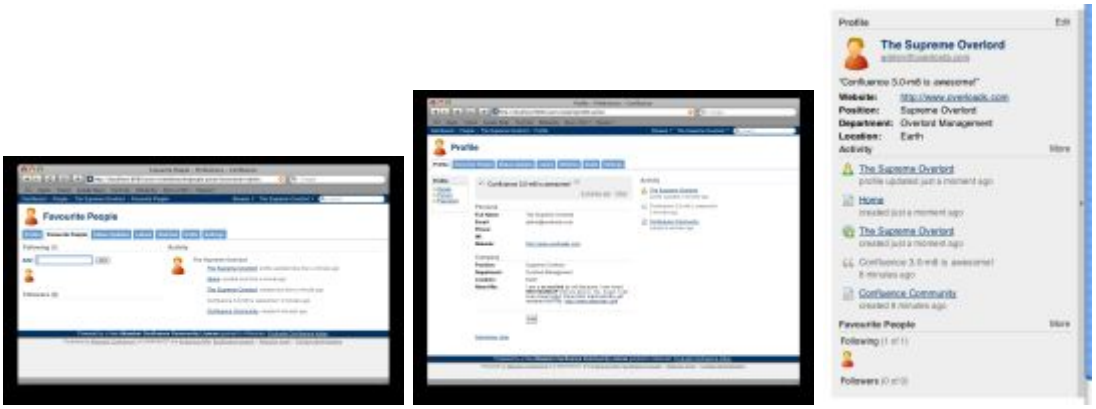
The Confluence team is excited to bring you our latest milestone release: Confluence 3.0-m8, "Milestone 8". This release includes a huge amount of improvements to the community features, additional performance improvements and many bug fixes.

Community

Be sure to check out the great improvements in Confluence's new community features:

- Improved Activity list on Profile Page - The user status list has been moved to its own tab and the new Recently Updated Macro has been put in its place. It's also been given a face lift and fits nicely into the profile page. Facebook eat your heart out.
- Status list moved to seperate tab with pagination - The user status list now has a proper location, it's own tab. This tab is the central place where you manage and view your historical status updates. Links to delete old status updates will be found on this tab.
- Each status item has its own page so you can link to a status - You can now create links to old status objects allowing for historical references as each status can be displayed on its own page.
- Personal space sidebar has been updated to include an activity list and your follower information - The personal space sidebar is coming along nicely and looks great with the new recently updated macro and the follow information.
- Adding a follower on your Favourite People tab (or using the hover) will refresh the page and update the list immediately - We changed how the follower information was looked up as to not use the lucene index and hooked into the hover so most changes on this page will trigger a refresh and the updates will be seen immediately! No more wondering if your new follower actually worked

- Those extra fields on your User Profile page are now run through the wiki renderer allowing links and simple formating



Macro Browser

The macro browser team spent the past two weeks working behind the scenes on placeholders for the rich text editor. Unfortunately, this change is too complex and risky to get into 3.0, so we have just a few small improvements in the macro browser this milestone:

- More parameters and documentation are available for macros in the macro browser. We now have 100% of the bundled macros with documented parameters, and will be just fine-tuning the documentation for the remainder of the release.
- The Office Connector {viewfile} macro shows up in the macro browser as separate macros for Word, Excel, PDF and Powerpoint documents.

PDF Export

Balsamiq Mockups and several other complex macros now export to PDF correctly (CONF-14792).

Why don't you try out the new PDF export right now? On any page, go to *Tools*, *Export as PDF*.

Engine Room

The Engine Room team implemented several important performance improvements for this milestone. Specifically:

- Searching should be faster and not blocked by index updates every minute. This should also make the dashboard and macros that use the search index slightly faster. CONF-7749, CONF-14803
- Image attachments embedded in a page are now cached by your web browser. This should make pages with lots of images faster to load and reduce the load on the server. CONF-8034
- Upgraded the WebDAV plugin to 2.0-beta2. This plugin has been released to the community for a while now, to very positive response. It should be more compatible with more clients now, though there is still some special configuration needed for Windows Clients. Read the documentation about client-configuration in order to test it.

20% Projects

Administrators can now customise the size of Confluence's caches from the web interface. Previously, you needed to edit an XML file and restart Confluence whenever you wanted to adjust the caches. This addresses one of the most popular supportability feature requests in the Confluence JIRA project, with 34 votes: CONF-12836.



Bugs Fixed

Several macros which were broken in M7 are now fixed: {flowchart}, {rsvp} and {pagetree}. CONF-14615, CONF-14829, CONF-14581

The Gliffy plugin and Beanshell macros now work with Confluence 3.0; previously customisation was required to get them to work properly with our latest development versions.

Below is a complete list of issues resolved in Confluence 3.0-m8.

Errors were reported by the JIRA trusted connection.

- APP_UNKNOWN; Unknown Application: {0}; ["confluence:4557196"]

JIRA Issues (1 issues)

Type	Key	Summary	Assignee	Reporter	Priority	Status	Resolution	Created	Updated	Due
	CONF-15102	Allow macro Body Text label to be overridden and add a description	David Taylor [Atlassian]	David Taylor [Atlassian]		Resolved	Fixed	Apr 05, 2009	May 16, 2009	

Release Notes 3.0-m9 ("Milestone 9")

Milestone release advisory

Do not use this release to upgrade your production systems.

For all production use and testing of Confluence, please use the latest official release.
This release is a public **development release** ('milestone') leading up to Confluence 3.2, which will probably ship in mid March 2010. Development releases are a snapshot of our work in progress, allowing our customers and especially plugin developers to see what we're up to.

Who should upgrade?

Please note the following

- **Development releases are not safe** — Development releases are snapshots of the ongoing Confluence development process. As such:
  - While we try to keep these releases stable, they have not undergone the same degree of testing as a full release.
  - Features in development releases may be incomplete, or may change or be removed before the next full release.
- **No upgrade path** — Because development releases represent work in progress, we **can not** provide a supported upgrade path between development releases, or from any development release to the eventual final release. Thus, it is possible that you will not be able to migrate any data you store in a Confluence development release to a future Confluence release.

In supplying milestone releases, our aim is to provide plugin developers with an opportunity to see the latest changes in the code.

Each milestone release has passed all our automatic tests and has been used for one week on our official internal Confluence server. Most of the issues solved have been reviewed too, and all of our milestone releases even have been performance tested for a while.

However, since our milestones releases are timeboxed (i.e. they get released every two weeks, no matter how far we have come implementing features and bugfixes), there is always a chance that we have new known bugs that are scheduled to be fixed in the next milestone, or completely new bugs unknown even to us.

Additionally, we have not completed our performance testing and compatibility testing for databases and application servers. So, for example, a milestone release may behave well on a small installation but show severe problems when subjected to many users.

Upgrade Procedure

Follow the normal upgrade instructions to upgrade from Confluence 3.1.x to this release. *We strongly recommend that you backup your Confluence home directory and database before upgrading!*

Downloads

All development releases are available from the development releases page on the Atlassian website.

Summary

We are one week away from feature freeze, so this is the last major milestone before we enter beta next week. We are still working on some wording, on some plugin upgrades, and a few UI tweaks, but M9 is pretty much how Confluence 3 will look like. From now on, it's mainly about fixing up non-critical bugs and doing code reviews. Many Code reviews!

Enjoy:

Community

Followed Users and Favourite Spaces

The terminology has been cleaned up and made consistent across the application. If you want to see what a user is doing you *follow* them, if you want to see what happens in a space, you mark that space as *favourite*. You can mark a personal space as *favourite*, and this is independent of *following* the user who owns the space.

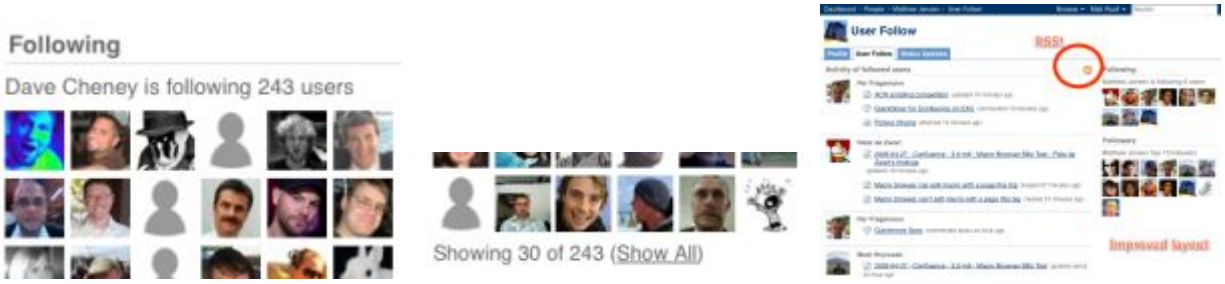
One area we still want to improve is the 'User Follow' menu item and profile tab, which needs a better name.

Follow Tab RSS and Design Improvements

You can now get an RSS feed on any user's 'User Follow' tab. This feed will include all activity of the users followed by that person. This is a great new way of monitoring content on the Extranet.

We've redesigned the follow tab to highlight the activity information more and handle large lists of followers. We also added some clarifying text to this tab to better describe the different sections of this page.

There is now a maximum number of followers shown on this page, and separate pages to show the complete list. This addresses the performance and usability problems discovered in the previous milestone when you had a large number of followers or people following you.



Profile Page Updates

The layout of the profile page has been changed to try to improve the visiblity of the activity. This puts the user's personal information over on the right hand side of the page.

On the left hand side, we now have new links to follow a user and favourite their space.



Update User Status Permission

A new global permission was introduced to give system administrators control over who can use the user status feature.

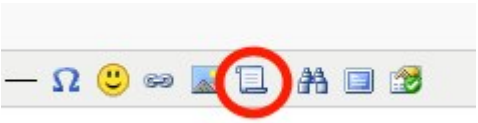
If the permission is turned off for a user (or their group) then they will not have the ability to set their user status, nor will they have an 'Status Updates' tab in their profile.

Macro Browser

Macro Browser in the RTE

The Macro Browser is now finally available in the Confluence rich text editor (RTE)! It will assist you in inserting and editing macros, without having to understand the wiki markup.

Editing a macro can be done by placing the cursor in the start or end tags of the macro wiki markup. If you have context menus enabled, you will also see a new option for insert/editing macros.



Upgraded Macro Icons

Thanks to Stephen, we have even better icons for the macros in the macro browser. Check 'em out!



AUI 1.0

Confluence is now using the first release of AUI, 1.0.

PDF Export

Many bug fixes have been done for the new PDF export:

CONF-15220 - PDF Export converts a bulleted sublist into a numbered sublist
CONF-15148 - Follow macro should not include the 'Add' action when exported to PDF
CONF-15042 - PDF Space Export text uses US spelling instead of Australian
CONF-15038 - PDF Export from a page with non-default theme is trying to use the old method, and so it fails with an error
CONF-14909 - column and section macros render differently in PDF export when they have borders
CONF-14904 - IM presence macro shows a roque nbsp when exporting to PDF
CONF-14901 - The PDF export doesn't recognize all literal colors when used in the color macro
CONF-14889 - PDF Export from the Info pages is trying to use the old method, and so it fails with an error

Engine Room

Confluence has been upgraded to the latest releases of Plugins 2 and SAL.

Release Notes 3.0-rc1 ("Release Candidate 1")

Hi everyone,

we are approaching the 3.0 release date. We have built our first release candidate, and we hope to turn this into the final release unchanged soon.

There have been a few changes since the last two betas: The translations are now up to date, more icons have been added to the macro browser, and we have tweaked a few screens, like the profile page. And of course we have fixed dozens of bugs (see JIRA URLs below). We are very pleased with the quality of the release and we are currently not aware of any showstoppers or critical issues. Chances are that the final release will be just a rebranded RC, however you should still not use the RC on production systems, since we still have some final testing planned. We encourage you to use it for staff training purposes and for your upgrade testing though, since no more UI changes or changes to the upgrade process are scheduled.

Milestone release advisory

Do not use this release to upgrade your production systems.

For all production use and testing of Confluence, please use the latest official release.
This release is a public **development release** ('milestone') leading up to Confluence 3.2, which will probably ship in mid March 2010. Development releases are a snapshot of our work in progress, allowing our customers and especially plugin developers to see what we're up to.

Who should upgrade?

Please note the following

- **Development releases are not safe** — Development releases are snapshots of the ongoing Confluence development process. As such:
  - While we try to keep these releases stable, they have not undergone the same degree of testing as a full release.
  - Features in development releases may be incomplete, or may change or be removed before the next full release.
- **No upgrade path** — Because development releases represent work in progress, we **can not** provide a supported upgrade path between development releases, or from any development release to the eventual final release. Thus, it is possible that you will not be able to migrate any data you store in a Confluence development release to a future Confluence release.

In supplying milestone releases, our aim is to provide plugin developers with an opportunity to see the latest changes in the code.

Each milestone release has passed all our automatic tests and has been used for one week on our official internal Confluence server. Most of the issues solved have been reviewed too, and all of our milestone releases even have been performance tested for a while.

However, since our milestones releases are timeboxed (i.e. they get released every two weeks, no matter how far we have come implementing features and bugfixes), there is always a chance that we have new known bugs that are scheduled to be fixed in the next milestone, or completely new bugs unknown even to us.

Additionally, we have not completed our performance testing and compatibility testing for databases and application servers. So, for example, a milestone release may behave well on a small installation but show severe problems when subjected to many users.

Upgrade Procedure

Follow the normal upgrade instructions to upgrade from Confluence 3.1.x to this release. *We strongly recommend that you backup your Confluence home directory and database before upgrading!*

Downloads

All development releases are available from the development releases page on the Atlassian website.

Issues resolved between Beta2 and Beta3: http://jira.atlassian.com/secure/IssueNavigator.jspa?reset=true&pid=10470&fixfor=14602
Issues resolved between Beta3 and RC1: http://jira.atlassian.com/secure/IssueNavigator.jspa?reset=true&pid=10470&fixfor=14625

Thanks for your feedback so far, it has helped us track down a few bugs, which will result in a smoother upgrade for everyone!

Cheers,
The Confluence Team

Release Notes 3.1-m1 ("Milestone 1")

Outdated page

These are not the release notes you are looking for

This page is outdated, but has high Google ranking

Please refer to the official Confluence 3.1 release notes over here.

Milestone release advisory

Do not use this release to upgrade your production systems.

For all production use and testing of Confluence, please use the latest official release.
This release is a public **development release** ('milestone') leading up to Confluence 3.2, which will probably ship in mid March 2010. Development releases are a snapshot of our work in progress, allowing our customers and especially plugin developers to see what we're up to.

Who should upgrade?

Please note the following

- **Development releases are not safe** — Development releases are snapshots of the ongoing Confluence development process. As such:
  - While we try to keep these releases stable, they have not undergone the same degree of testing as a full release.
  - Features in development releases may be incomplete, or may change or be removed before the next full release.
- **No upgrade path** — Because development releases represent work in progress, we **can not** provide a supported upgrade path between development releases, or from any development release to the eventual final release. Thus, it is possible that you will not be able to migrate any data you store in a Confluence development release to a future Confluence release.

In supplying milestone releases, our aim is to provide plugin developers with an opportunity to see the latest changes in the code.

Each milestone release has passed all our automatic tests and has been used for one week on our official internal Confluence server. Most of the issues solved have been reviewed too, and all of our milestone releases even have been performance tested for a while.

However, since our milestones releases are timeboxed (i.e. they get released every two weeks, no matter how far we have come implementing features and bugfixes), there is always a chance that we have new known bugs that are scheduled to be fixed in the next milestone, or completely new bugs unknown even to us.

Additionally, we have not completed our performance testing and compatibility testing for databases and application servers. So, for example, a milestone release may behave well on a small installation but show severe problems when subjected to many users.

Upgrade Procedure

Follow the normal upgrade instructions to upgrade from Confluence 3.1.x to this release. *We strongly recommend that you backup your Confluence home directory and database before upgrading!*

Downloads

All development releases are available from the development releases page on the Atlassian website.

Yay, we are back!

After 3.0 came out, we spent some time in bugfix-mode, then we had a 20% week, 2 devspeed weeks, and the last week was mainly spent planning the next release. So that's why there has not been much visible progress recently lately. Development on core 3.1 features has started, but in this milestone we mainly ship all the bugfixes we did for 3.0.1, some small improvements, and three 20% projects.

New Webstart Installer

We are trying something new. While we will provide the downloadable files too, we encourage you to try our webstart based evaluation installer. Visit http://webstart.atlassian.com/confluence-3.1-m1-r2/confluence.jnlp for a snappy installation experience.

New Native Mac Installer

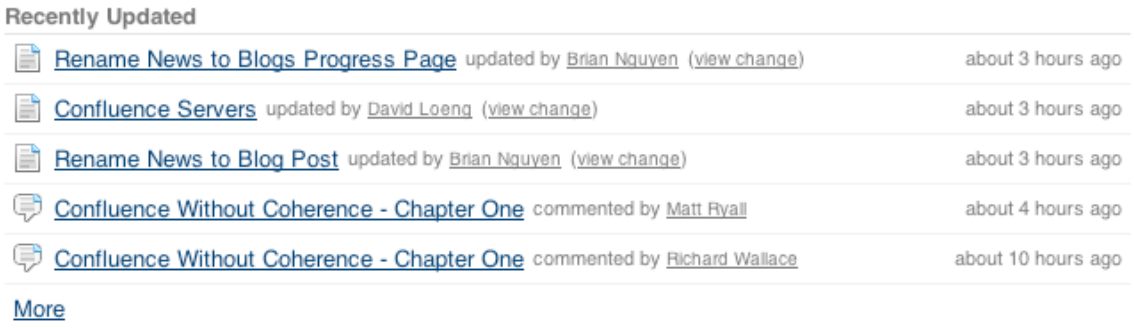
When you go to the download section, you will notice a new DMG file. This contains our new native experimental Mac installer. Give it a try and tell us what you think. Instructions coming soon!

And here are all the cool features in M1 you will get.

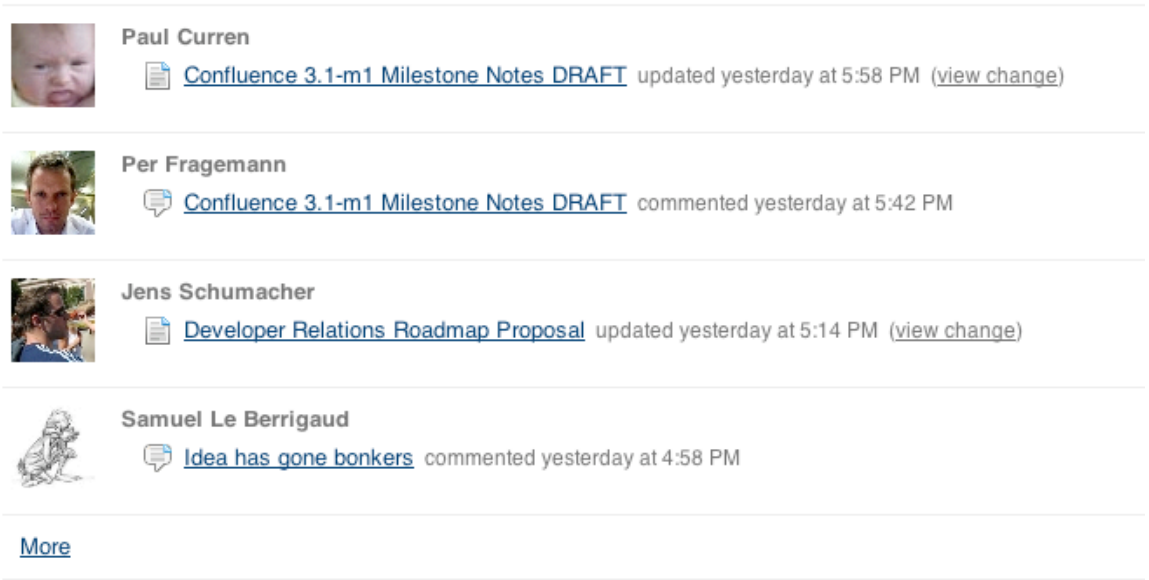


Interactive "Get more"

Recently updated macros are now interactive. You will get a little "More" link to load new updates dynamically within the page. This small but awesome feature can also be found on your network page, and in the space's recent update tab. It doesn't screenshot too well, so try it out for yourself. Anyway here is what the macro looks like:



And here is what the network tab looks like:



Caveats:

- We feel that the link is a bit small and can be mistaken for a normal link. Especially on the static pages (network, updates in space) it makes sense to make the button much larger and more prominent. We will work with Design to find a solution
- The link also needs to go into the profile page (so you can see more updates on everyone's stream), and into the status tab. Dave will add that soon

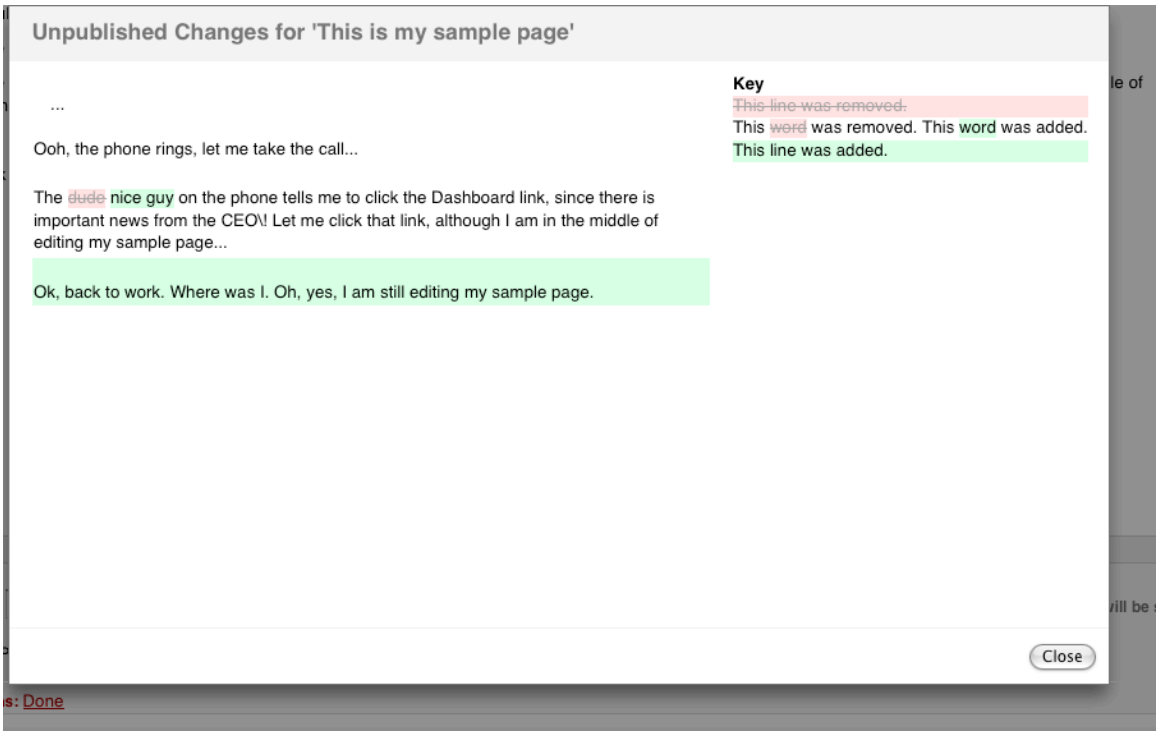

This is a Dave Loeng 20% production



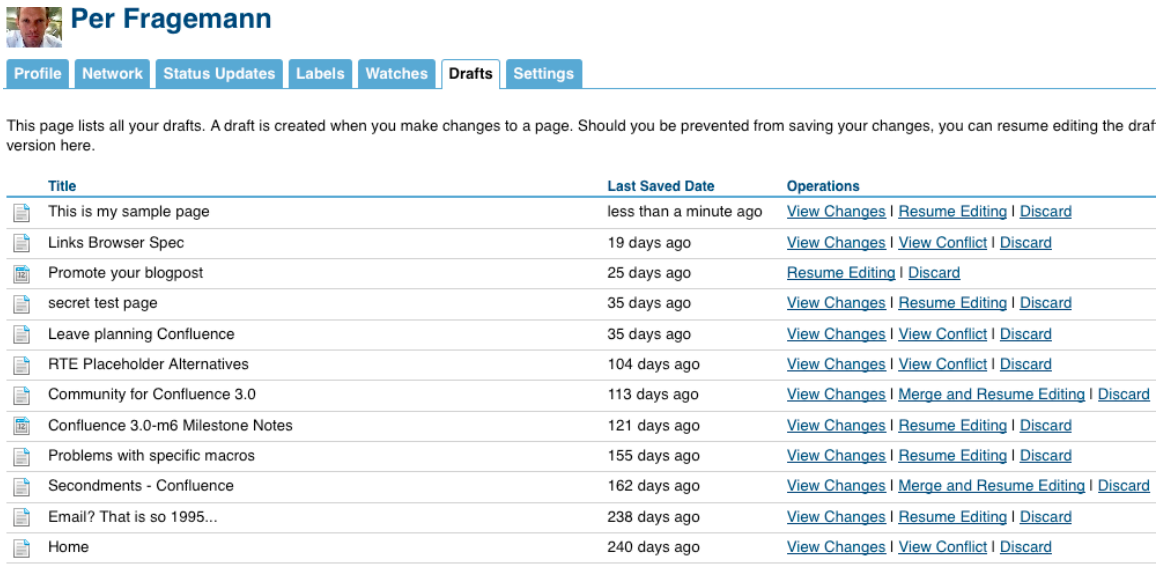
Draft diffs and visibility

Always wondered what your old drafts were about? Or when a draft is being saved, what part of that huge document you're editing has actually changed? Well, now you can, thanks to DraftDiffs.

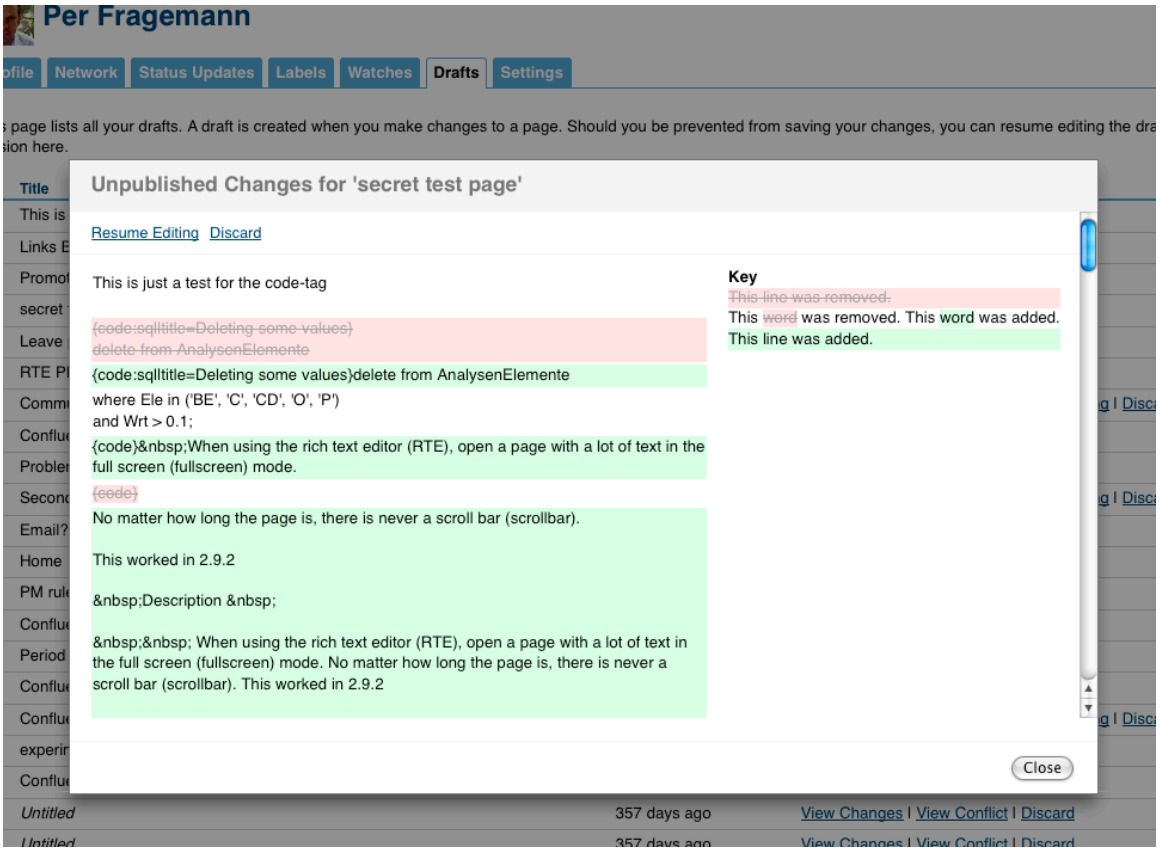
View the diff right from the edit page:



View what your old drafts are about from the drafts overview page:



Which takes you to this screen:



Known issues: Annoying Javascript alert when there are no changes at all.This feature unfortunately also highlights our roundtripping bugs... ahem...



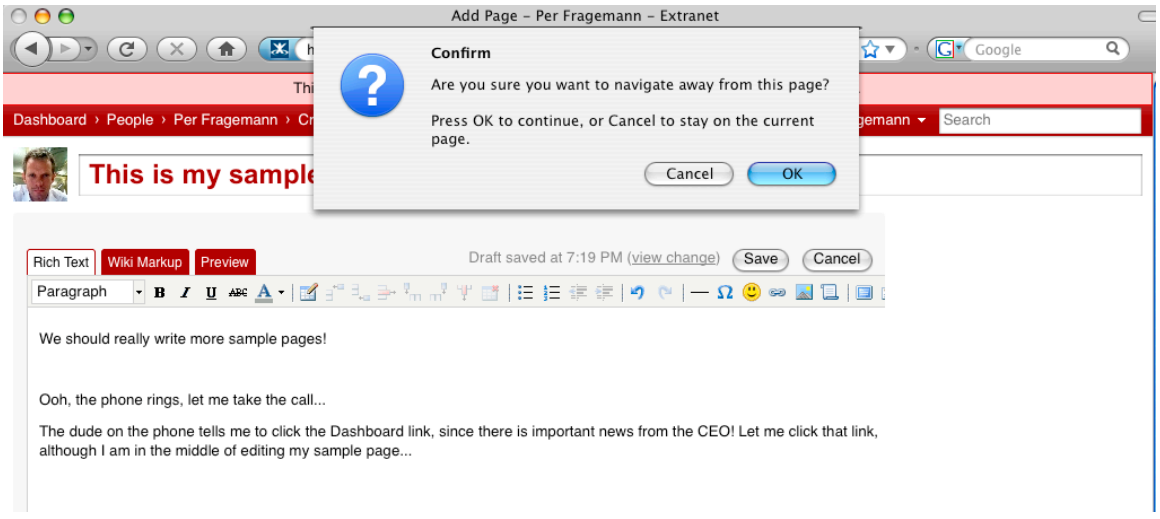
This is a Brian Nguyen 20% production



Small Improvements

Edit loss warning

Whenever you are in edit mode and decide to leave the page or close it, you will now get a warning popup. This is especially useful when writing comments (no drafts...) and for technically challenged people who don't understand drafts in the first place. As raised in the JIRA issue, we may have to provide a setting to disable these warnings (it can be a bit annoying if you like to change your mind often), so watch CONF-16075 for updates. The idea has been raised that we may only want this feature for comments when no drafts are being saved. We will dogfood this over the next few weeks, and see how we like it or not.



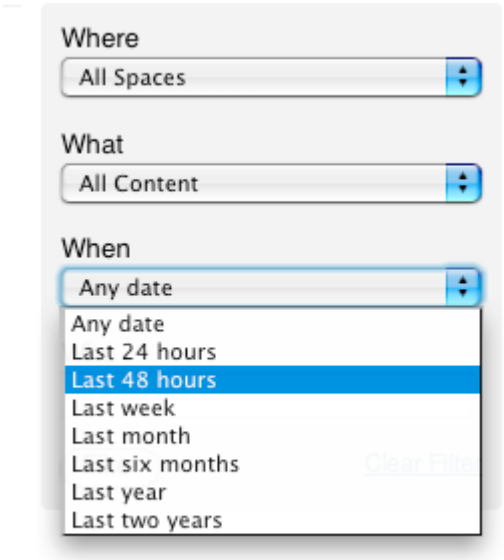
Gradient

This is just an experiment, probably not shipping as a default for 3.1, but worth a look. Do you like it?



More options to search for by time

We added two more options to filter by time: Last 6 months, Last year, Last 2 years, and fixed the calculation for today/yesterday by making it work by 24h instead of what the server thinks is midnight. 48h obviously encompasses the last 24 hours two, which was akward in the old system, where yesterday would really just mean yesterday, but not today.



Plugin points in the editor

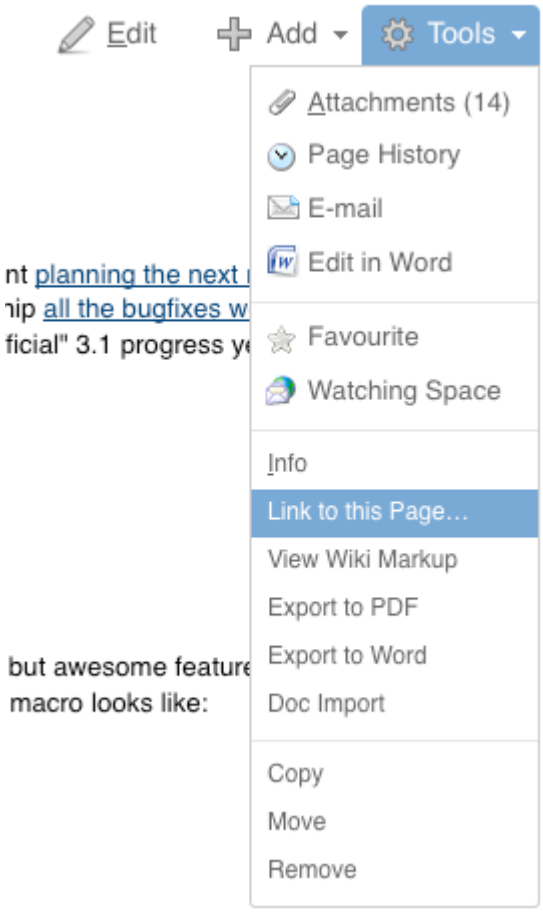
We were approached by a developer who wants to write a spellchecker plugin for Confluence, so Agnes volunteered to pluginpointinize the editor. You won't see anything right now, but it will make many plugin authors happy. Read the documentation for how to write plugins for the editor.


This is an Agnes Ro 20% production

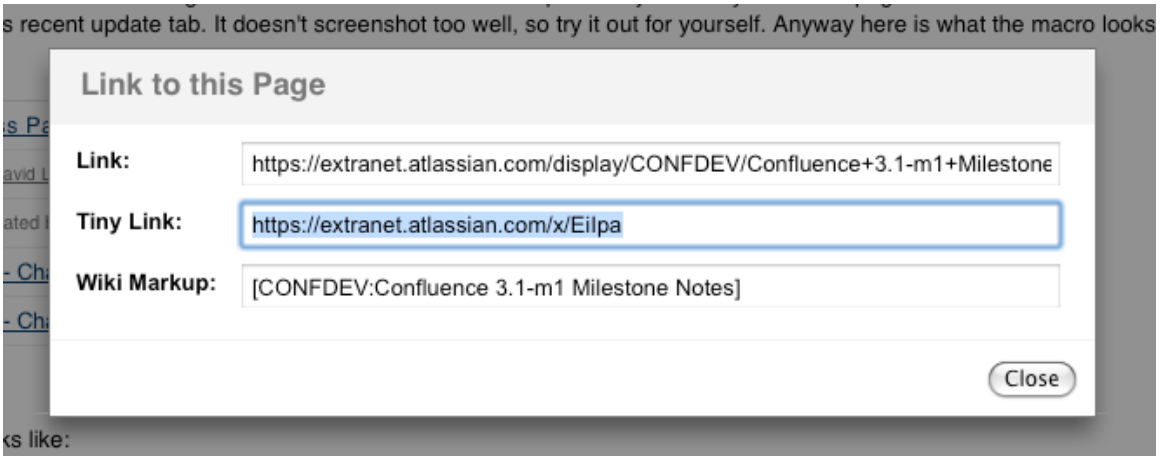


"Link to this page" Dialog

In order to make Tiny Links more accessible, we added a new Menu option which opens a new dialog, which shows the three ways to link to a page. Full URL, TinyURL, and Wiki Link. Probably the Full URL is not as important, but the dialog looked a bit unbalanced, and it makes it clearer that both URL's are equivalent.



Selecting the menu opens up this dialog, with the Tiny Link conveniently selected.



We didn't rename Tiny Links to Permalinks, as has been suggested. What do you think, should we?



This is the final Chris Broadfoot 20% production.

Release Notes 3.1-m3 ("Milestone 3")

Milestone release advisory

Do not use this release to upgrade your production systems.

For all production use and testing of Confluence, please use the latest official release.
This release is a public **development release** ('milestone') leading up to Confluence 3.2, which will probably ship in mid March 2010. Development releases are a snapshot of our work in progress, allowing our customers and especially plugin developers to see what we're up to.

Who should upgrade?

Please note the following

- **Development releases are not safe** — Development releases are snapshots of the ongoing Confluence development process. As such:
  - While we try to keep these releases stable, they have not undergone the same degree of testing as a full release.
  - Features in development releases may be incomplete, or may change or be removed before the next full release.
- **No upgrade path** — Because development releases represent work in progress, we **can not** provide a supported upgrade path between development releases, or from any development release to the eventual final release. Thus, it is possible that you will not be able to migrate any data you store in a Confluence development release to a future Confluence release.

In supplying milestone releases, our aim is to provide plugin developers with an opportunity to see the latest changes in the code.

Each milestone release has passed all our automatic tests and has been used for one week on our official internal Confluence server. Most of the issues solved have been reviewed too, and all of our milestone releases even have been performance tested for a while.

However, since our milestones releases are timeboxed (i.e. they get released every two weeks, no matter how far we have come implementing features and bugfixes), there is always a chance that we have new known bugs that are scheduled to be fixed in the next milestone, or completely new bugs unknown even to us.

Additionally, we have not completed our performance testing and compatibility testing for databases and application servers. So, for example, a milestone release may behave well on a small installation but show severe problems when subjected to many users.

Upgrade Procedure

Follow the normal upgrade instructions to upgrade from Confluence 3.1.x to this release. *We strongly recommend that you backup your Confluence home directory and database before upgrading!*

Downloads

All development releases are available from the development releases page on the Atlassian website.

Maintenance mainly

Milestones 2 and 3 don't have many feature-changes, because a lot of the development happens on branches and under the hood. We have been doing some UI polish, but nothing large is to be seen yet.

We did however split up resource-loading, so CSS is now at the top of the page, and JS is at the bottom, resulting in faster page views. If you are a plugin author using resource bundles and JS/CSS (which you should, obviouslu), then definitely check out this milestone to see if your plugin works as expected.

Stay tuned for M4 in about two weeks, which will show off some nice new feature improvements.

Cheers,
Per and the Confluence development team

Confluence 3.1 Deprecated Code Cleanup

This document contains information for Confluence plugin developers regarding changes to the Confluence codebase in the upcoming 3.1 release.

Confluence 3.1 Deprecation Cleanup

Every major Confluence release, we clean up the codebase by removing classes and methods that have been deprecated in previous releases. These changes were included in the Milestone 2 release of Confluence 3.1. We strongly encourage plugin developers to test their plugins against our milestone releases.

If the removal of these classes or methods causes you significant problems maintaining your plugins, please let us know. We will work with you to find a workaround, or possibly reinstate the code before the final release. (Note that the longer the code in question has been deprecated, the less inclined we will be to replace it. Some of the stuff we've removed here will have been giving you compiler warnings for *years*.)

More information on our guidelines for removing deprecated code: Deprecation Guidelines

Changes in Confluence 3.1 Milestone 2



For the sake of brevity, the `com.atlassian.confluence` part of package names has been omitted.

Removed Classes

Class	Deprecated Since
.pages.BreadcrumbsManager .pages.DefaultBreadcrumbsManager	2.7
.plugin.descriptor.web.conditions.user.UserHasHistoryCondition	2.8
.util.VelocityUtils	2.0
.util.PageComparator	2.3
.util.ResourceManager com.atlassian.confluence.util.DefaultResourceManager	2.8
.event.EventListener com.atlassian.confluence.event.DeprecatedEventListenerWrapper	2.3
.servlet.download.ResourceDownload	2.10[1]
.renderer.radeox.macros.include.AbstractHttpRetrievalMacro	2.7[2]
.macro.macros.AbstractHtmlGeneratingMacro	2.7[2]
bucket.search.EntityObjectDateExtractor	1.4
bucket.container.ContainerManager bucket.container.ContainerContext	2.3

[1] While the ResourceDownload class was only deprecated in 2.10, it could only ever be used in the context of the ResourceManager which was deprecated earlier. It is very unlikely that code exists that relies on one but not the other.
[2] These Radeox macros were superseded by the equivalent V2Renderer macros in Confluence 1.4, but only marked as deprecated in the 2.7 release

Removed Constants



Since Java inlines constant references during compilation, the removal of constants should not break binary compatibility with any plugin that references them.

Class	Constant	Deprecated Since
.search.lucene.extractor.PageContentEntityObjectExtractor	PAGE_REAL_TITLE	2.8
.security.SpacePermission	ADMINISTER_CONFLUENCE_PERMISSION	2.7
.setup.Bandana.ConfluenceDaoBandanaPersister	GLOBAL_BANDANA_CONTEXT	2.8

Removed Methods

Class	Method	Deprecated Since	Notes
.core.ConfluenceActionSupport	ThemeHelper getGlobalHelper()	2.0	Support for pre-2.0 themes
	ThemeHelper getSpaceHelper()	2.0	Support for pre-2.0 themes

.core.ContentEntityManager	Iterator getRecentlyModifiedEntities(int maxResults)	2.0	use the SearchManager for this kind of query
.core.ContentEntityObject	String getRealTitle()	2.8	use `getDisplayTitle()` instead
.core.ContentPermissionManager	List getInheritedViewContentPermissions(Page page)	2.5	
.core.persistence.ContentEntityObjectDao	Iterator getRecentlyModifiedEntitiesByType(ListQuery query, int firstResult)	2.8	use the SearchManager for this kind of query
.importexport.ExportContext	DateFormatter getDateFormatter()	2.8.2	
.pages.AttachmentUtils	static File getOldContainingFolder(Attachment attachment)	2.2	
.pages.actions.ViewPageAttachmentsAction	AttachmentHelper getTargetHelper(Attachment attachment)	2.8	use getWebInterfaceContext()
.plugin.editor.Editor	String getEditorSpecificCss()	2.8	
.security.ContentPermission	String getGroup()	2.4	use getGroupName()
.security.PermissionManager	boolean isGlobalAdministrator(User user)	2.7	use isConfluenceAdministrator
.setup.BootstrapManager	boolean isConfluenceHomeValid() boolean isWebdavEnabled() void setBaseUrl(String baseUrl)	2.8	use SettingsManager
.spaces.SpaceManager	boolean isValidSpaceKey(String key)	2.3	use Space.isValidGlobalSpaceKey
	boolean.isValidPersonalSpaceKey(String key)	2.3	use Space.isValidPersonalSpaceK
	List getPages(Space space, boolean currentOnly) List getPagesStartingWith(Space space, String prefix) List getBlogPosts(Space space, boolean currentOnly) List getMail(Space space, boolean currentOnly) List getSpaces() list getSpacesCreatedOrUpdatedSinceDate(Date date)	2.3	
.spaces.Space	List getPages() List getCurrentPages() List getBlogPosts() List getCurrentBlogPosts() List getMail() List getCurrentMail()	2.3	use PageManager
.user.PersonalInformationManager	PersonalInformation getPersonalInformation(String username)	2.3	use getPersonalInformation(Us user)
.user.SearchEntitiesManager	SearchResult findGroups(TermQuery query)	2.8	use findGroupsAsList()
	SearchResult findGroups(TermQuery query, boolean filter)	2.8	use findGroupsAsList()
	SearchResult findUsers(Query query)	2.8	use findUsersAsList()
.util.GeneralUtil	String format(Date date) formatTime(Date date) formatBlogDate(Date date) formatDateTime(Date date)	2.3	use $dateFormatter
	isGlobalAdministrator(Object notUsedAnyMore, User user)	2.0	use PermissionManager or $permissionHelper
	getProperties(String resource, Class caller) getPropertiesFromFile(File file) getPropertiesFromStream(InputStream is)	2.3	use PropertyUtils

.upgrade.AbstractUpgradeTask	addError(String message) addError(Throwable throwable)	1.2	use addError(UpgradeError err

Release Notes 3.1-m4 ("Milestone 4")

Milestone release advisory

Do not use this release to upgrade your production systems.

For all production use and testing of Confluence, please use the latest official release.
This release is a public **development release** ('milestone') leading up to Confluence 3.2, which will probably ship in mid March 2010. Development releases are a snapshot of our work in progress, allowing our customers and especially plugin developers to see what we're up to.

Who should upgrade?

Please note the following

- **Development releases are not safe** — Development releases are snapshots of the ongoing Confluence development process. As such:
  - While we try to keep these releases stable, they have not undergone the same degree of testing as a full release.
  - Features in development releases may be incomplete, or may change or be removed before the next full release.
- **No upgrade path** — Because development releases represent work in progress, we **can not** provide a supported upgrade path between development releases, or from any development release to the eventual final release. Thus, it is possible that you will not be able to migrate any data you store in a Confluence development release to a future Confluence release.

In supplying milestone releases, our aim is to provide plugin developers with an opportunity to see the latest changes in the code.

Each milestone release has passed all our automatic tests and has been used for one week on our official internal Confluence server. Most of the issues solved have been reviewed too, and all of our milestone releases even have been performance tested for a while.

However, since our milestones releases are timeboxed (i.e. they get released every two weeks, no matter how far we have come implementing features and bugfixes), there is always a chance that we have new known bugs that are scheduled to be fixed in the next milestone, or completely new bugs unknown even to us.

Additionally, we have not completed our performance testing and compatibility testing for databases and application servers. So, for example, a milestone release may behave well on a small installation but show severe problems when subjected to many users.

Upgrade Procedure

Follow the normal upgrade instructions to upgrade from Confluence 3.1.x to this release. *We strongly recommend that you backup your Confluence home directory and database before upgrading!*

Downloads

All development releases are available from the development releases page on the Atlassian website.

Features!

As you may know, we have started a major hiring campaign, and this has slowed us down a little in the past few weeks. M4 however is back with a vengeance, boasting quite a few improvements. See for yourself

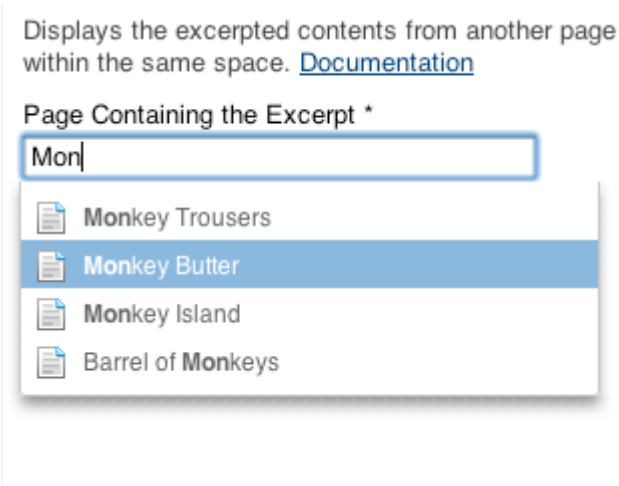
Editor Team

Move Page

We are reworking the way you can move pages around. It is now possible to move pages without editing the page, and you don't drag&drop anymore, you simply pick your parent in the tree. The next Milestones will allow you to search for a parent, and to use your recently viewed pages to move. At the moment, you can not re-order while moving the page (is this a problem though? What do you think?

Macro Browser

David Taylor's "Smart Fields" 20% project has been included in the Macro Browser although only for macro parameters that take single entries. i.e. if a macro takes a single username, space, or page then a suggestion drop down will be presented (See the screenshot). If a macro takes multiple usernames or spaces or pages then there will be no assistance.



David's "Smart Fields" work also included some custom fields for the Office Connector viewfile macros to provide assistance with referencing the MS Office document you are trying to view. Instead of separate fields for space and page name these have been combined into a singe field with a suggest drop down. The file name field has been converted to a select box showing the appropriate attachments for the selected page.

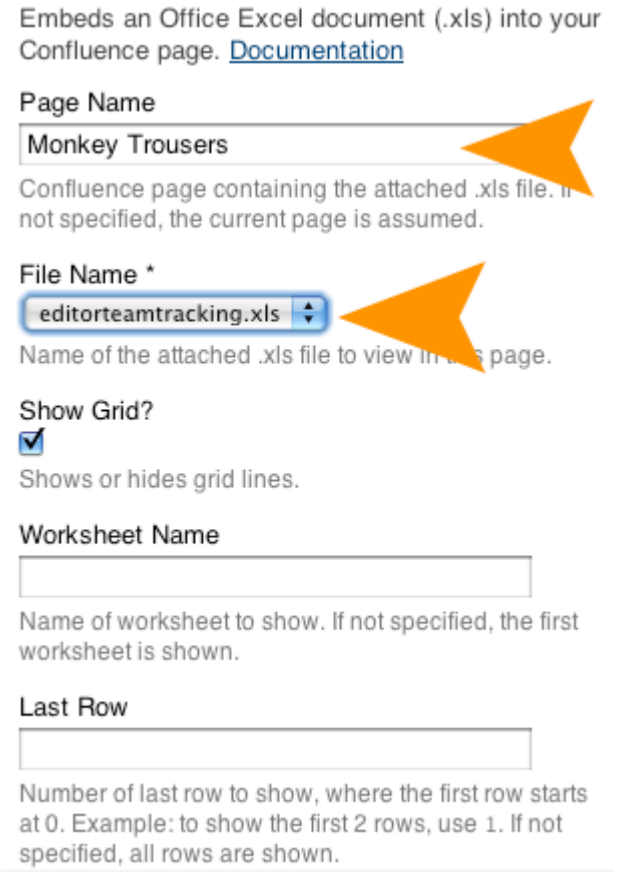


Image Browser

Some further progress has been made on the second "insert image" button on the RTE and wiki editor toolbars. You can now view and insert images that are already attached to the page.



This dialog is still missing some obvious functionality such as image configuration options and image upload hence the reason the original dialog is still present for now.

Known Issue: You should also note that in the current milestone this new dialog is only working correctly when editing pages, not when creating new pages.

Engine Room

JS/CSS rearrangements

As you saw on our previous milestone notes, we have rearranged CSS to the top and JS to the bottom of the page, making the rendering experience better. This might have consequences for your macros/plugins, so DO have a look, and tell us what you think.

REST

The Confluence REST plugin (prototype API) has been added, however it will be disabled for this release. Hopefully we will be able to enable this in next milestone when the plugin is a bit more stable.

Some space level resources have been implemented.

Getting a List of Spaces
Resource: `/space`
Description: List all spaces (maximum page size of 50) visible to the current user
Parameters:

`start-index`	start offset of the list
`max-results`	maximum number of results to return
`type`	space type

Result:

```
1.<spaces>
2.<space name="Demonstration Space" key="ds">
3.<link rel="self" href="http://localhost:8080/rest/prototype/1/space/ds"/>
4.</space>
5.</spaces>
```

Looking up Details of a space
Resource: `/space/{key}`
Description: Displays the details of the space identified by {key}
Parameters:

`expand`	set expansion options for children elements

Result:

```
01.<space name="Demonstration Space" key="ds" expand="children">
02.<link rel="self" href="http://localhost:8080/rest/prototype/1/space/ds"/>
03.<children size="2"/>
04.<home>
05.<id>32799</id>
06.<link rel="self" href="http://localhost:8080/rest/prototype/1/content/32799"/>
07.<type>page</type>
08.<title>Home</title>
09.<url>/display/ds/Home</url>
10.<children size="6"/>
11.</home>
12.</space>
```



Some of the names of the elements in the space details example are subject to change. The `children` element for example will likely be changed to something more appropriate (it refers to the top level pages of a space).

Small Improvements Team

Driving on the left shouldn't be so hard, so we've updated the left hand navigation panel.



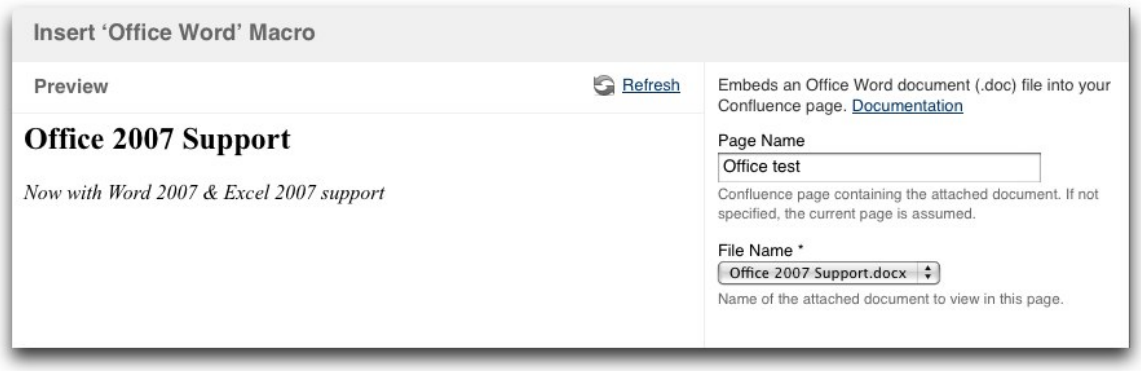
Bugfix Team

Fixed multiple bugs. Most notably:
CONF-12864 - improved performance of PageNotFound action
CONF-9575 - fixed concurrency issue that was breaking reindexing job

Gadgets & Office 2007 Team

Word 2007 & Excel 2007

We added Word 2007 support for document import and the view-file macro, as well as Excel 2007 support for the view-file macro. We don't support PowerPoint 2007 for the view-file macro yet and can't index any Office 2007 documents yet. Work in progress.



Support for importing other document formats

You can now import RTF and ODT files in addition to DOC and DOCX.

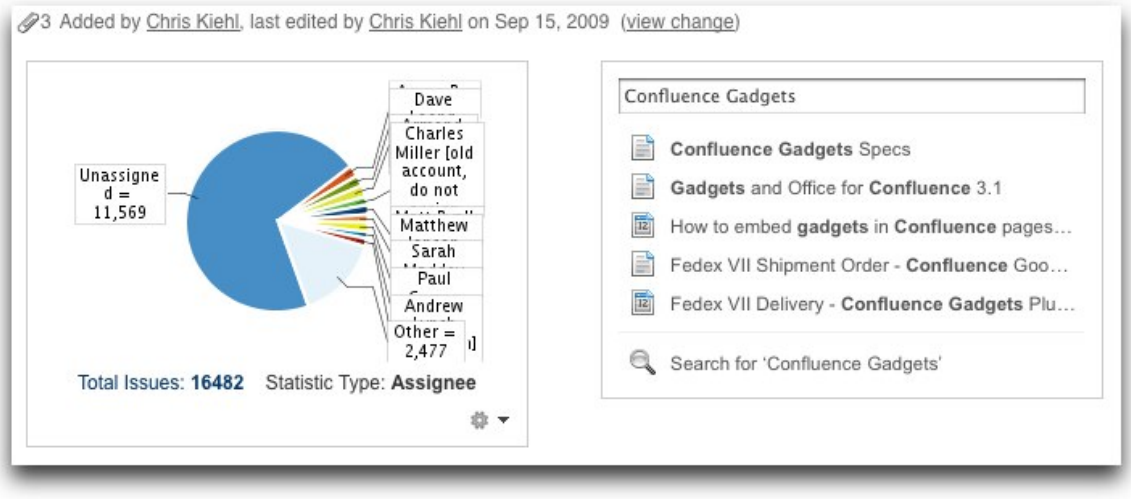
Gadget Macro

We are happy to introduce the new gadget macro! It's currently not very usable as we don't have any UI yet to change the user preferences for a gadget, which means you have to enter them manually.

Example Markup:

```
{gadget:url=[baseurl]/rest/gadgets/1.0/g/com.atlassian.confluence.plugins.gadgets:gadget-search/gadge
```

Jira chart gadget and Confluence QuickNav gadget on a Confluence page:



QuickNav gadget

We are also stoked to show off our first Confluence gadget which can be embedded in other gadget containers. The gadget URL is:

```
[baseurl]/rest/gadgets/1.0/g/com.atlassian.confluence.plugins.gadgets:gadget-search/gadgets/gadget-se
```

Unfortunately you can't embed that gadget into iGoogle or GMail yet, because we have to upgrade to a newer version of the OAuth plugin first. This will work starting in 3.1-m5.

Gadget Directory

For every gadget you want to embed on a page, it needs to be added to the gadget directory first. This is to make sure you can't just embed any third party gadgets because those might contain malicious JavaScript which will then be executed as coming from the same security domain as Confluence.

OAuth Integration

We have also integrated the OAuth admin plugin which allows you to configure your consumer information as well as adding OAuth consumers and service providers. Detailed documentation will be available on CAC in the future.

20% and miscellaneous

Thumbnail Dialogs

Thumbnail images previously opened the full image in a new popup window, when clicked. Now clicking on a thumbnail opens the image in a  fancy dialog (same as the image gallery).

Sprite Image for Macro Browser

On the edit page, the macro browser is built up in the background with its pretty macro icons. Although the macro icons are cached, there are quite a number of requests made on a cold browser cache (we have 52 macro icons, so 52 requests are made just for the macro browser). This has been reduced to one request by generating a sprite image for all the icons and using it with css and background-image positioning.

hod has been added for macros to more clearly define their desire to be kicked out of paragraphs, and code macros do that. This change is still awaiting review, so it may be temporary.

AUI upgraded to 1.2.1!

Release Notes 3.1-m5 ("Milestone 5")

Milestone release advisory

Do not use this release to upgrade your production systems.

For all production use and testing of Confluence, please use the latest official release.
This release is a public **development release** ('milestone') leading up to Confluence 3.2, which will probably ship in mid March 2010. Development releases are a snapshot of our work in progress, allowing our customers and especially plugin developers to see what we're up to.

Who should upgrade?

Please note the following

- **Development releases are not safe** — Development releases are snapshots of the ongoing Confluence development process. As such:
  - While we try to keep these releases stable, they have not undergone the same degree of testing as a full release.
  - Features in development releases may be incomplete, or may change or be removed before the next full release.
- **No upgrade path** — Because development releases represent work in progress, we **can not** provide a supported upgrade path between development releases, or from any development release to the eventual final release. Thus, it is possible that you will not be able to migrate any data you store in a Confluence development release to a future Confluence release.

In supplying milestone releases, our aim is to provide plugin developers with an opportunity to see the latest changes in the code.

Each milestone release has passed all our automatic tests and has been used for one week on our official internal Confluence server. Most of the issues solved have been reviewed too, and all of our milestone releases even have been performance tested for a while.

However, since our milestones releases are timeboxed (i.e. they get released every two weeks, no matter how far we have come implementing features and bugfixes), there is always a chance that we have new known bugs that are scheduled to be fixed in the next milestone, or completely new bugs unknown even to us.

Additionally, we have not completed our performance testing and compatibility testing for databases and application servers. So, for example, a milestone release may behave well on a small installation but show severe problems when subjected to many users.

Upgrade Procedure

Follow the normal upgrade instructions to upgrade from Confluence 3.1.x to this release. *We strongly recommend that you backup your Confluence home directory and database before upgrading!*

Downloads

All development releases are available from the development releases page on the Atlassian website.

Gadgets & Office Team

Gadgets in the Macro Browser

Want to embed one of those fancy new JIRA gadgets into your page? This has become really easy with the new Macro Browser integration. Just have a look in the "External Content" category to find available gadgets. Confluence comes with the Confluence QuickNav gadget by default, but the administrator can add as many gadgets as you like.



Configure your JIRA gadgets in the Macro Browser

You can now configure your JIRA gadgets using the macro browser. Just pick a JIRA gadget from the Macro Browser, configure it and add it to your Confluence page or blog post.





Office 2007 support

We've fixed a couple of minor issues with the Word 2007 integration. But the really good news is that we've started working on PowerPoint 2007 support! Nothing to see yet, but stay tuned.

Dialogs Team

The new image browsing dialog has been restyled and as a bonus now works in new pages and blogs.

Engine Room

The Confluence REST plugin (prototype API) has now been enabled for this milestone. In addition to spaces (in previous milestone) we now have an API for viewing Confluence content (pages, blogs, comments) with page children expansion supported.

Some very basic documentation can be found here: Prototype REST API

Bugfix Team

Fixed multiple bugs .

Misc

Faster editor load

With a few improvements we have made to the way we load TinyMCE, it should appear faster across all the browsers.

This is an Agnes Ro 20% production

Prototype REST API

This page is obsolete
Please refer to the REST API documentation in the Confluence developer space.

This page documents the Prototype REST API supported for Confluence 3.1.

Space

Getting a List of Spaces

Resource: `/space`
Description: List all spaces (maximum page size of 50) visible to the current user
Parameters:

`start-index`	start offset of the list
`max-results`	maximum number of results to return
`type`	space type e.g. personal, all
`expand`	use 'space' to expand the spaces details of each space listed

Result:

```
01.<spaceEntityList expand="space">
02.<space name="FOO" key="FOO">
03.<link rel="self" href="http://localhost:8080/confluence/rest/prototype/1/space/FOO"/>
04.<rootpages size="1"/>
05.</space>
06.<space name="Test Space" key="TST">
07.<link rel="self" href="http://localhost:8080/confluence/rest/prototype/1/space/TST"/>
08.<rootpages size="2"/>
09.</space>
10.</spaceEntityList>
```

Looking up Details of a space

Resource: `/space/{key}`
Description: Displays the details of the space identified by {key}
Parameters:

`expand`	use 'rootpages' to list all the root level pages of the space

Result:

```
01.<space name="Demonstration Space" key="ds" expand="rootpages">
02.<link rel="self" href="http://localhost:8080/rest/prototype/1/space/ds"/>
03.<rootpages size="2"/>
04.<home>
05.<id>32799</id>
06.<link rel="self" href="http://localhost:8080/rest/prototype/1/content/32799"/>
07.<type>page</type>
08.<title>Home</title>
09.<url>/display/ds/Home</url>
10.<children size="6"/>
11.</home>
12.</space>
```

Content

Getting Content

Resource: `/content/{id}`
Description: Gets confluence content by id
Parameters:

`expand`	use 'children' to list all the child pages of the page

Result:

```
1.<content type="page" id="1180211" expand="children">
2.<link rel="self" href="http://localhost:8080/confluence/rest/prototype/1/content/1180211"/>
3.<title>Home</title>
4.<url>/display/FOO/Home</url>
5.<spaceKey>FOO</spaceKey>
6.<children size="1"/>
7.<body>This is the home of the FOO space.</body>
8.</content>
```

Release Notes 3.1-m6 ("Milestone 6")

Milestone release advisory

Do not use this release to upgrade your production systems.

For all production use and testing of Confluence, please use the latest official release.
This release is a public **development release** ('milestone') leading up to Confluence 3.2, which will probably ship in mid March 2010. Development releases are a snapshot of our work in progress, allowing our customers and especially plugin developers to see what we're up to.

Who should upgrade?

Please note the following

- **Development releases are not safe** — Development releases are snapshots of the ongoing Confluence development process. As such:
  - While we try to keep these releases stable, they have not undergone the same degree of testing as a full release.
  - Features in development releases may be incomplete, or may change or be removed before the next full release.
- **No upgrade path** — Because development releases represent work in progress, we **can not** provide a supported upgrade path between development releases, or from any development release to the eventual final release. Thus, it is possible that you will not be able to migrate any data you store in a Confluence development release to a future Confluence release.

In supplying milestone releases, our aim is to provide plugin developers with an opportunity to see the latest changes in the code.

Each milestone release has passed all our automatic tests and has been used for one week on our official internal Confluence server. Most of the issues solved have been reviewed too, and all of our milestone releases even have been performance tested for a while.

However, since our milestones releases are timeboxed (i.e. they get released every two weeks, no matter how far we have come implementing features and bugfixes), there is always a chance that we have new known bugs that are scheduled to be fixed in the next milestone, or completely new bugs unknown even to us.

Additionally, we have not completed our performance testing and compatibility testing for databases and application servers. So, for example, a milestone release may behave well on a small installation but show severe problems when subjected to many users.

Upgrade Procedure

Follow the normal upgrade instructions to upgrade from Confluence 3.1.x to this release. *We strongly recommend that you backup your Confluence home directory and database before upgrading!*

Downloads

All development releases are available from the development releases page on the Atlassian website.

Overview

We are pretty close to Beta now, so if you are a plugin author and haven't checked out the latest milestones of 3.1, you should do it now. The API won't change much anymore, from now on it's mainly bugfixing.

Gadgets & Office Team

Editing user preferences of OpenSocial gadgets in the Macro Browser

You can now edit user preferences of OpenSocial gadgets directly in the Macro Browser. Let's take our all time favourite Hamster gadget as an example:



> ⚠ The UI will be slightly different in the next milestone, as we are going to move the form fields in the center of to the right hand side and make them part of the normal macro parameters.

Gadget Directory

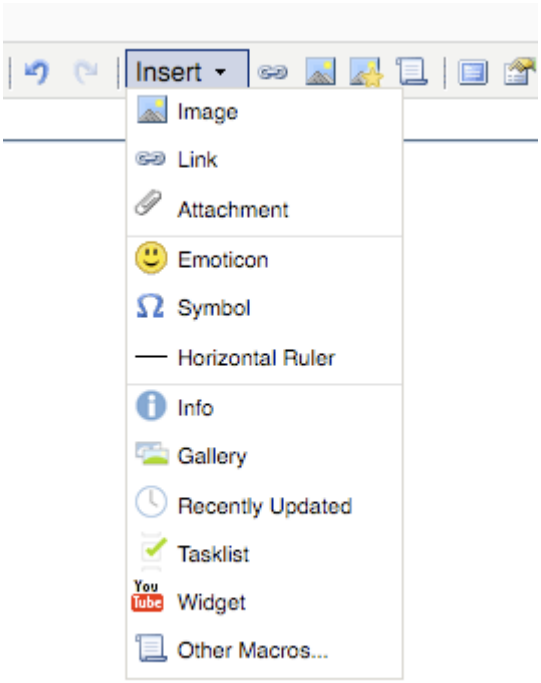
You can now discover which gadgets are exposed by Confluence and can be used in other OpenSocial containers like JIRA 4 or iGoogle. There is a new "Gadget Directory" menu item in the browse menu which allows you to browse through the list of available gadgets.





Dialogs Team

There's a new "Insert" menu in the editor that helps you insert links, attachments, images and some default macros without having to already know which little button you're meant to press. (On IE the macro icons look odd, that's a bug and will be adressed soon)



Help us choose the right macros!

We picked a few macros that we thought users should know about, which are simple enough to use right away, and which represent a decent spread across a range of macros to stimulate your interest in the "add more" buttton. Our constraint is 5 macros max, since the menus gets too long on ~~stupid~~ ~~goddamn~~ good old IE6 on a 1024x768 screen.

Apart from the edit menu, some bug fixes of note: [Problems with links] in the RTE have been fixed; pressing escape in the new image dialog doesn't prevent it working next time; and move page works in IE8.

A lot of work has gone into the new page move dialog but it's not quite ready, so we didn't put it into M6.

Engine Room

A setting (under General Config) to serve the Javascripts back in the header has been introduced. This has been disabled by default for this milestone but will eventually be turned on as the default for 3.1. Reasoning behind this is that quite a few plugins have been broken and unusable by moving the scripts to the bottom of the page.

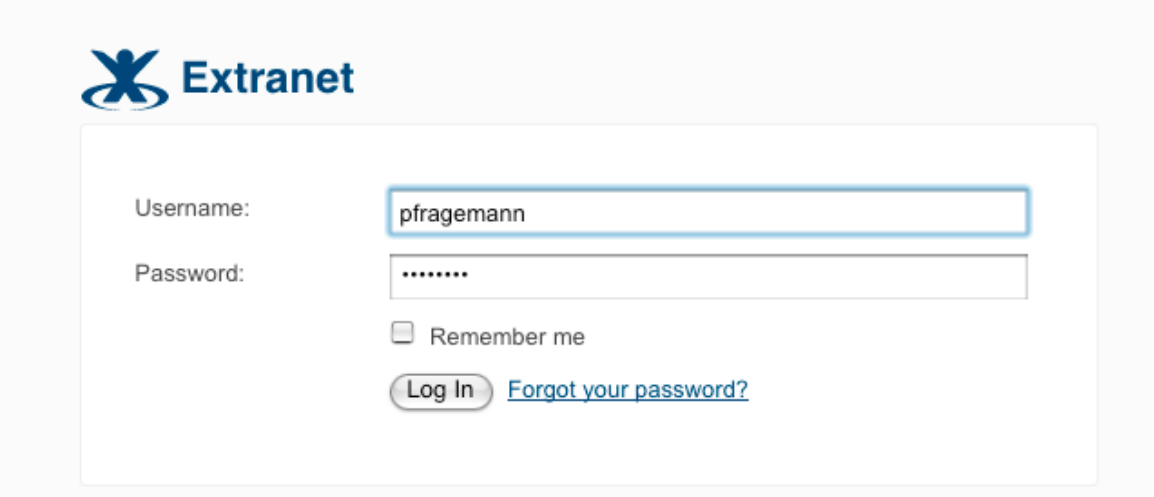
Work has been continued on the REST API however these were purely back end changes and the API itself has not changed since m5.

Bugfix Team

CONF-17171 - made selecting a page version on a page history view easier
General bug fixing.

Small Improvements

New Login Screen design, which looks more consistent with other products like JIRA. Another step forward in terms of making Confluence look better



Known Issues

Errors were reported by the JIRA trusted connection.

- APP_UNKNOWN; Unknown Application: {0}; ["confluence:4557196"]

JIRA Issues (6 issues)										
Type	**Key**	**Summary**	**Assignee**	**Reporter**	**Priority**	**Status**	**Resolution**	**Created**	**Updated**	**Due**
	CONF-17444	The move-page dialog can be cancelled with ESC but the move operation is still executed	Unassigned	Per Fragemann [Atlassian]		Open	Unresolved	Nov 02, 2009	Dec 08, 2009	
	CONF-17332	Can click Preview multiple times on "From the web" tab in image dialog, produces many spinners	Unassigned	Per Fragemann [Atlassian]		Open	Unresolved	Oct 26, 2009	Dec 08, 2009	
	CONF-17329	When Adding or Editing a page, the top-right 'Save' and 'Cancel' buttons are partially overlapped by the toolbar in Wiki Markup mode.	Unassigned	Giles Gaskell [Atlassian]		Open	Unresolved	Oct 26, 2009	Dec 08, 2009	
	CONF-17235	Error is given on IE7 and 8 when clicking on an uploaded image	Dave Loeng [Atlassian]	Mark Hrynczak [Atlassian]		Open	Unresolved	Oct 14, 2009	Mar 08, 2010	
	CONF-17229	Drag and drop should do permission check before it uploads files	Unassigned	Mark Hrynczak [Atlassian]		Open	Unresolved	Oct 14, 2009	Mar 02, 2010	
	CONF-17226	Drag & Drop plugin shows incorrect file name during upload for files containing '<'	Unassigned	Mark Hrynczak [Atlassian]		Open	Unresolved	Oct 14, 2009	Dec 08, 2009	

Release Notes 3.1-m7 ("Milestone 7")

Milestone release advisory

Do not use this release to upgrade your production systems.

For all production use and testing of Confluence, please use the latest official release.
This release is a public **development release** ('milestone') leading up to Confluence 3.2, which will probably ship in mid March 2010. Development releases are a snapshot of our work in progress, allowing our customers and especially plugin developers to see what we're up to.

Who should upgrade?

Please note the following

- **Development releases are not safe** — Development releases are snapshots of the ongoing Confluence development process. As such:
  - While we try to keep these releases stable, they have not undergone the same degree of testing as a full release.
  - Features in development releases may be incomplete, or may change or be removed before the next full release.
- **No upgrade path** — Because development releases represent work in progress, we **can not** provide a supported upgrade path between development releases, or from any development release to the eventual final release. Thus, it is possible that you will not be able to migrate any data you store in a Confluence development release to a future Confluence release.

In supplying milestone releases, our aim is to provide plugin developers with an opportunity to see the latest changes in the code.

Each milestone release has passed all our automatic tests and has been used for one week on our official internal Confluence server. Most of the issues solved have been reviewed too, and all of our milestone releases even have been performance tested for a while.

However, since our milestones releases are timeboxed (i.e. they get released every two weeks, no matter how far we have come implementing features and bugfixes), there is always a chance that we have new known bugs that are scheduled to be fixed in the next milestone, or completely new bugs unknown even to us.

Additionally, we have not completed our performance testing and compatibility testing for databases and application servers. So, for example, a milestone release may behave well on a small installation but show severe problems when subjected to many users.

Upgrade Procedure

Follow the normal upgrade instructions to upgrade from Confluence 3.1.x to this release. *We strongly recommend that you backup your Confluence home directory and database before upgrading!*

Downloads

All development releases are available from the development releases page on the Atlassian website.

Not quite a beta yet!

While this release is almost feature complete there are still one major technical improvement (superbatching) and one feature (the dashboard-widget) outstanding, so we can't call it beta just yet. Apart from that though, M7 is feature complete. Anything that you don't see here, you won't see in the final product. We will keep fixing bugs (yes, there are still quite a few unfortunately), but if something else in the functionality really annoys you, please raise it with us now! You can add a comment, or you can send us a mail, or of course raise JIRA issues.

The files are not available for download at the moment, they will be added on Thursday. This page is a teaser, and for use by our translators.

Gadgets & Office Team

Indexing and Searching of Office 2007 files

The contents of Confluence attachments created by PowerPoint 2007, Excel 2007, and Word 2007 are now fully searchable inside Confluence. This includes files with pptx, xlsx, and docx extensions as well as other extensions like potx (PowerPoint template) and dotx (Word template).



PowerPoint 2007 previews

You can now preview PowerPoint 2007 files directly in Confluence. We're proud to say that this completes the goal of supporting the new 2007 file formats for the three main Microsoft Office products (Word, Excel, and PowerPoint) in Confluence 3.1. This is in addition to already supporting the legacy binary format.



Keyboard navigation of full-screen PowerPoint and PDF slideshows

If you are using Flash player 10 or higher, you will no longer see any controls when previewing Powerpoint or PDF files in full screen mode. Instead, you must use the spacebar or the arrow keys to move between the slides. Press escape to exit full screen mode.



Launch Microsoft Office from the Attachments screen.

Previously, you had to preview a file to get the option to edit it. We've added a new 'Edit' option to each attachment that will launch the appropriate desktop editor for that attachment. The old 'Edit' option has been renamed to 'Properties'.



Editing user preferences of OpenSocial gadgets in the Macro Browser

As promised in the last release notes we moved the editing of non hidden user preferences to the usual location for macro parameters on the right hand side:



Activity Stream Gadget

We've added an activity stream gadget for Confluence. This means you can now have your Confluence Activity on your JIRA dashboard or on a Confluence page! You can also directly comment from the activity stream.

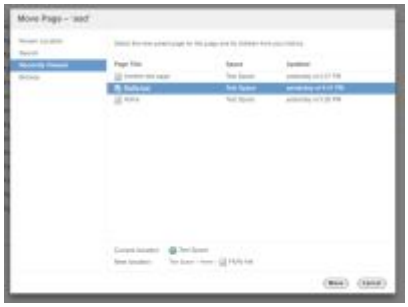
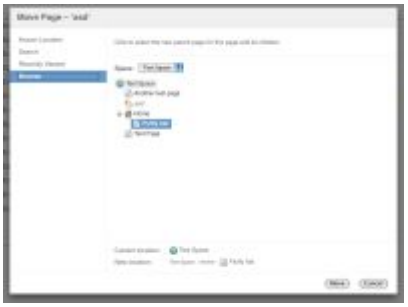
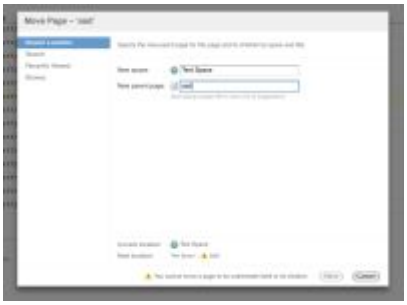


Dialogs Team

Page move dialog

We've added new tabs and improved the design of the new Move Page dialog based on feedback from earlier iterations. Some key improvements:

- A quick way of specifying a known location in the wiki using space and page title autocomplete (quick-nav style)
- Search and Recently Viewed tabs for locating a parent page in different ways
- An improved tree component which gives better loading feedback
- Fixed description panels so helpful text doesn't disappear when you scroll the tree
- A new location panel in the dialog so you know where you're moving a page from and to
- Immediately highlight problems with the new location, such as attempting to move a page beneath itself or its children.

The move dialog is feature-complete for this release, but we still have a few remaining improvements to make. In particular, scrolling the tree to the correct location when it opens and making the tree completely undraggable will be fixed for the next release.

Image browser

The new image browser dialog now supports uploading files. The upload proceeds faster than before (because the entire window doesn't need reloading), and you'll get a thumbnail of your image preselected, ready to insert.

We've fixed the bugs with image borders that prevented them working properly in the rich text editor, and enabled borders by default for newly inserted images. The dialog also includes some simple keyboard navigation that should make it more intuitive to interact with. All features of the image dialog have been tested with all our supported browsers.

We also spent some time on improved validation of uploaded attachments and thumbnail generation. The old insert image window is now disabled, so please let us know if you have any problems with the new functionality.

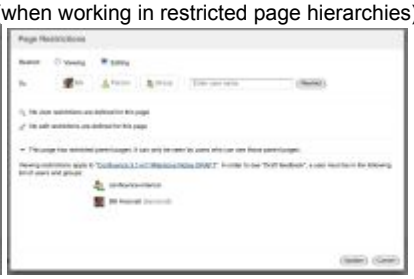


The image dialog is feature-complete for this release, but also has a few issues which will be fixed before the final release. Better handling of server outages and not being able to attach due to permissions is coming in the next release.

Page "Permissions" dialog

We have reverted to the two row layout to avoid permission-restriction-inconsistencies, and fixed a few bugs along the way. Most notably, inherited restrictions are now displayed separately from the current level's restrictions, and hidden by default. This avoids cluttering the page with stuff you usually know about already (when working in restricted page hierarchies).





Engine Room

Nothing big went into M7, the super-batching work is ongoing and will mean faster page loads from 3.0-beta1 onwards.

Small Improvements

The footer is now always at the bottom at the viewport, rather than the bottom of the content. This makes the layout a lot cleaner on short pages:



Open bugs

Plenty.

What's next?

Beta 1 is due next week

Confluence 3.1 Beta 1 Release Notes

> **Confluence 3.1 Beta 2 has been released!**
>
> Please refer to the Confluence 3.1 Beta 2 Release Notes for updated information on the enhancements available in this version.

Confluence 3.1 Beta 1 is a public **development release** ("Beta") leading up to the official release of Confluence 3.1, which we aim to ship in Q4/2009.

Both "Milestone" and "Beta" versions of Confluence are development releases, which are preliminary releases leading up to the official release of a major Confluence version. They are a snapshot of our work in progress and provide an advance preview of new features to our customers and the general public. Confluence plugin developers can also use development releases to test and fix their plugins in advance of an official release.

The main distinction between a beta and a milestone release is that milestone releases typically acquire new features with each subsequent milestone version, whereas beta releases are predominantly feature-complete. Beta releases still undergo bug fixing and occasionally, existing features may be enhanced or added in subsequent beta versions.

Do not use in production

Beta releases should not be used in production environments as they are not officially supported.

For all production use and testing of Confluence, please use the latest official release.

Who should try this out?

With beta releases, the Confluence development team aims to provide plugin developers with an opportunity to see the latest changes in the code.

Furthermore, if you are a Confluence customer who is eager to see the new features and provide us with feedback on our upcoming major release, we encourage you to try out our beta releases.

Please note the following

- **Development releases are not safe** — Development releases are snapshots of the ongoing Confluence development process. As such:
  - While we try to keep these releases stable, they have not undergone the same degree of testing as a full release.
  - Features in development releases may be incomplete, or may change or be removed before the next full release.
- **No upgrade path** — Because development releases represent work in progress, we **can not** provide a supported upgrade path between development releases, or from any development release to the eventual final release. Thus, it is possible that you will not be able to migrate any data you store in a Confluence development release to a future Confluence release.

Each beta release has passed all our automated tests, has undergone some performance testing and has been used for one week on our official internal Confluence server. Furthermore, most of the solved issues have been reviewed.

Be aware that our beta releases are still undergoing final performance and compatibility testing for databases and application servers. Hence, we recommend that you use beta releases on installations with small (as opposed to full production-level) user bases.

Upgrade Procedure

If you wish to upgrade your existing Confluence installation with this version, ensure you have created a separate copy of your current Confluence production installation *first* and using that copy, follow the normal upgrade instructions to upgrade it to this beta release. If you have also implemented customised site- or space-specific layouts, you will need to re-implement them after the upgrade. Otherwise, some of the new features in Confluence (or possibly existing features) may not function correctly.

Downloads

All development releases are available from Development Releases on the Atlassian website.

Known Issues

There are several bugs outstanding which will be resolved before the official release. Some bugs you are most likely to come across are:

- **Page Preview** — When editing a page or blog post and you click the '**Preview**' tab, only the first section of content is shown. You will not be able to scroll down to view the remaining content on a page.
- **Widget Connector Icon** — The current icon has not yet been finalised and will change before the official release of Confluence 3.1.
- **Quick Navigation Aid** — The Quick Navigation Aid does not work in Internet Explorer 6.0.

Refer to our JIRA site for a list of Confluence 3.1-specific bugs.

Highlights of Confluence 3.1

Highlights of this Release:

- Introducing Gadgets
- Office 2007 Support
- New 'Move Page' Feature
- New Image Browser
- Draft Comparisons
- New Page Restrictions Dialog Box
- New Rich Text Editor Insert Menu
- Other Editor Enhancements
  - Macro Browser Smart Fields
  - Edit Mode Exit Notification
- Other Improvements

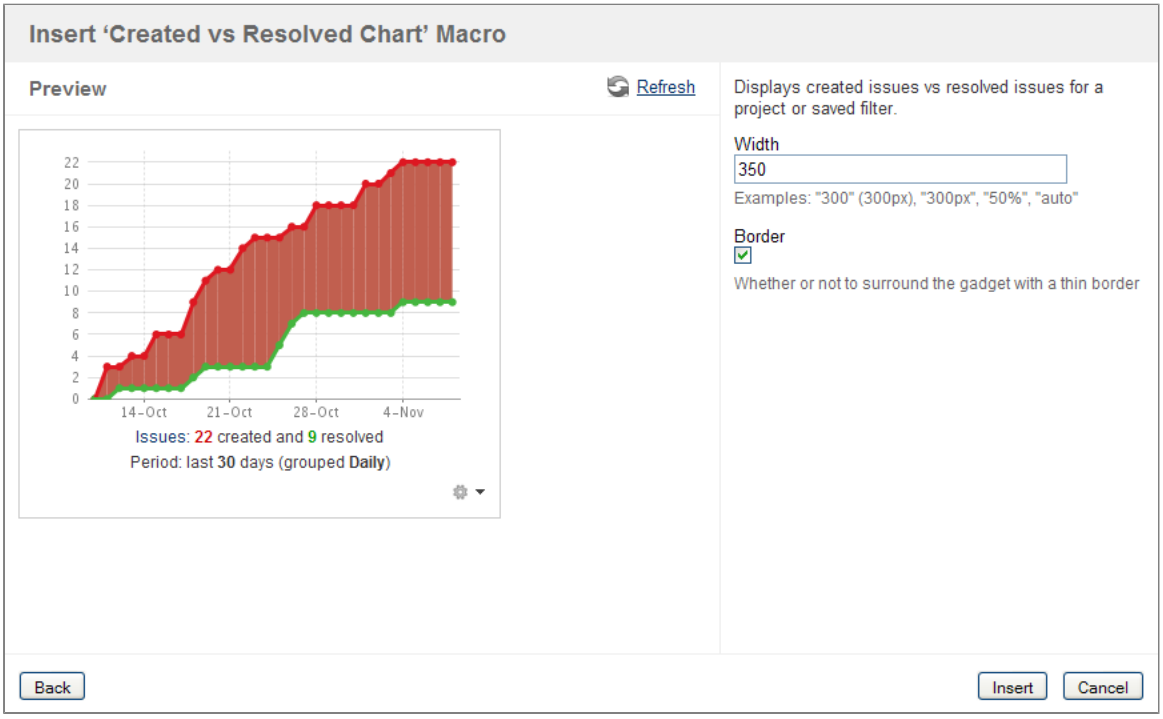


Introducing Gadgets

Gadgets are small objects that offer dynamic content and functionality which may be served by any OpenSocial-compliant web application, such as JIRA 4.0+, the same or another Confluence installation, or non-Atlassian applications such as iGoogle and Gmail.

- Confluence supports the use of gadgets in pages and blog posts, which are accessible through the macro browser.
- Confluence can also serve its own gadgets, for use in any other OpenSocial-compliant web application including the same or another Confluence installation. Two such gadgets are bundled with Confluence:
  - **Activity Stream** — This gadget shows a list of recent activities that have occurred on the Confluence server, such as the addition of new pages, blog posts or comments, content edits, status updates and so on.
  - **Quick Navigation Aid** — This gadget provides heading and content search capabilities on a Confluence server.

Inserting a JIRA Gadget onto a Confluence Page

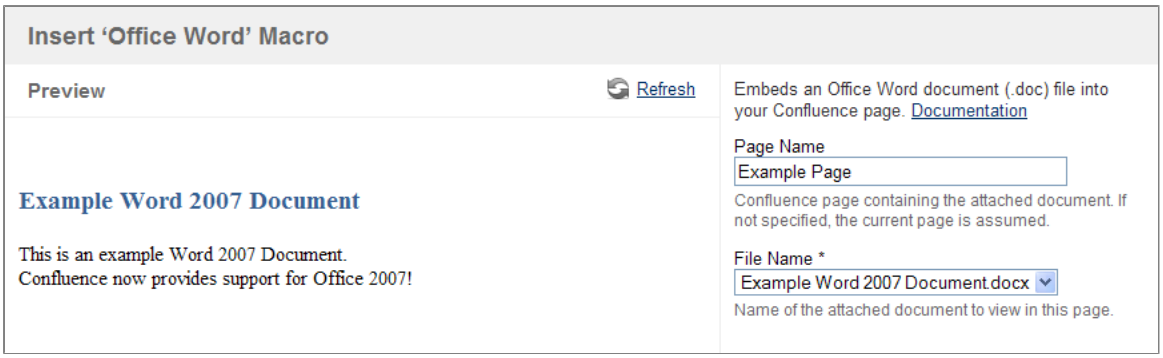




Office 2007 Support

Confluence now provides full support for Office 2007 files, allowing you to view and edit content from Microsoft Word 2007 (`.docx` and `.dotx`), PowerPoint 2007 (`.pptx` and `.potx`) and Excel 2007 (`.xlsx`) files.

- Along with existing Microsoft Office versions, Confluence now fully indexes Microsoft Office 2007 files and their content can be searched by Confluence.
- Office files can be edited directly from any page's or blog post's list of attachments.
- Using Confluence's Office connector macros, you can insert Word, PowerPoint or Excel 2007 files directly into your Confluence page.

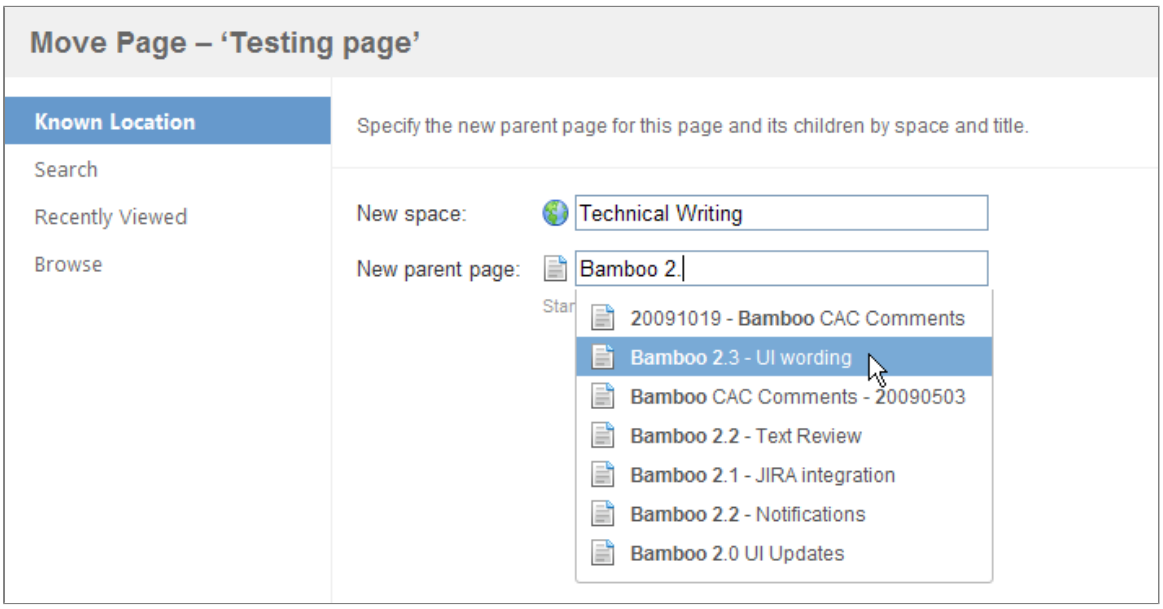




New 'Move Page' Feature

Confluence introduces a new page moving feature, that easily allows you to move the page you are currently viewing, adding or editing to another page elsewhere in the same or another space of your Confluence site. This feature is available through a new 'page move' dialog box, which provides the following flexible methods for moving pages:

- **Known Location** – Allows you to type the name of a space and within that space, the 'parent' page under which to move your page.
- **Search** – Allows you to search for a 'parent' page (within a selected space or set of spaces) under which to move your page.
- **Recently Viewed** – Allows you to select one of your recently viewed pages to be the 'parent' of your page to be moved.
- **Browse** – Allows you to select a space and page (within the tree of pages in the space) that will be the 'parent' of your page to be moved.
  For more information, refer to Moving a Page.

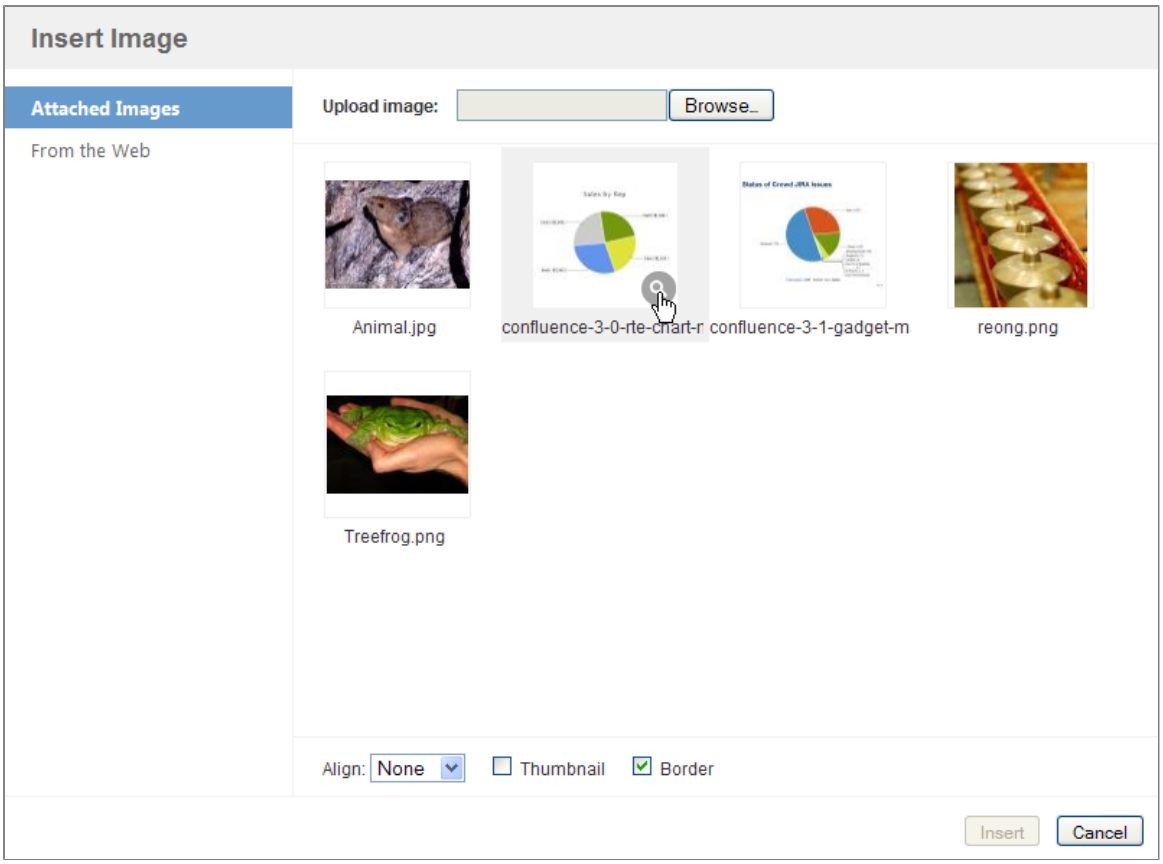




New Image Browser

A new 'Image Browser' has been introduced to replace the old 'Insert Image' dialog box. The image browser provides a less-cluttered and enhanced interface that allows you to:

- Hover over any image in the browser and expand the image (to preview it in detail) by clicking its 'magnifying glass' icon in the lower-right corner, before inserting it onto a page.
- Add a link to an image elsewhere on the web via the image's URL.



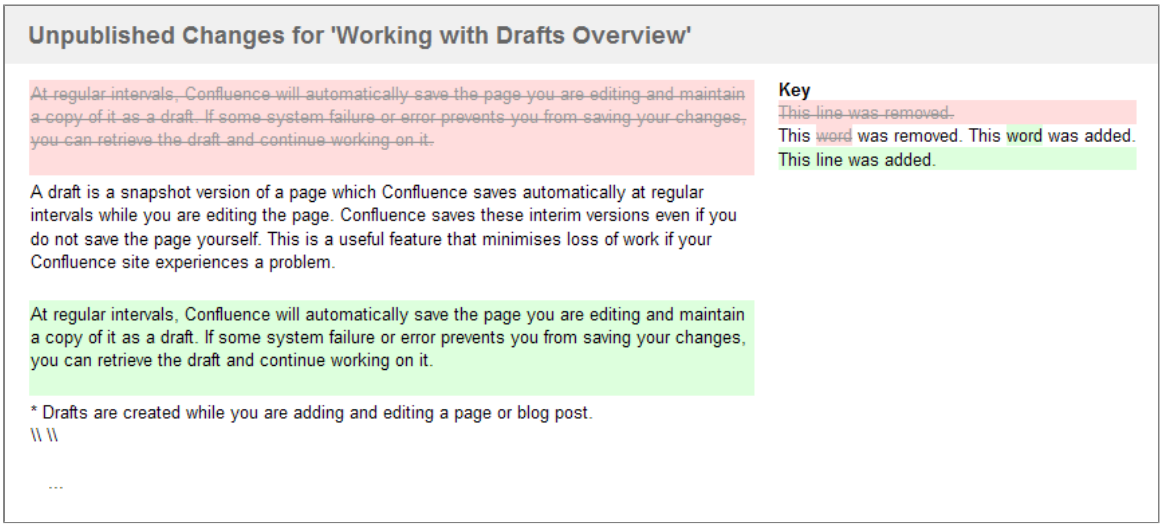


Draft Comparisons

Confluence's drafts features have been improved, such that you can now view any of your unsaved draft changes before deciding to resume editing them. This nifty feature comes in handy, particularly when other people have made subsequent changes to a page or blog post in your drafts list and you need to merge changes or resolve a conflict.



Title	Last Saved Date	Operations
Page with Partial Content	less than a minute ago	Resume Editing | Discard
Page with Content to Merge	2 minutes ago	View Changes | Merge and Resume Editing | Discard
Page with a Conflict	12 minutes ago	View Changes | View Conflict | Discard

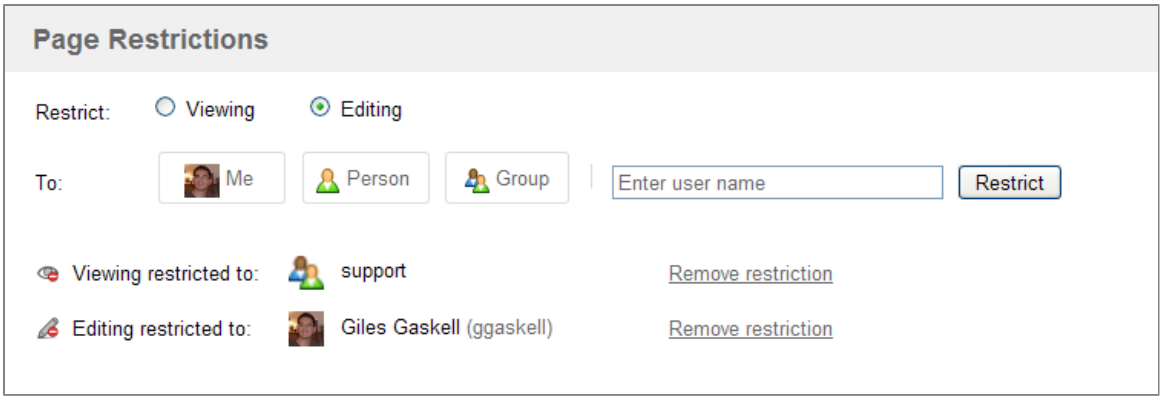




New Page Restrictions Dialog Box

Confluence's page restrictions feature has been incorporated into a convenient and accessible dialog box that is now easier to use than before.

- The page restrictions dialog box can be accessed from the padlock icon or the '**Tools**' -> '**Restrictions**' menu item on any Confluence page. From this dialog box, you can view all viewing and editing restrictions associated with the page you are viewing. You no longer need to view the page's associated 'Info' page in order to view these restrictions.
- It is no longer necessary for a page to be in edit mode in order to modify its page restrictions. You can edit all page restrictions from this dialog box.
- The page restrictions dialog box is still accessible when when a page is in edit mode.

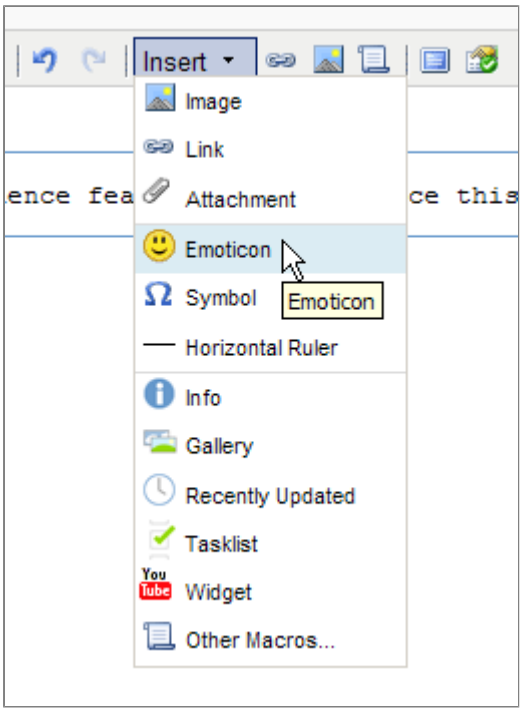




New Rich Text Editor Insert Menu

Confluence's rich text editor now combines a number of its commonly used editing features into a new convenient '**Insert**' menu.

- The Horizontal line, Insert Symbol and Insert Emoticon Toolbar icons have been moved into the new insert menu.
- The functionality to insert images, links or attachments into a page can now also be accessed from this menu.
- The macro browser, as well as a number of commonly-used macros are conveniently accessible from this menu too.

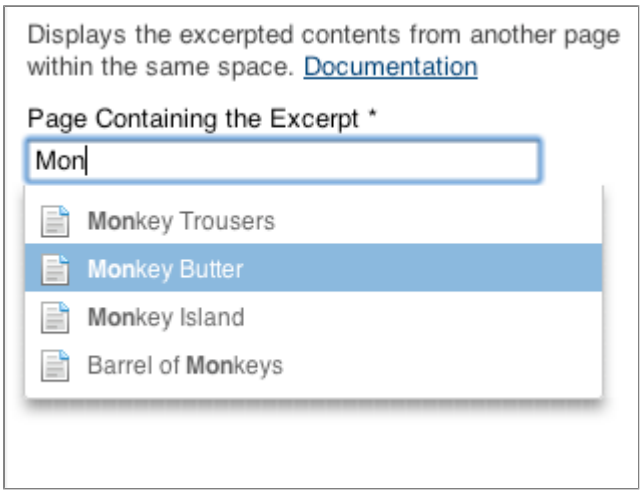




Other Editor Enhancements

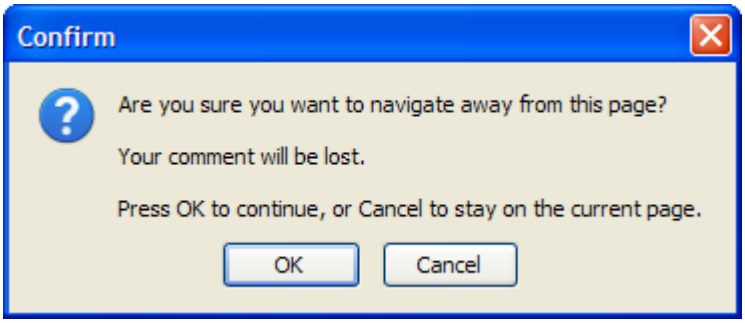
Macro Browser Smart Fields

When using the Macro Browser, an 'auto-complete' feature is now provided on any parameters that require the entry of a single item, such as a page title, username or space key. This greatly facilitates the customisation of macros and minimises the need to know the exact item names in advance.



Edit Mode Exit Notification

Whenever you add or edit a page, comment or blog post and then click onto another Confluence feature that navigates away from your unsaved content, a message box appears, warning that your content will be saved as a draft (if it is a page) or lost (if a comment). This allows you to cancel out of this action if it was accidental.





Other Improvements

Other small enhancements and improvements to Confluence include:

- Support for Internet Explorer 8 — Confluence 3.1 now fully supports Internet Explorer 8, released around mid-2009.
- Support for OAuth — With the introduction of gadgets (above) in this release, Confluence 3.1 now allows you to establish OAuth relationships with other web applications such as JIRA 4.0+, iGoogle, Gmail etc., thereby allowing them to share resources via gadgets.
- New 'Link to this page' feature — If you wish to link to a Confluence page from any other location on the web, use the convenient '**Link to this Page**' feature (available from any page's or blog post's '**Tools**' menu). Upon selecting this feature, the 'Link to this Page' dialog box opens, from which you can copy three versions of the link to embed elsewhere:
  - **Link** – Standard URL which should work from any other accessible location on the web.
  - **Tiny Link** – A reduced-length version of the 'Link', which can be used in text fields of limited length, such as tweets or Confluence Status Updates.
  - **Wiki Markup** – A wiki markup version of the link, which can be used in any other location within your Confluence site.
- 'More' links on activity streams — 'More' links have been added to various activity streams throughout the Confluence interface, including the profile sidebar, a user's profile page and via the recently updated macro. Clicking on a 'More' link expands the list of results, providing a convenient means of accessing progressively more distant user activities.
- New Log In and Log Out screens.

- In an aim to minimise confusion, 'News Items' are now consistently called 'Blog Posts' throughout the Confluence interface and a list of blog posts is collectively referred to as a 'Blog'.
- Other minor interface improvements.

Confluence 3.1 Beta 2 Release Notes

Confluence 3.1 Beta 2 is a public **development release** ("Beta") leading up to the official release of Confluence 3.1, which we aim to ship in Q4, 2009.

Both "Milestone" and "Beta" versions of Confluence are development releases, which are preliminary releases leading up to the official release of a major Confluence version. They are a snapshot of our work in progress and provide an advance preview of new features to our customers and the general public. Confluence plugin developers can also use development releases to test and fix their plugins in advance of an official release.

The main distinction between a beta and a milestone release is that milestone releases typically acquire new features with each subsequent milestone version, whereas beta releases are predominantly feature-complete. Beta releases still undergo bug fixing and occasionally, existing features may be enhanced or added in subsequent beta versions.



Do not use in production

Beta releases should not be used in production environments as they are not officially supported.

For all production use and testing of Confluence, please use the latest official release.

Who should try this out?

With beta releases, the Confluence development team aims to provide plugin developers with an opportunity to see the latest changes in the code.

Furthermore, if you are a Confluence customer who is eager to see the new features and provide us with feedback on our upcoming major release, we encourage you to try out our beta releases.



Please note the following

- **Development releases are not safe** — Development releases are snapshots of the ongoing Confluence development process. As such:
  - While we try to keep these releases stable, they have not undergone the same degree of testing as a full release.
  - Features in development releases may be incomplete, or may change or be removed before the next full release.
- **No upgrade path** — Because development releases represent work in progress, we **can not** provide a supported upgrade path between development releases, or from any development release to the eventual final release. Thus, it is possible that you will not be able to migrate any data you store in a Confluence development release to a future Confluence release.

Each beta release has passed all our automated tests, has undergone some performance testing and has been used for one week on our official internal Confluence server. Furthermore, most of the solved issues have been reviewed.

Be aware that our beta releases are still undergoing final performance and compatibility testing for databases and application servers. Hence, we recommend that you use beta releases on installations with small (as opposed to full production-level) user bases.

Upgrade Procedure

If you wish to upgrade your existing Confluence installation with this version, ensure you have created a separate copy of your current Confluence production installation *first* and using that copy, follow the normal upgrade instructions to upgrade it to this beta release. If you have also implemented customised site- or space-specific layouts, you will need to re-implement them after the upgrade. Otherwise, some of the new features in Confluence (or possibly existing features) may not function correctly.

Downloads

All development releases are available from Development Releases on the Atlassian website.

Known Issues

Below are some known issues associated with Confluence 3.1.

On this page:

- JIRA/Crowd and Confluence deployment
- JIRA Gadgets in Confluence
- Bamboo integration

- PDF exports only render gadgets as links
- Problematic Confluence Gadgets window when running Confluence on Java 6
- Other issues

JIRA/Crowd and Confluence deployment

Confluence will not start up or will display strange behaviour (drop down menus not working) if JIRA 4.0/4.0.1 or Crowd 2.0.x is running on the same application server installation, for example, by attempting to run Confluence and JIRA 4.0 in the same Apache Tomcat server installation. This problem results from a bug in JIRA (tracked as JRA-19894) that is scheduled to be fixed in JIRA 4.0.2. Crowd will be fixed in the 2.1 release. In the meantime, please read our KB article on how to resolve this issue.

In the meantime, you can run JIRA or Crowd and Confluence in different 'instances' of the same application server, for example two separate Apache Tomcat server installations. If you do so already or have installed standalone distributions of both Confluence and JIRA 4.0 or Crowd 2.0.x, you can ignore this known issue.

JIRA Gadgets in Confluence

As of Confluence 3.1, users can embed gadgets into Confluence pages. When integrating JIRA gadgets into Confluence pages, you may encounter UI problems like a missing "Login"-button which is required to make the gadget authenticate with the JIRA server. This problem can be circumvented by setting up JIRA and Confluence to use Trusted Apps communication (since it removes the need for manual authentication). See KB article for details.

Bamboo integration

Our continuous integration product Bamboo exposes gadgets which can be embedded into Confluence pages. However, some of these exhibit problems once embedded onto a Confluence page:

- BAM-4900 : Unable to edit Bamboo gadgets in Confluence
- BAM-4890 : Bamboo gadget added in JIRA dashboard is not saving the preferences

These bugs are being fixed in Bamboo 2.5, which will ship in January 2010.

PDF exports only render gadgets as links

If you place any gadget on a Confluence page and export the page to PDF, the gadget output will not be rendered in the PDF output. Instead, each gadget is rendered on a page as a box containing the name of the gadget, the latter of which is hyperlinked. Clicking this hyperlink, opens the gadget contents itself in a new browser window or tab.

Problematic Confluence Gadgets window when running Confluence on Java 6

The Confluence Gadgets window may indicate that '**An error has occurred while trying to load the Gadget Directory**' and prevent you from accessing the URLs of your Confluence gadgets. This problem can occur if you are running Confluence on Java 6. After you install Confluence 3.1 or upgrade an existing Confluence installation to this version, please check the Confluence Gadgets window immediately after starting the Confluence server.

If you see this error message and cannot access your Confluence gadgets, it can be resolved by restarting Confluence. (You may need to do this more than once.)

For more information about this issue, please refer to CONF-17417.

Other issues

Refer to our JIRA site for a list of Confluence 3.1-specific bugs.

Highlights of Confluence 3.1

Highlights of this Release:

- Introducing Gadgets
- Drag & Drop
- Office 2007 Support
- New 'Move Page' Feature
- Enhanced Image Browser
- Draft Comparisons
- Page Restrictions Dialog Box
- Web Browser Version Support
- Other Editor Enhancements
    - Edit Mode Exit Notification
    - New Rich Text Editor Insert Menu
    - Macro Browser Smart Fields
    - Editor speed
- Other Improvements
    - Add Pages or Blog Posts from the Dashboard
    - New 'Link to this page' feature
    - 'More' option on Activity Streams
    - User Interface Performance Improvements
    - Other Small Enhancements and Improvements to Confluence



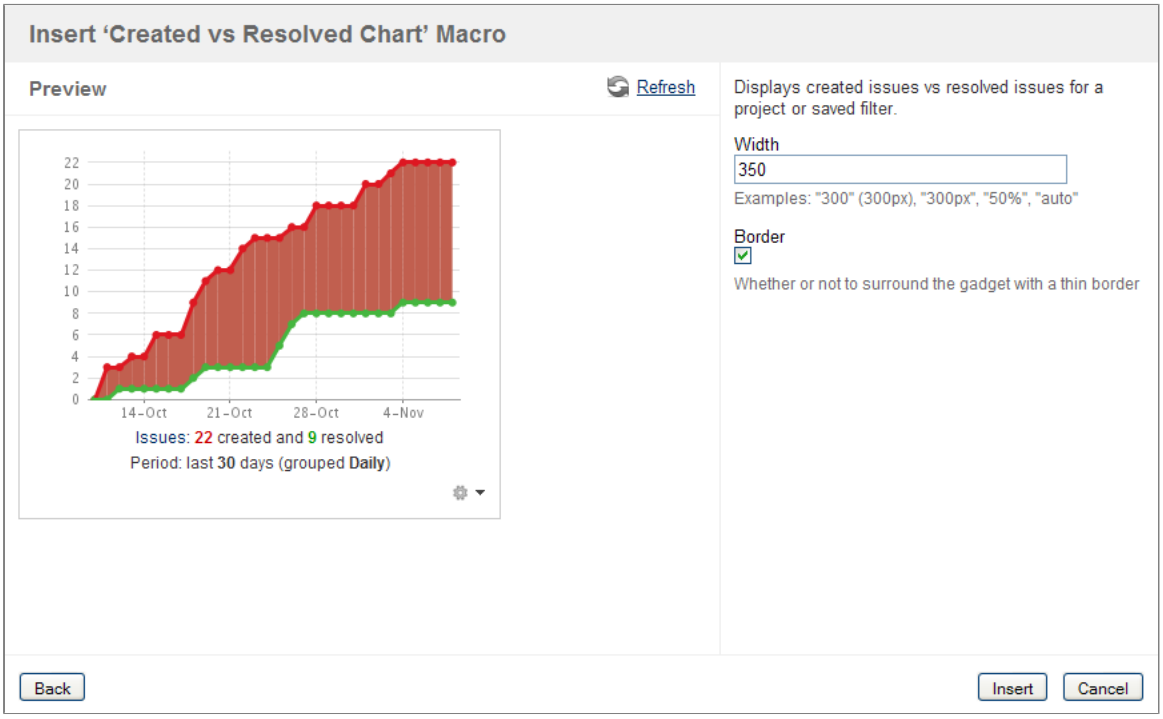
Introducing Gadgets

Gadgets are small objects that offer dynamic content and functionality which may be served by any OpenSocial-compliant web application, such as JIRA 4.0+, Confluence or non-Atlassian applications such as iGoogle and Gmail.

- Confluence supports the use of gadgets in pages and blog posts, which are accessible through the macro browser.
- Confluence can also serve its own gadgets, for use in any other OpenSocial-compliant web application. Two such gadgets are bundled with Confluence:
  - **Activity Stream** — This gadget shows a list of recent activities that have occurred on the Confluence server, such as the addition of new pages, blog posts or comments, content edits, status updates and so on.
  - **Quick Navigation Aid** — This gadget provides heading and content search capabilities on a Confluence server.

    Your Confluence installation can also serve these gadgets in any of its own pages or blog posts.

Inserting a JIRA Gadget onto a Confluence Page





Drag & Drop

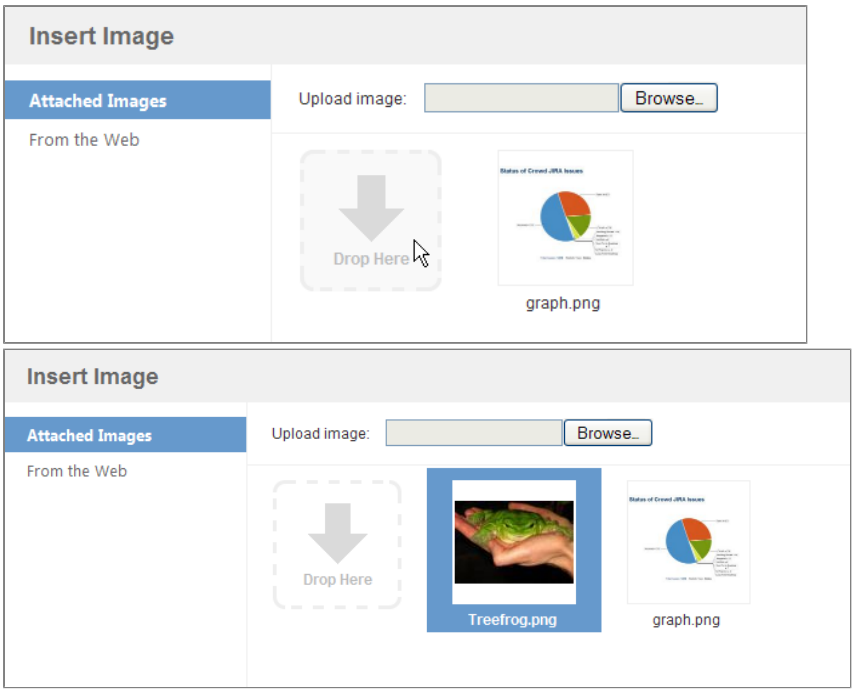
The new 'drag and drop' feature allows you to drag one or more file(s) which are accessible from your computer and drop them directly into a Confluence page or blog post.

- Files can be attached to a page or blog post by dropping them directly onto the page view or the 'Attachments' list associated with the page.
- Image files can be attached to a page or blog post by dragging them from your computer directly onto the Image Browser.
- Image and Office files can be added directly into your Confluence page or blog post content by dropping them into the rich text editor's editor window.

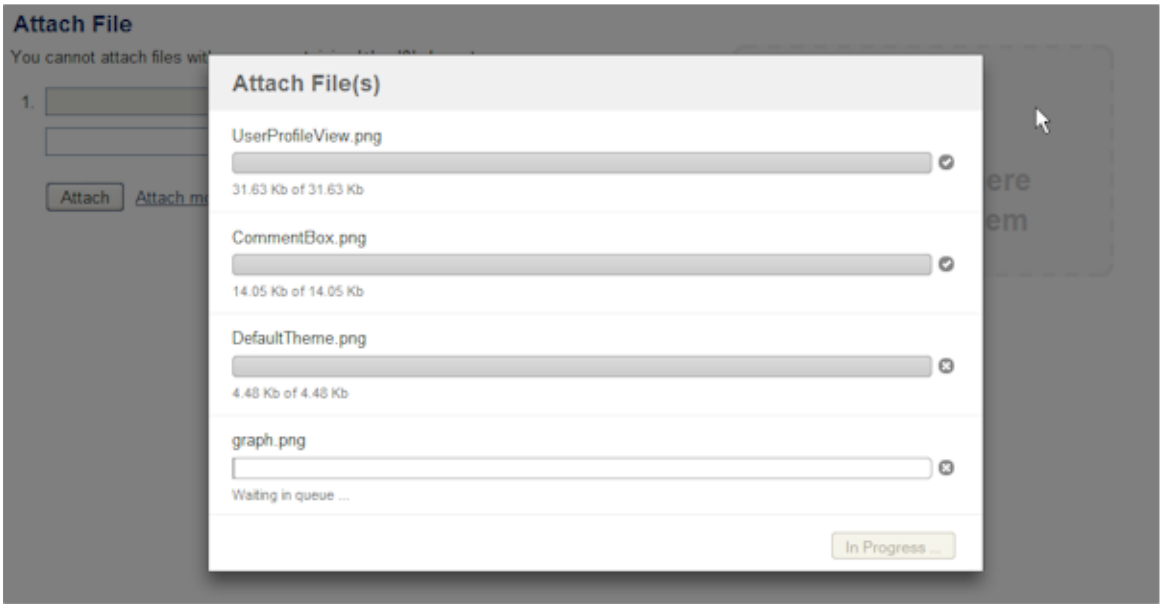
Download Video

For more information about this feature and on how to set it up, refer to the Drag-and-Drop documentation.

Screenshots: Attaching an Image to the Image Browser



Screenshot: Attaching Multiple Files to an 'Attachments' list



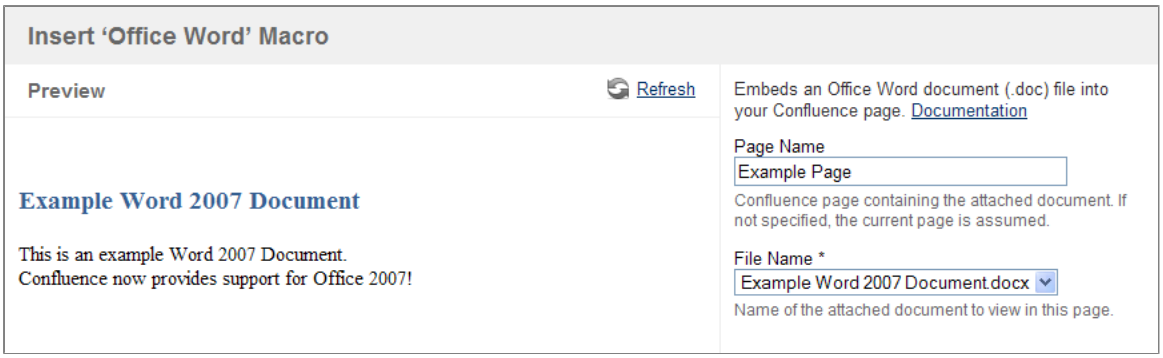


Office 2007 Support

Confluence now provides full support for the new Office 2007 file formats, allowing you to view and edit content from Microsoft Word 2007 (`.docx` and `.dotx`), PowerPoint 2007 (`.pptx` and `.potx`) and Excel 2007 (`.xlsx`) files.

- Along with existing Microsoft Office versions, Confluence now fully indexes Microsoft Office 2007 files and their content can be searched by Confluence.
- Using Confluence's Office connector macros, you can insert Word, PowerPoint or Excel 2007 files directly into your Confluence page or blog post.
- Office files can be edited directly from any page or blog post or their list of attachments.

  If you use the Firefox browser to work with Confluence, don't forget to reconfigure the Firefox add-on (WebDAV Launcher options) to handle the new Office 2007 file extensions. Otherwise, you will not be able to edit these new Office 2007 file formats from Confluence.

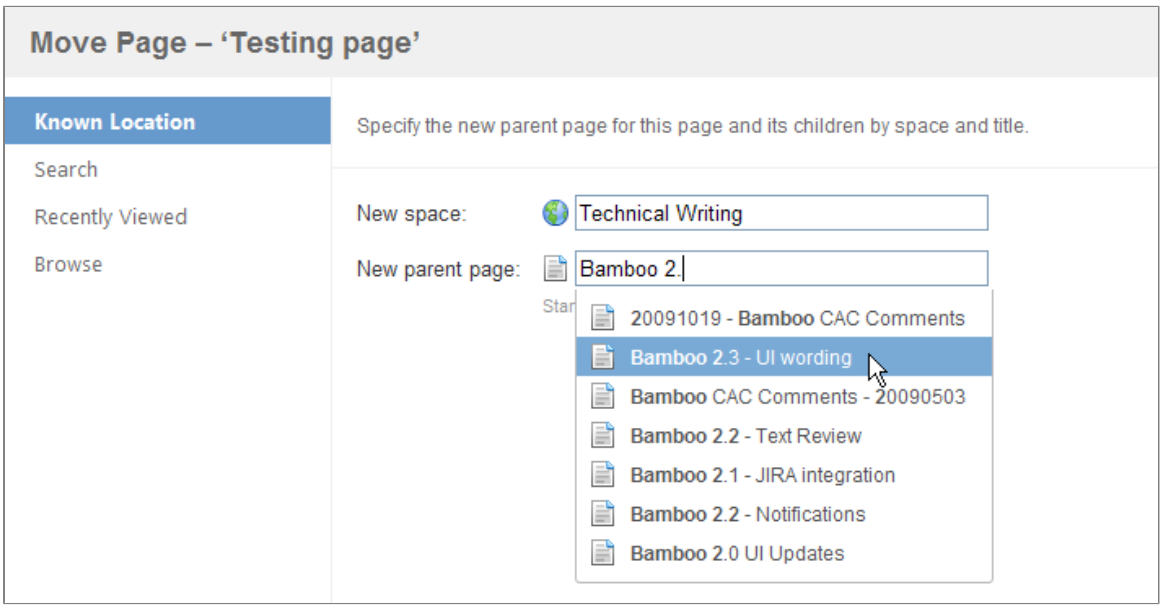




New 'Move Page' Feature

Confluence introduces a new page moving feature, that easily allows you to move the page you are currently viewing, adding or editing to another page elsewhere in the same or another space of your Confluence site. This feature is available through a new 'Move Page' dialog box, which provides the following flexible methods for moving pages:

- **Known Location** – Allows you to type the name of a space and within that space, the 'parent' page under which to move your page.
- **Search** – Allows you to search for a 'parent' page (within a selected space or set of spaces) under which to move your page.
- **Recently Viewed** – Allows you to select one of your recently viewed pages to be the 'parent' of your page to be moved.
- **Browse** – Allows you to select a space and page (within the tree of pages in the space) that will be the 'parent' of your page to be moved.
  For more information, refer to Moving a Page.

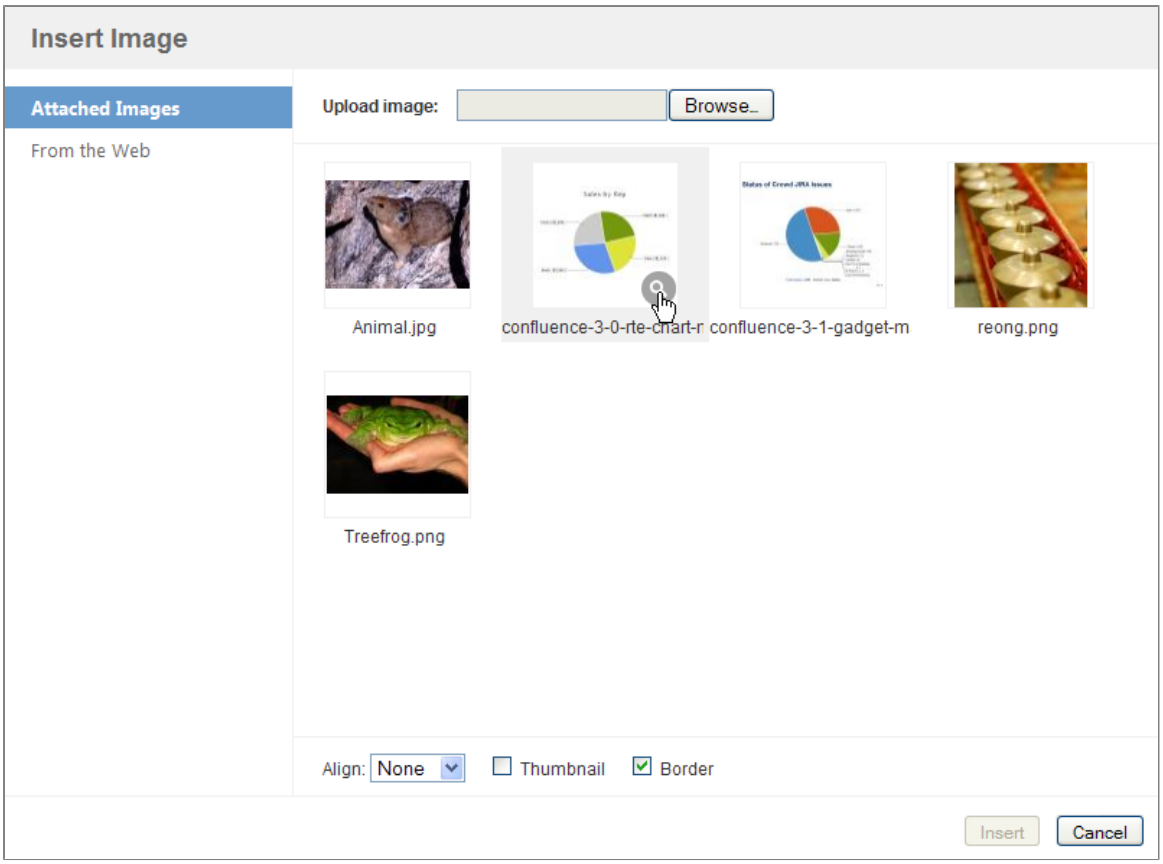




Enhanced Image Browser

A new 'Image Browser' has been introduced to replace the old 'Insert Image' window. The image browser provides a less-cluttered and enhanced interface that allows you to:

- Preview an image in detail before inserting it into a page. This is done by hovering over any image in the browser and clicking the 'magnifying glass' icon in the lower-right corner.
- Preview an image elsewhere on the web via its URL before inserting it into a page.



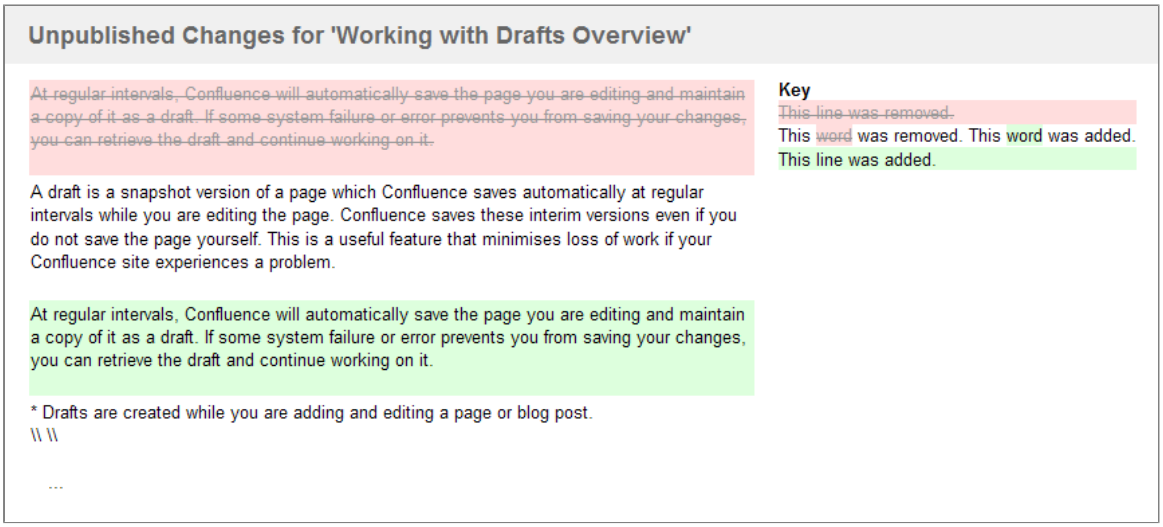


Draft Comparisons

Confluence's drafts features have been enhanced, such that you can now view any of your unsaved draft changes before deciding to resume editing them. This nifty feature comes in handy, particularly when other people have made subsequent changes to a page or blog post in your drafts list and you need to merge changes or resolve a conflict.



Title	Last Saved Date	Operations
Page with Partial Content	less than a minute ago	Resume Editing | Discard
Page with Content to Merge	2 minutes ago	View Changes | Merge and Resume Editing | Discard
Page with a Conflict	12 minutes ago	View Changes | View Conflict | Discard

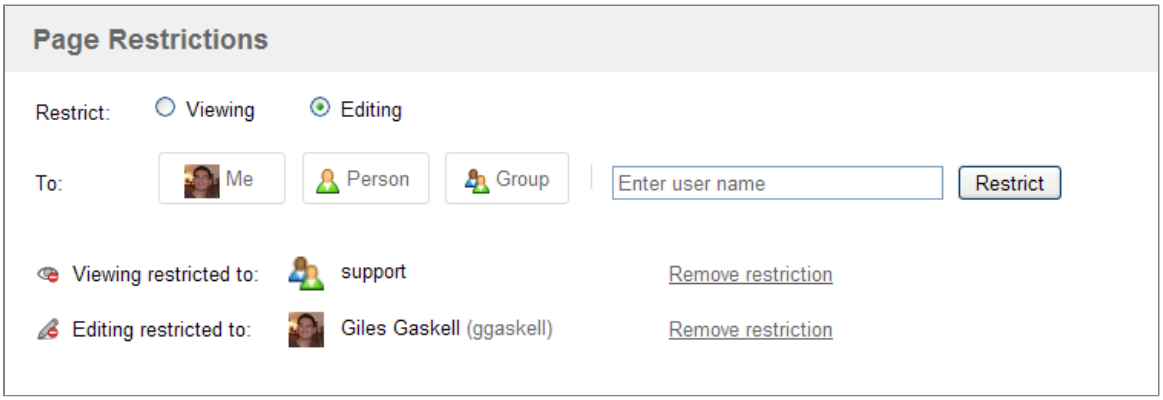




Page Restrictions Dialog Box

Confluence's page restrictions feature has been incorporated into a convenient and accessible dialog box that is now easier to use than before.

- The page restrictions dialog box can be accessed from the padlock icon or the '**Tools**' -> '**Restrictions**' menu item whilst viewing any Confluence page. From this dialog box, you can see all viewing and editing restrictions associated with the current page. You no longer need to view the page's associated 'Info' page to see the page's restrictions.
- You no longer have to edit a page to modify its page restrictions. You can edit all page restrictions from this easily accessible dialog box.
  The page restrictions dialog box is still accessible when when a page is in edit mode.
- In addition to user and group names, the name field also accepts a user's full name. Full names are 'auto-completed' to help you find the relevant person more rapidly.





Web Browser Version Support

Confluence 3.1 now fully supports the following recent web browser versions:

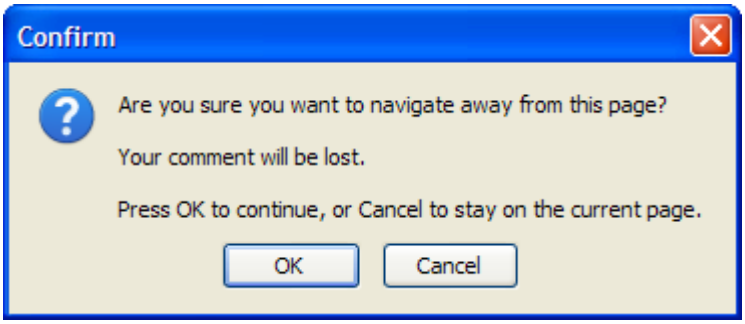
- Internet Explorer 8
- Safari 4
- Firefox 3.5



Other Editor Enhancements

Edit Mode Exit Notification

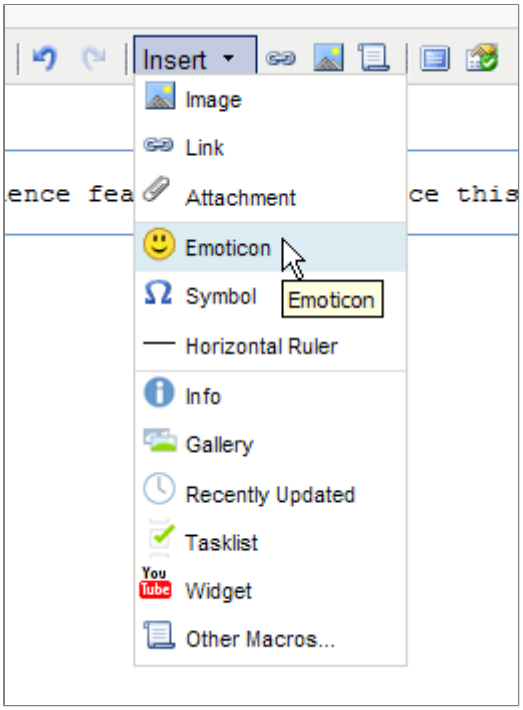
Whenever you add or edit a page, comment or blog post and then click onto another Confluence feature that navigates away from your unsaved content, a message box appears, warning that your content will be saved as a draft (if it is a page) or lost (if a comment). This allows you to cancel out of this action if it was accidental.



New Rich Text Editor Insert Menu

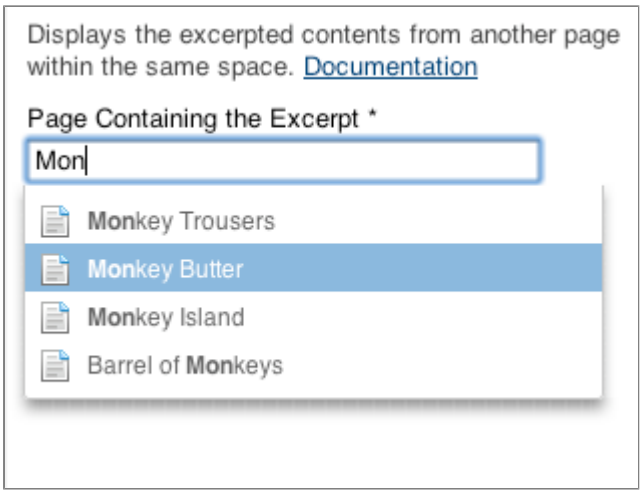
Confluence's rich text editor now combines a number of its commonly used editing features into a new convenient '**Insert**' menu.

- The Horizontal line, Insert Symbol and Insert Emoticon Toolbar icons have been moved into the new insert menu.
- The functionality to insert images, links or attachments into a page can now also be accessed from this menu.
- The macro browser, as well as a number of commonly-used macros are conveniently accessible from this menu too.



Macro Browser Smart Fields

When using the Macro Browser, an 'auto-complete' feature is now provided on any parameters that require the entry of a single item, such as a page title, username or space key. This greatly facilitates the customisation of macros and minimises the need to know the exact item names in advance.



Editor speed

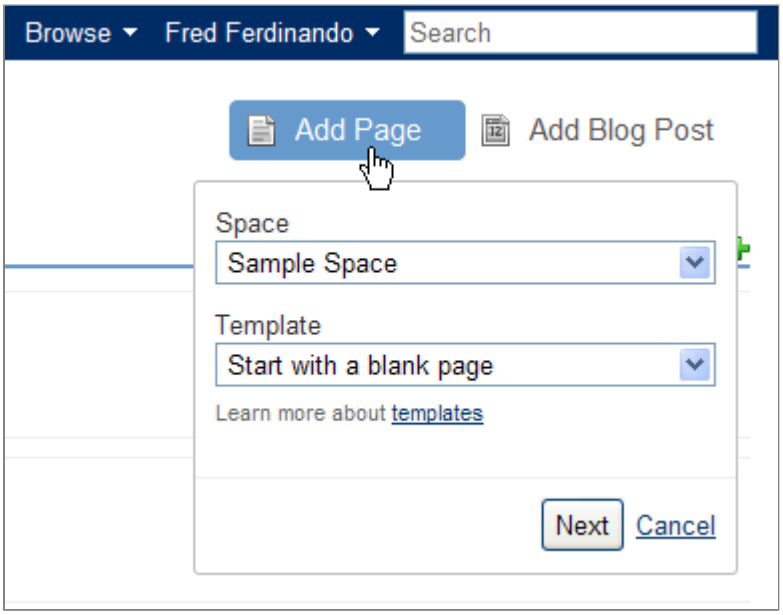
Thanks to many individual technical improvements, the Rich Text Editor opens up a lot faster than in previous Confluence releases. On a local network it appears almost instantly, and even when accessing a Confluence server on a different continent (in our case, accessing a Confluence server in the US from Australia) it takes less than 3 seconds to start editing a page.



Other Improvements

Add Pages or Blog Posts from the Dashboard

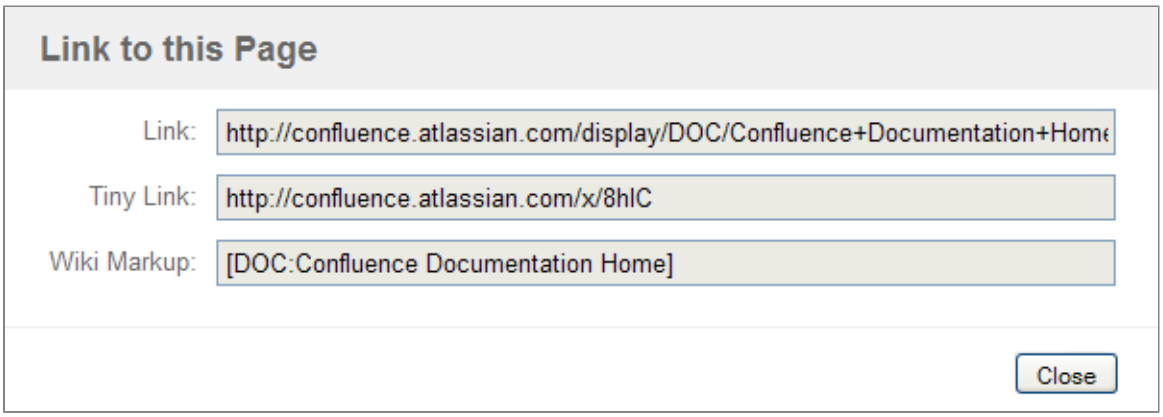
You can now add pages or blog posts directly from the Dashboard without having to browse to a specific space first. To do this, click on either the '**Add Page**' or '**Add Blog Post**' buttons to open the pop-up balloon, which allows you to choose the space in which to add the new page or blog post and in the case of pages, a template on which to base the page content.



New 'Link to this page' feature

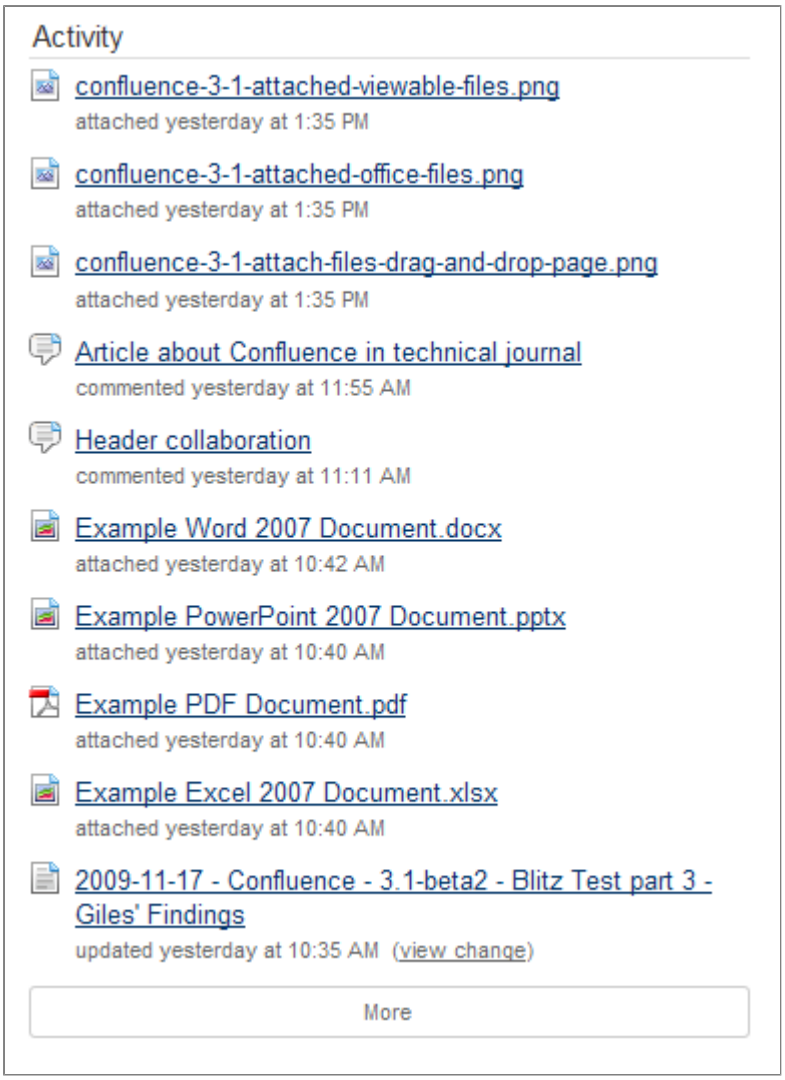
If you wish to link to a Confluence page from any other location on the web, use the convenient '**Link to this Page**' feature (available from any page's or blog post's '**Tools**' menu). Upon selecting this feature, the 'Link to this Page' dialog box opens, from which you can copy three versions of the link to embed elsewhere:

- **Link** – Standard URL which should work from any other accessible location on the web.
- **Tiny Link** – A reduced-length version of the 'Link', which can be used in text fields of limited length, such as tweets or Confluence Status Updates.
- **Wiki Markup** – A wiki markup version of the link, which can be used in any other location within your Confluence site.



'More' option on Activity Streams

A 'More' option has been added to various activity streams throughout the Confluence interface, including the profile sidebar, a user's profile page and via the recently updated macro. Clicking 'More' expands the list of results, providing a convenient means of accessing progressively more distant user activities.



User Interface Performance Improvements

Most JavaScript and Cascading Style Sheet (CSS) files are now downloaded in one batch, greatly improving the performance of Confluence's editing features and general page rendering.

Other Small Enhancements and Improvements to Confluence

- Support for OAuth — With the introduction of gadgets (above) in this release, Confluence 3.1 now allows you to establish OAuth relationships with other web applications such as JIRA 4.0+, iGoogle, Gmail etc., thereby allowing them to share resources via gadgets.
- New Log In and Log Out screens.
- In an aim to minimise confusion, 'News Items' are now consistently called 'Blog Posts' throughout the Confluence interface and a list of blog posts is collectively referred to as a 'Blog'.
- Macro developers are now able to specify whether the macro body should or should not be displayed in Rich Text editor. For more information, please refer to CONF-12149.
- Other minor interface improvements.

Confluence 3.1 Release Candidate 1 Release Notes

Milestone release advisory

Do not use this release to upgrade your production systems.

For all production use and testing of Confluence, please use the latest official release.
This release is a public **development release** ('milestone') leading up to Confluence 3.2, which will probably ship in mid March 2010. Development releases are a snapshot of our work in progress, allowing our customers and especially plugin developers to see what we're up to.

Who should upgrade?

Please note the following

- **Development releases are not safe** — Development releases are snapshots of the ongoing Confluence development process. As such:
  - While we try to keep these releases stable, they have not undergone the same degree of testing as a full release.
  - Features in development releases may be incomplete, or may change or be removed before the next full release.
- **No upgrade path** — Because development releases represent work in progress, we **can not** provide a supported upgrade path between development releases, or from any development release to the eventual final release. Thus, it is possible that you will not be able to migrate any data you store in a Confluence development release to a future Confluence release.

In supplying milestone releases, our aim is to provide plugin developers with an opportunity to see the latest changes in the code.

Each milestone release has passed all our automatic tests and has been used for one week on our official internal Confluence server. Most of the issues solved have been reviewed too, and all of our milestone releases even have been performance tested for a while.

However, since our milestones releases are timeboxed (i.e. they get released every two weeks, no matter how far we have come implementing features and bugfixes), there is always a chance that we have new known bugs that are scheduled to be fixed in the next milestone, or completely new bugs unknown even to us.

Additionally, we have not completed our performance testing and compatibility testing for databases and application servers. So, for example, a milestone release may behave well on a small installation but show severe problems when subjected to many users.

Upgrade Procedure

Follow the normal upgrade instructions to upgrade from Confluence 3.1.x to this release. *We strongly recommend that you backup your Confluence home directory and database before upgrading!*

Downloads

All development releases are available from the development releases page on the Atlassian website.

While you were away...

You can see all the feature details in the beta 2 announcement. This update about the Release Candidate is just for developers and evaluators who want to know what changed between the Beta 2 and the Release Candidate. Not much!

Changes in RC1

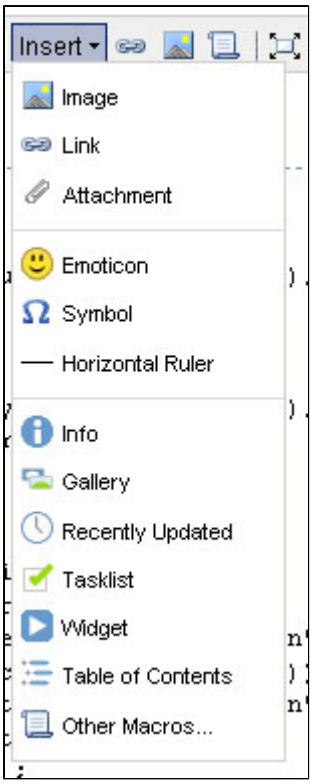
It's should come as no surprise that we have mainly been bugfixing since the Beta 2.

Faster Editor

But, in order to make our Selenium builds pass more reliably, Agnes had to spend some time on making the editor faster. And she did!

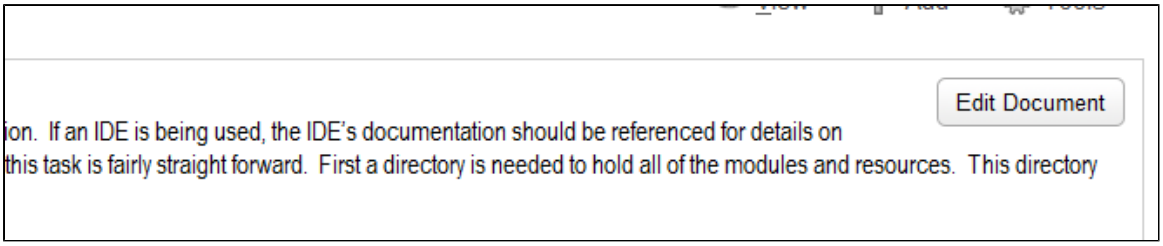
UI tweaks

Matt came up with a better widget connector icon, we've added the table-of-content which had been forgotten, and Stephen came up with a great fullscreen icon that really grabs your attention and means something.



Office file preview has a more prominent edit button

A common complaint at Atlassian Summit was that people didn't realize that you could launch your editor directly from the file preview so we've made it more obvious.



Road ahead

We aim at turning the RC into the Final Release and ship it on Thursday the 3rd December.

Release Notes 3.2-m3 ("Milestone 3")

Milestone release advisory

Do not use this release to upgrade your production systems.

For all production use and testing of Confluence, please use the latest official release.
This release is a public **development release** ('milestone') leading up to Confluence 3.2, which will probably ship in mid March 2010. Development releases are a snapshot of our work in progress, allowing our customers and especially plugin developers to see what we're up to.

Who should upgrade?

Please note the following

- **Development releases are not safe** — Development releases are snapshots of the ongoing Confluence development process. As such:
  - While we try to keep these releases stable, they have not undergone the same degree of testing as a full release.
  - Features in development releases may be incomplete, or may change or be removed before the next full release.
- **No upgrade path** — Because development releases represent work in progress, we **can not** provide a supported upgrade path between development releases, or from any development release to the eventual final release. Thus, it is possible that you will not be able to migrate any data you store in a Confluence development release to a future Confluence release.

In supplying milestone releases, our aim is to provide plugin developers with an opportunity to see the latest changes in the code.

Each milestone release has passed all our automatic tests and has been used for one week on our official internal Confluence server. Most of the issues solved have been reviewed too, and all of our milestone releases even have been performance tested for a while.

However, since our milestones releases are timeboxed (i.e. they get released every two weeks, no matter how far we have come implementing features and bugfixes), there is always a chance that we have new known bugs that are scheduled to be fixed in the next milestone, or completely new bugs unknown even to us.

Additionally, we have not completed our performance testing and compatibility testing for databases and application servers. So, for example, a milestone release may behave well on a small installation but show severe problems when subjected to many users.

Upgrade Procedure

Follow the normal upgrade instructions to upgrade from Confluence 3.1.x to this release. *We strongly recommend that you backup your Confluence home directory and database before upgrading!*

Downloads

All development releases are available from the development releases page on the Atlassian website.

Welcome back

it's been a while since we shipped Confluence 3.1 in December. Since then, we've been having heaps of fun at our second Lab Week, during which we experimented with all sorts of new technology and crazy ideas, mainly around the editor. We've been busy coding up new features over the last weeks, but also various back-end changes. Milestones 1 and 2 were internal only and didn't show much visual anyway, but now we're getting there.

The Confluence 3.2 development cycle is going to be a short one. We are currently aiming at shipping in Mid March. There will probably be one a M4 milestone in two weeks, but then we're in Beta mode already.

The features

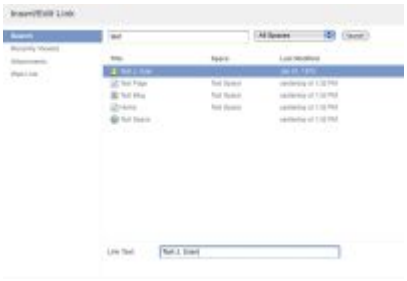
New Links Browser

The new link browser is now available in the editor toolbar (the old one is still there as a backup).



Note: this funky icon is temporary

This new browser now has all the functionality present in the current link popup, except for the recently modified tab. You can now add links to attachments and recently viewed pages, which weren't available in m2. Please raise any bugs found against the linking component in Confluence.



Know issues

- quick search only returns pages
- error messages from the server aren't displayed yet
- usability work still to come (like cursor placement and keyboard navigation)

Documentation theme

We are going to bundle the new Documentation Theme,

Fixed Width Theme

We are introducing a theme that by default has sidebars. The main work here is behind the scenes, fixing up our HTML and CSS so that 3rd-party theme developers will benefit and be able to write such themes more easily.



If you're wondering where the 'sudden green' has come from, it has originated from Confluence's product colour on our website.

Naturally the colour scheme is customisable:

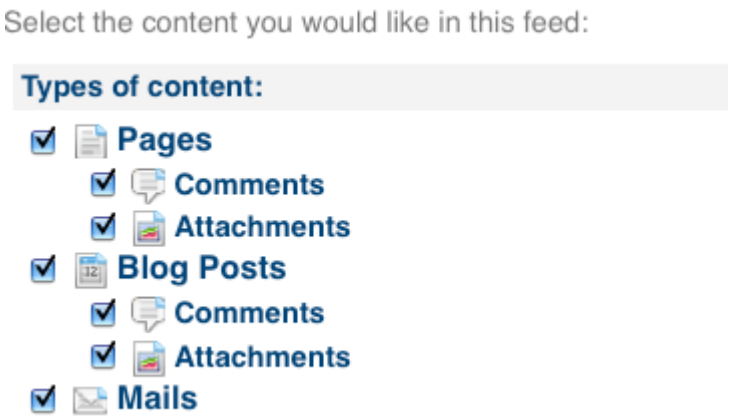


Known issues

- Space Admin screens still need work (site admin screens don't get themes applied, but space admin does)
- There's more polish to go into this theme - mostly small spacing and font size adjustments
- Our minimum resolution is 1024x768; content that doesn't fit at that size won't fit in the fixed width theme either.
- There is a certain discussion around the colours, and wether we should be a bit more conservative in terms of the background. We will keep discussing.
- The main point here is that many bugs have been fixed, and that 3rd-party theme developers will benefit from many under-the-hood bugfixes to our html and css.

Anatoli's 20% project: RSS improvements

- CONF-18372 : changed the feed builder UI so that now you can chose comments/attachments for either pages, blogs, both or none. The new UI will generate slightly different parameters so that they don't interfere with old RSS query format.



There are plans to further improve the UI.

- CONF-9312: changed the behavior of rss feed when filtering by label. Now when you ask for comments and pages with a particular label you will get the labeled pages and comments on those pages.

Backend changes

The Engine Room has upgraded plenty of libaries:

- Many common modules were upgraded.
- Migrated our Sal Plugin into Confluence
- REST Service for retrieving Recently Viewed pages
- Start of REST service for Confluence searches

More documentation about REST will soon follow.

Next steps

There are still heaps of improvements that will make it into 3.2, so stay tuned for M4 and Beta1, both due this month.

Cheers,
The Confluence Development Team

Release Notes 3.2-m4 ("Milestone 4")

Milestone release advisory

Do not use this release to upgrade your production systems.

For all production use and testing of Confluence, please use the latest official release.
This release is a public **development release** ('milestone') leading up to Confluence 3.2, which will probably ship in mid March 2010. Development releases are a snapshot of our work in progress, allowing our customers and especially plugin developers to see what we're up to.

Who should upgrade?

Please note the following

- **Development releases are not safe** — Development releases are snapshots of the ongoing Confluence development process. As such:
  - While we try to keep these releases stable, they have not undergone the same degree of testing as a full release.
  - Features in development releases may be incomplete, or may change or be removed before the next full release.
- **No upgrade path** — Because development releases represent work in progress, we **can not** provide a supported upgrade path between development releases, or from any development release to the eventual final release. Thus, it is possible that you will not be able to migrate any data you store in a Confluence development release to a future Confluence release.

In supplying milestone releases, our aim is to provide plugin developers with an opportunity to see the latest changes in the code.

Each milestone release has passed all our automatic tests and has been used for one week on our official internal Confluence server. Most of the issues solved have been reviewed too, and all of our milestone releases even have been performance tested for a while.

However, since our milestones releases are timeboxed (i.e. they get released every two weeks, no matter how far we have come implementing features and bugfixes), there is always a chance that we have new known bugs that are scheduled to be fixed in the next milestone, or completely new bugs unknown even to us.

Additionally, we have not completed our performance testing and compatibility testing for databases and application servers. So, for example, a milestone release may behave well on a small installation but show severe problems when subjected to many users.

Upgrade Procedure

Follow the normal upgrade instructions to upgrade from Confluence 3.1.x to this release. *We strongly recommend that you backup your Confluence home directory and database before upgrading!*

Downloads

All development releases are available from the development releases page on the Atlassian website.

With great pride we bring to you the best milestone of this release cycle. We are now more or less feature-complete. There are a few usability issues which we still want to fix, and yes, there are still plenty of bugs. We didn't manage to bundle a few of our new great plugins yet, and have large chunks of code to review. That's why we didn't call this Beta yet, but named it M4

But overall, this is how 3.2 will get rolled out to customers. We want your feedback: Speak now or forever hold your peace.

Note to partners and translators: The preliminary translation diff files can be found as a subpage of this overview page: Translations for Confluence Releases



Engine Room

New REST Resources

- Attachment Resource

Request: https://confluence.example.com/rest/prototype/1/attachment/161

```
01.<?xml version="1.0" encoding="UTF-8" standalone="yes"?>
02.<attachment niceType="Image" version="1"
03.comment="image of Sydney harbour"
04.fileName="harbour.jpg" type="attachment" id="161">
05.        <ownerId>69518087748064278</ownerId>
06.        <title>harbour.jpg</title>
07.        <lastModifiedDate date="2004-11-21T23:20:18-0600" friendly="Nov 21, 2004"/>
08.        <createdDate date="2004-11-21T23:20:18-0600" friendly="Nov 21, 2004"/>
09.</attachment>
```

Some fields removed

More in-depth documentation coming soon.

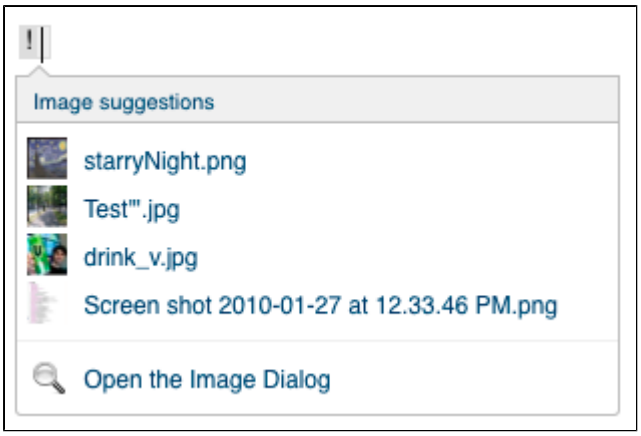
Other News

- Integrate auto-complete and Link Browser to use REST search and attachment resources
- Bug Fixes



Autocomplete

We are adding a new feature called Autocomplete to the Rich Text Editor, which currently only works in Firefox and (try your luck!) Safari. You can type "**[**" at the start of a line to display your browsing history and link to those pages, and if you keep typing you'll get recommendation for pages that match what you types. If you type "**!**" you'll see the attached images to the current page, and if you keep typing you'll be able to link to any image in your deployment.



This feature's goal is to make the RTE even faster than using wiki markup. Try it out and tell us what you think!

Note: May require a re-index of your system to show attachments

Known issues

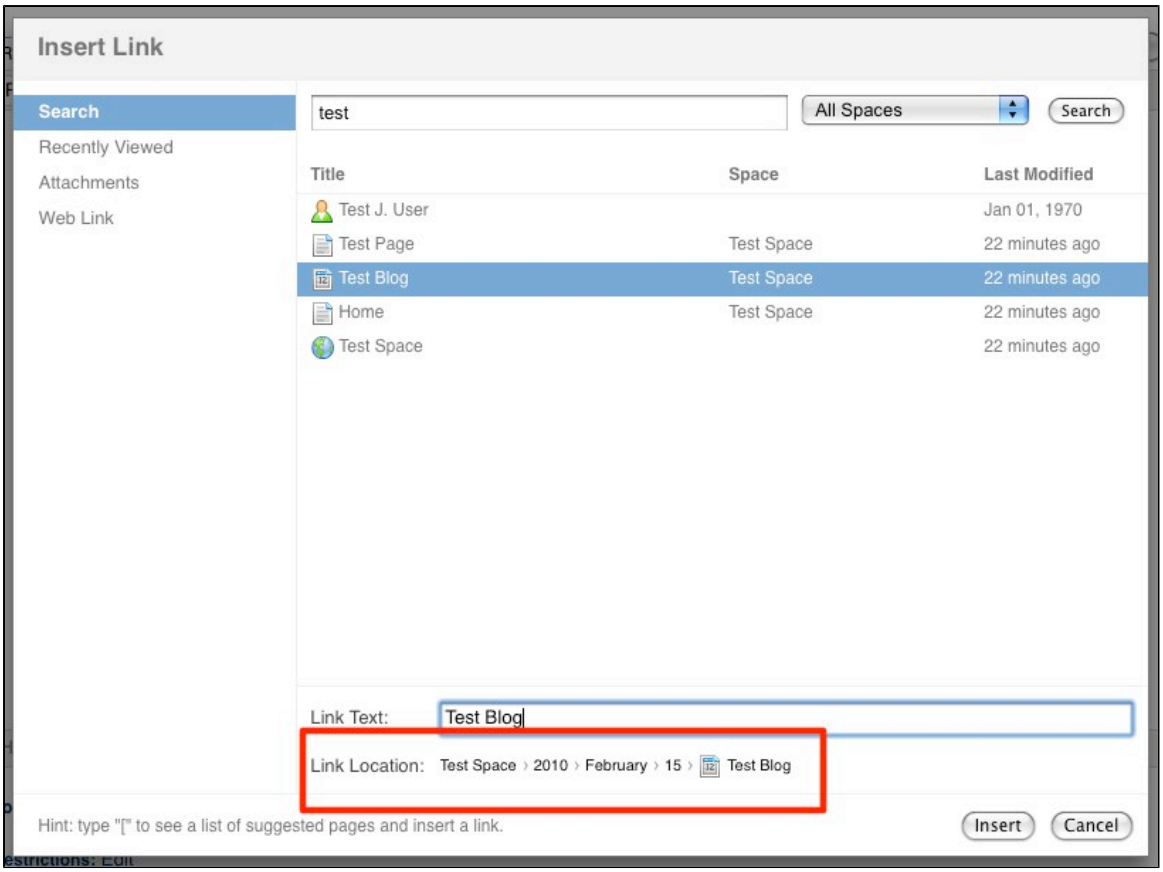
- CONF-18281 - doesn't work in IE yet
- CONF-18584 - when both the RTE and outer frames are scrolled the dropdown location is incorrect
- Plenty of small niggly bugs.



New Links Browser

The Links Browser is now feature complete and has been enabled as the default insert link button in the editor.

The breadcrumbs in the search panel of the dialog is new in this milestone. This is particularly helpful when editing existing links.



Know issues

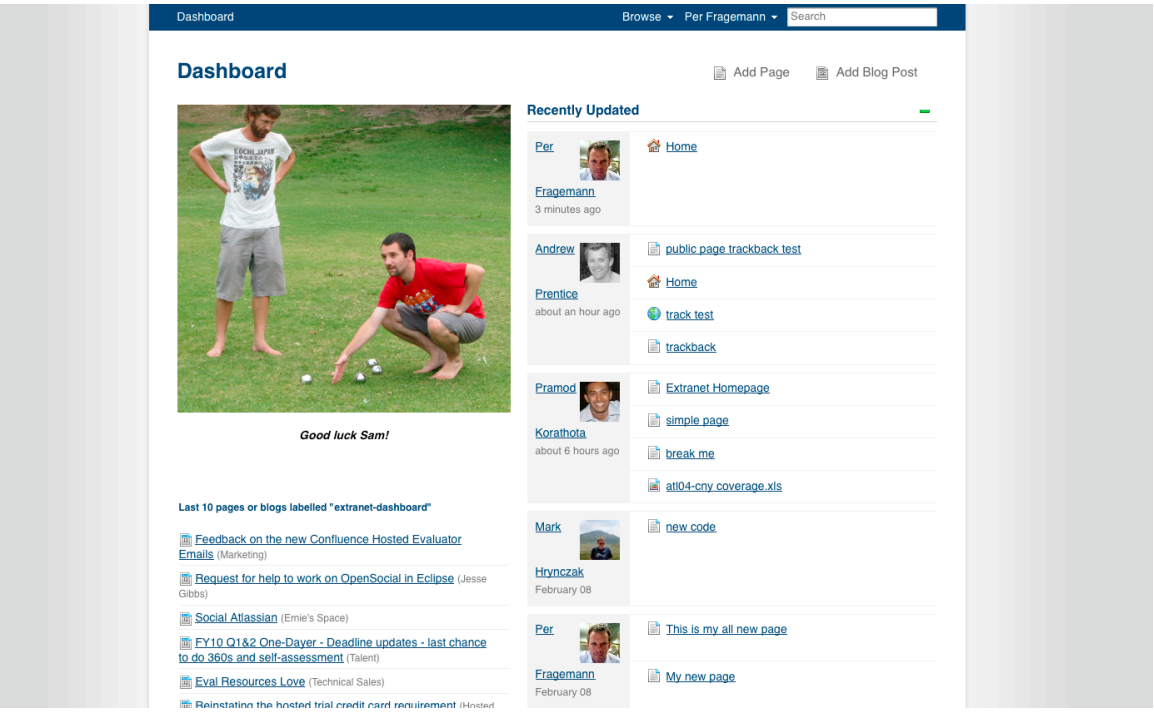
- CONF-18589 IE 6 caches responses for recently viewed and attachments tabs
- CONF-18575 Pressing enter on the search field leaves the page in IE7
- CONF-18455 After viewing a user profile, the recently viewed tab displays no pages



New Theme

Our new Fixed Width Theme (proper name still tba) now has

- new fancy gradient
- soothing blue colour scheme
- personal sidebar folds away neatly
    - we tried floating it over content, but people weren't so keen on that.
- tidier grid than last milestone
- comment threading now far more robust, comments RTE toolbar collapses or scrolls instead of truncating (this solution will be migrated to core code after trial use in fixed width - the problem can also happen in default)
- key bugfixes: space admin screens and RTE spacing



Known issues

- The header nav dropdowns don't line up with the bottom of the header yet
- Personal info sidebar has some bugs to resolve after final design change (mainly: IE6 breaks, sidebar overlaps "navigate space" in some browsers)
- Attachments page is getting truncated in some circumstances
- Some RTE messages wrap poorly

Note for bug tracking: bugs that also occur in the default Confluence at 1024x768 are not theme bugs per se, even if they're being investigated at the same time.



Bugfixes

Errors were reported by the JIRA trusted connection.

- APP_UNKNOWN; Unknown Application: {0}; ["confluence:4557196"]

JIRA Issues (11 issues)

Type	Key	Summary	Assignee	Reporter	Priority	Status	Resolution	Created	Updated	Due
	CONF-18489	The Download All as zip is using the wrong mime type	Andrew Lynch [Atlassian]	Brad Baker [Atlassian]		Resolved	Fixed	Feb 04, 2010	Feb 16, 2010	
	CONF-18397	Error message for invalid space labels is incorrectly encoded	Ryan Thomas [Atlassian]	Mark Hrynczak [Atlassian]		Resolved	Fixed	Jan 28, 2010	Mar 10, 2010	
	CONF-18218	LDAP users with uppercase usernames cannot add favourites	Xu-Heng Tjhin [Atlassian]	Gurleen Anand [Atlassian]		Resolved	Fixed	Jan 11, 2010	Mar 24, 2010	
	CONF-18195	Impossible to select a user using mouse from an autocomplete dropdown box in a search filter "Who"	Xu-Heng Tjhin [Atlassian]	Vincent Chang		Resolved	Fixed	Jan 07, 2010	Mar 27, 2010	
	CONF-16629	Uploading large fonts for PDF export fails with XSRF error	Ryan Thomas [Atlassian]	Igor Minar		Resolved	Won't Fix	Aug 12, 2009	Feb 17, 2010	
	CONF-16578	When Confluence license is invalid in confluence.cfg.xml, Confluence fails with "Upgrade failed" in the browser	Ryan Thomas [Atlassian]	Roy Hartono [Atlassian]		Resolved	Fixed	Aug 06, 2009	Mar 27, 2010	

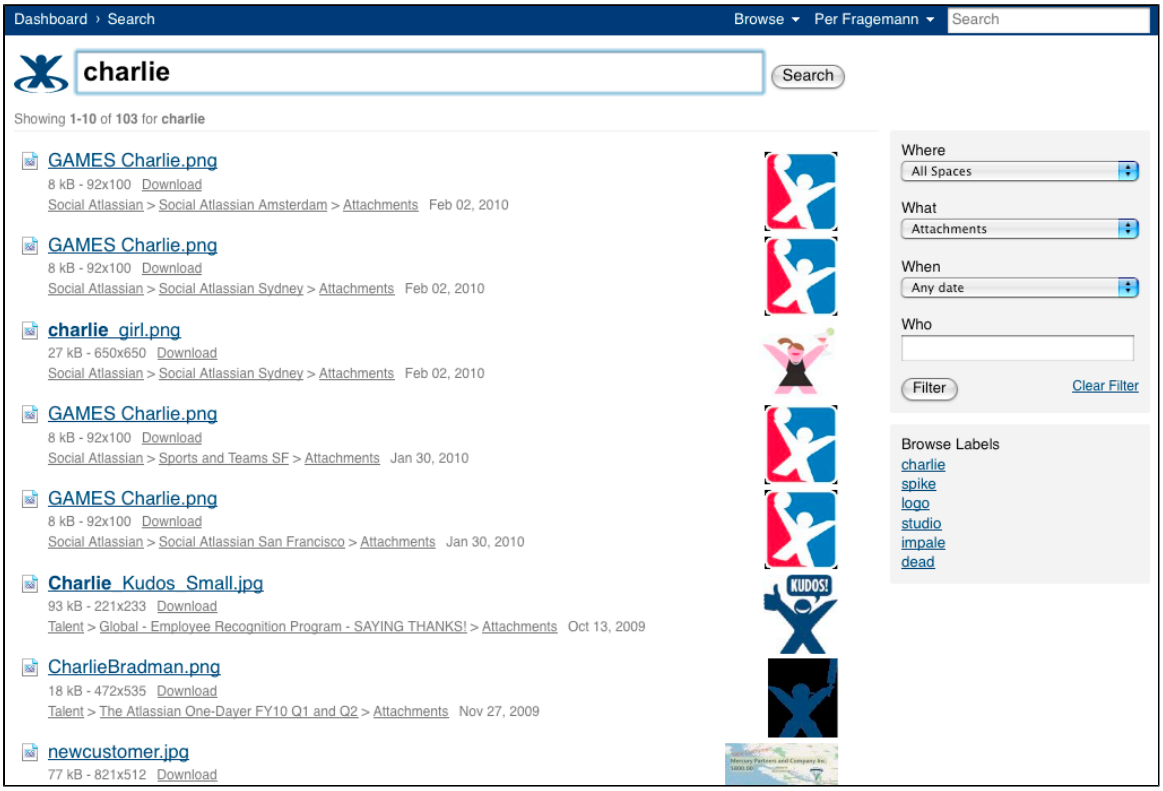
	CONF-16058	CreateRssFeedAction doesn't handle multiple RSS types in request nicely	Ryan Thomas [Atlassian]	David Taylor [Atlassian]		Resolved	Fixed	Jun 08, 2009	Mar 27, 2010	
	CONF-15317	Edit User Groups unreadable in Firefox	Brian Nguyen [Atlassian]	Maleko Taylor [Atlassian]		Resolved	Fixed	Apr 23, 2009	Mar 09, 2010	
	CONF-15180	Class Cast Exception being thrown when error encountered during mail queue flushing	Brian Nguyen [Atlassian]	Tim Wong [Atlassian]		Resolved	Fixed	Apr 13, 2009	Feb 16, 2010	
	CONF-13763	Moving attachments may silently fail to move the actual file	Ryan Thomas [Atlassian]	Chris Kiehl [Atlassian]		Resolved	Fixed	Nov 17, 2008	Mar 03, 2010	
	CONF-7473	Confluence uses the server's locale for dates and not the default language	Xu-Heng Tjhin [Atlassian]	Agnes Ro [Atlassian]		Resolved	Fixed	Dec 18, 2006	Feb 23, 2010	



Daniel's 20% project: Images in search results

If an image is found, it will be displayed in the search results. **The main point here is that we've made it possible to write plugins to spice up search results.** You could write custom display for any kind of attachment types, or even change what the regular pages look like.

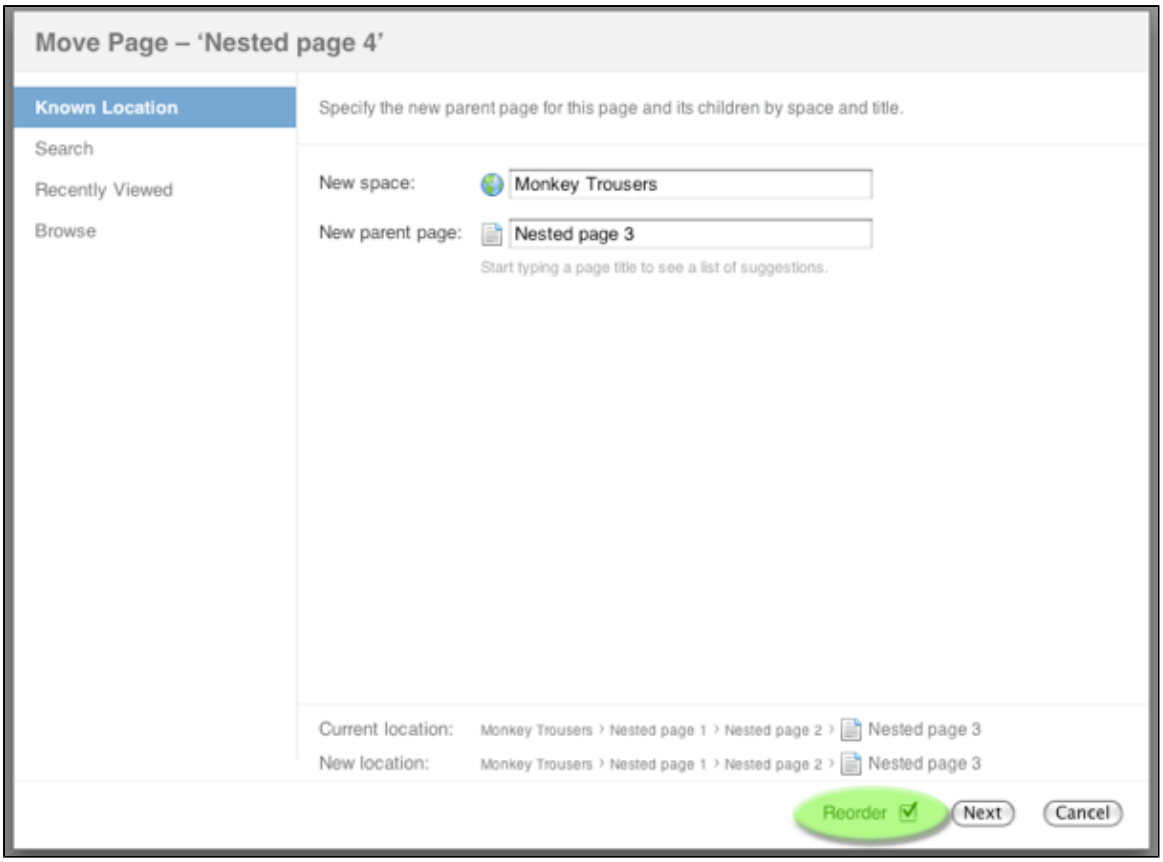
A tutorial for how to write your own renderer is located over at the Confluence Dev space



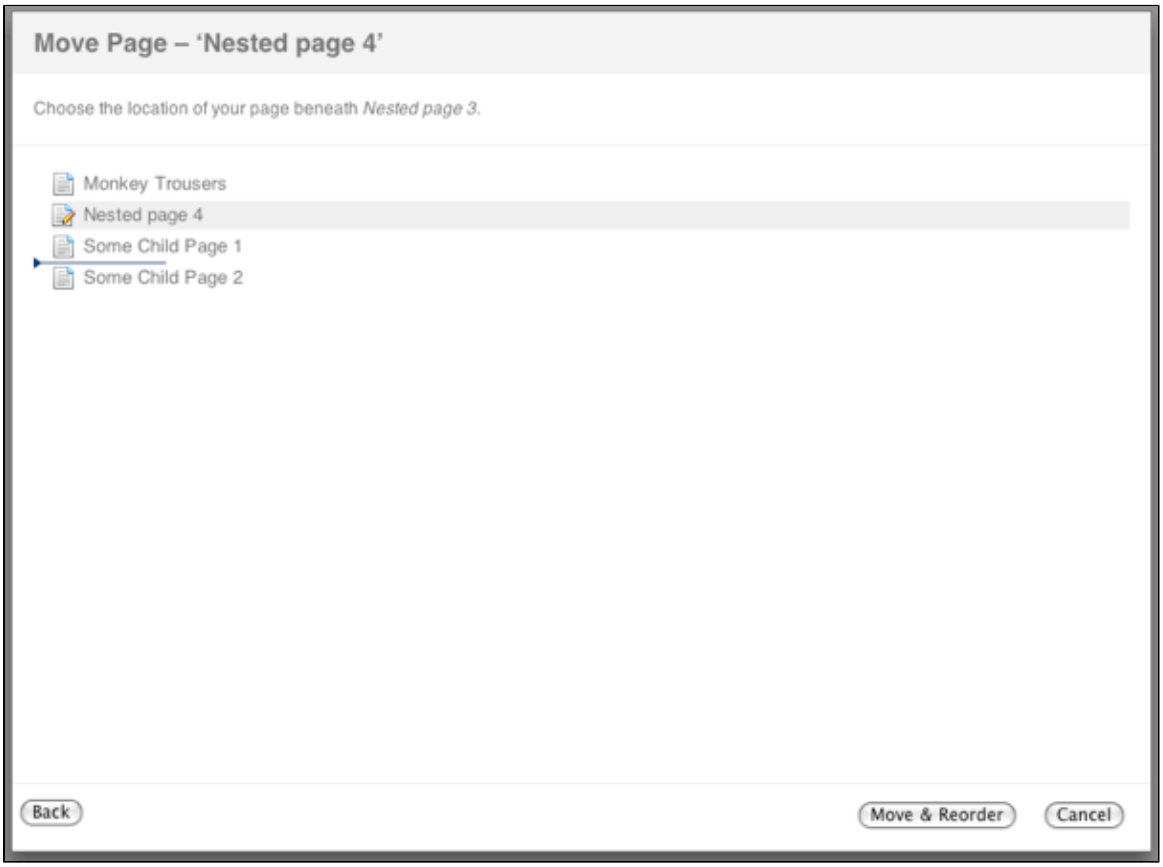


Paul's 20% project: Finishing the page move dialog

We introduced a new page move dialog in Confluence 3.0 but the ability to change the order of pages fell out of scope. This has now been implemented and is accessed as shown -



And looks like -





Small (yet awesome) Improvements

Page History



Say Hi to the new and (largely) improved Page History navigation.



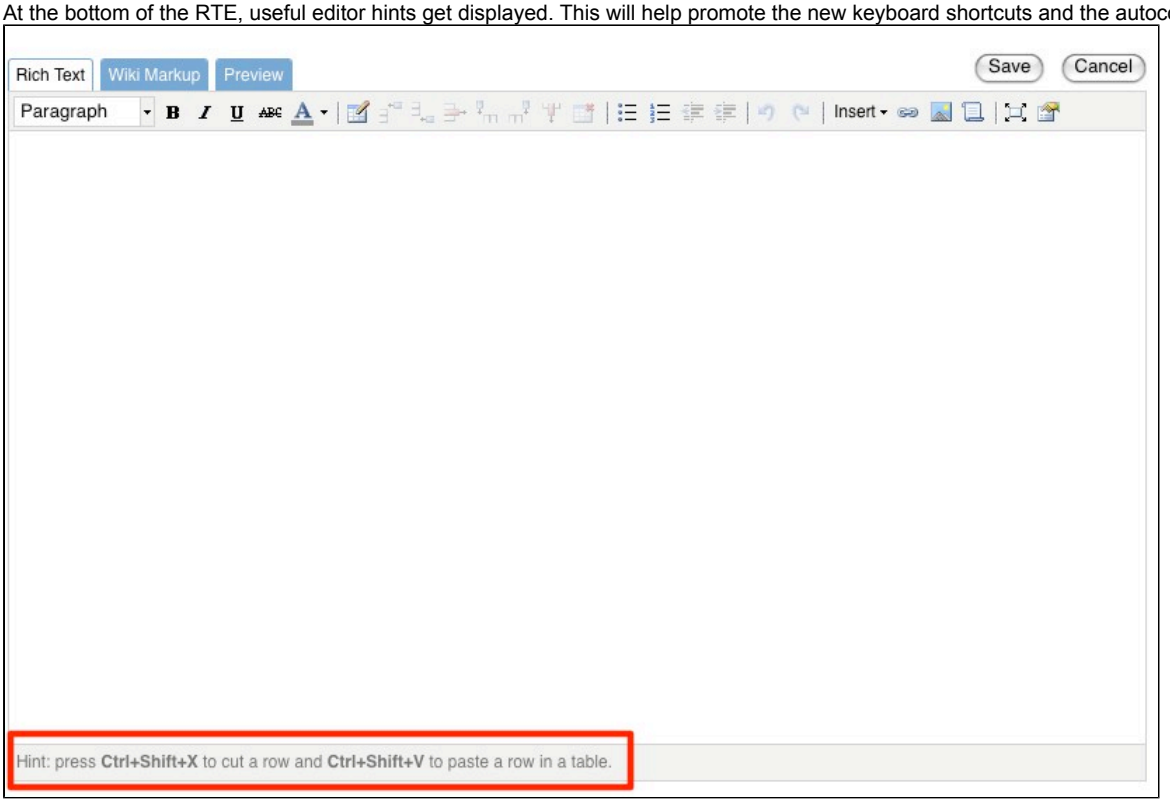
RTE Shortcuts Added

The following RTE shortcuts have been added, please try them out!

- Ctrl+Shift+C Copy table row
- Ctrl+Shift+X Cut table row
- Ctrl+Shift+V Paste table row
- Ctrl+Shift+B Bullet lists
- Ctrl+Shift+N Numbered lists
- Ctrl+Shift+F Full screen mode
- Ctrl+Shift+S Strikethrough
- Ctrl+Shift+T Insert Table
- Ctrl+Shift+A Macro Browser

RTE Hints

At the bottom of the RTE, useful editor hints get displayed. This will help promote the new keyboard shortcuts and the autocomplete feature.



Confluence 3.2 Beta Release Notes

These are not the release notes you are looking for! The final Confluence 3.2 Release Notes are over here.

Milestone release advisory

Do not use this release to upgrade your production systems.

For all production use and testing of Confluence, please use the latest official release.
This release is a public **development release** ('milestone') leading up to Confluence 3.2, which will probably ship in mid March 2010. Development releases are a snapshot of our work in progress, allowing our customers and especially plugin developers to see what we're up to.

Who should upgrade?

Please note the following

- **Development releases are not safe** — Development releases are snapshots of the ongoing Confluence development process. As such:
  - While we try to keep these releases stable, they have not undergone the same degree of testing as a full release.
  - Features in development releases may be incomplete, or may change or be removed before the next full release.
- **No upgrade path** — Because development releases represent work in progress, we **can not** provide a supported upgrade path between development releases, or from any development release to the eventual final release. Thus, it is possible that you will not be able to migrate any data you store in a Confluence development release to a future Confluence release.

In supplying milestone releases, our aim is to provide plugin developers with an opportunity to see the latest changes in the code.

Each milestone release has passed all our automatic tests and has been used for one week on our official internal Confluence server. Most of the issues solved have been reviewed too, and all of our milestone releases even have been performance tested for a while.

However, since our milestones releases are timeboxed (i.e. they get released every two weeks, no matter how far we have come implementing features and bugfixes), there is always a chance that we have new known bugs that are scheduled to be fixed in the next milestone, or completely new bugs unknown even to us.

Additionally, we have not completed our performance testing and compatibility testing for databases and application servers. So, for example, a milestone release may behave well on a small installation but show severe problems when subjected to many users.

Upgrade Procedure

Follow the normal upgrade instructions to upgrade from Confluence 3.1.x to this release. *We strongly recommend that you backup your Confluence home directory and database before upgrading!*

Downloads

All development releases are available from the development releases page on the Atlassian website.

Welcome to the Confluence 3.2 Beta program

With great pride we bring to you the Confluence 3.2 Beta release. We still have a few bugs to fix, but overall, this is how 3.2 will get rolled out to customers in mid March. And we want your feedback.

Note to partners and translators: The preliminary translation diff files can be found as a subpage of this overview page: Translations for Confluence Releases

Advance warning

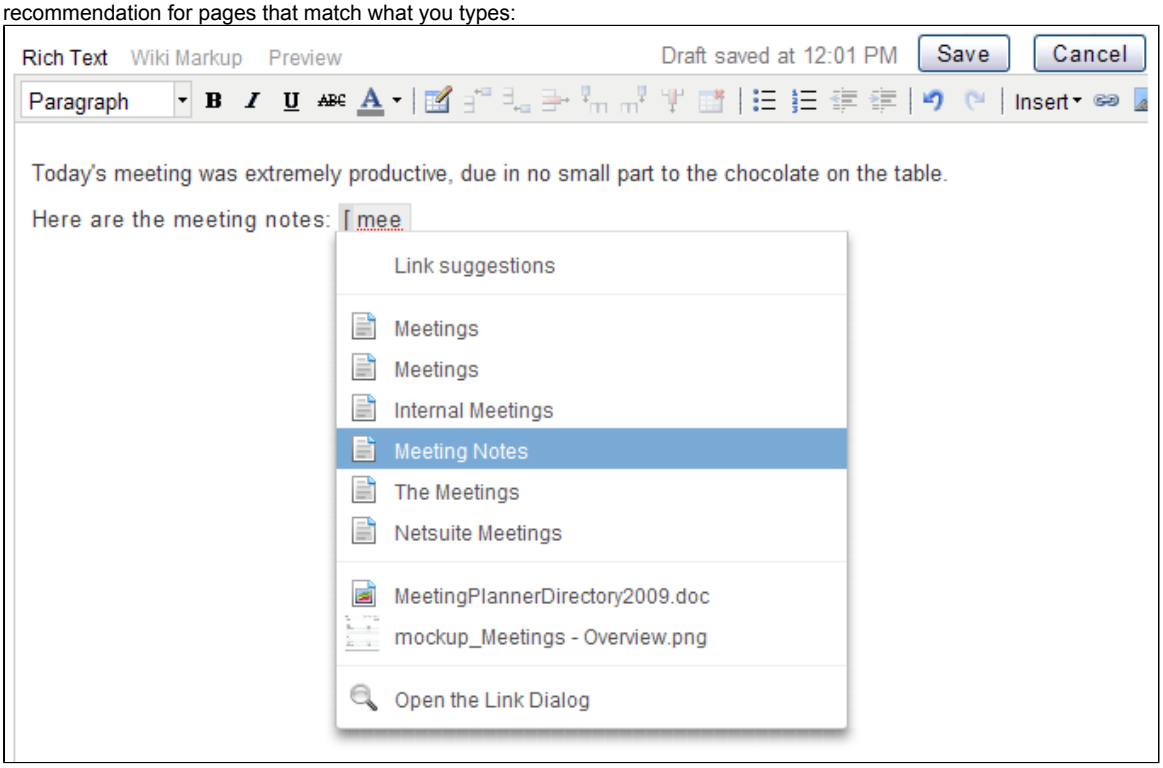
This release is the last one to officially support Internet Explorer 6. Read the upgrade notes and the original announcement over here.

This release is the last one to officially support application servers other than Tomcat. From the next release on, Websphere and Weblogic and JBoss will not be supported anymore. For further information on this, please see the end of support announcements page

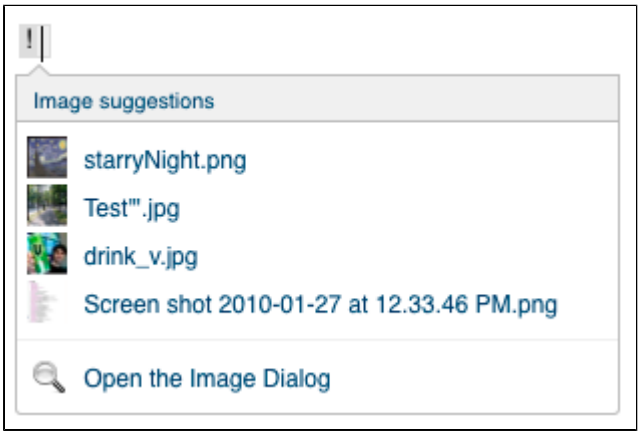


Autocomplete in the editor

We are adding a new feature called Autocomplete to the Rich Text Editor. In this beta it only works in Firefox and Safari, but the goal is to make it work on Internet Explorer as well by the time the 3.2 release ships.
You can type "**[**" at the start of a line to display your browsing history and link to those pages, and if you keep typing you'll get recommendation for pages that match what you types:



If you type "**!**" you'll see the attached images to the current page, and if you keep typing you'll be able to link to any image in your deployment.



This feature's goal is to make the RTE even faster than using wiki markup. Try it out and tell us what you think!

Note: May require a re-index of your system to show attachments

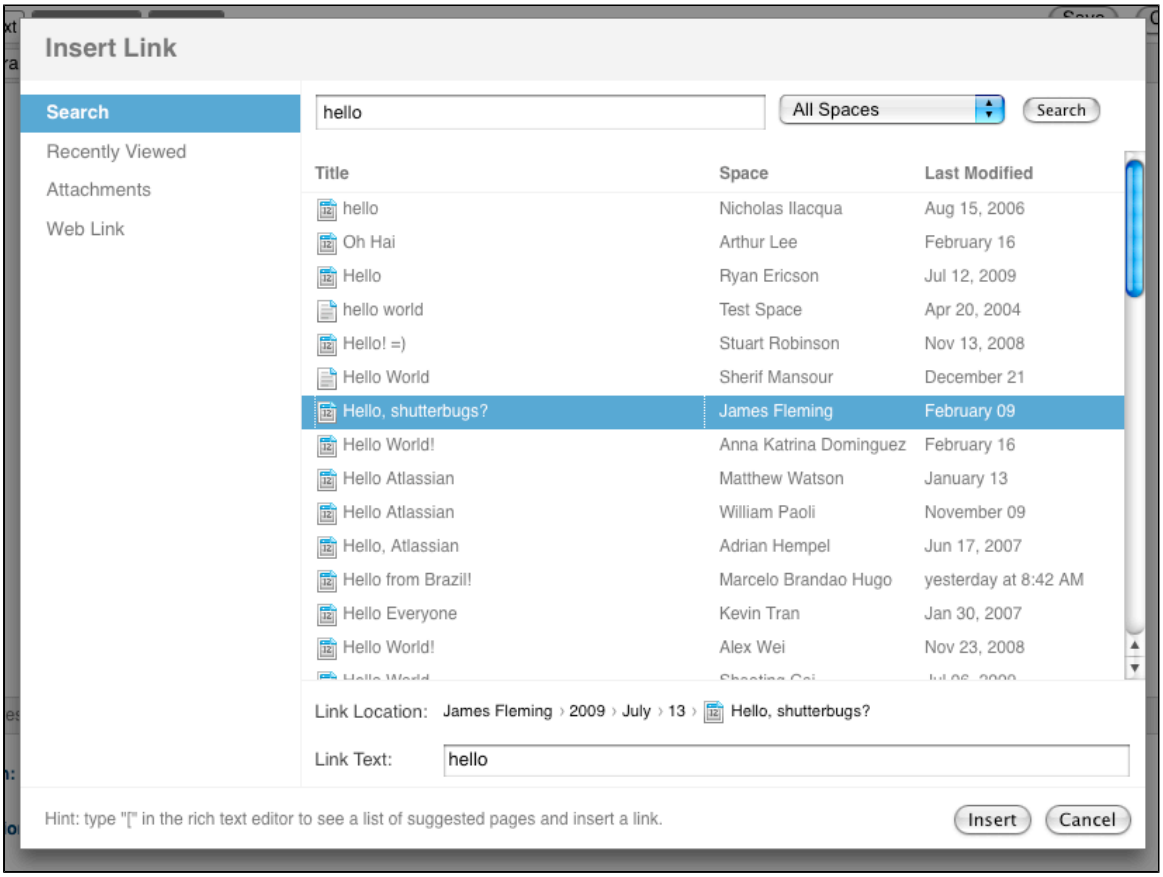
Known issues

- CONF-18281 - doesn't work in IE yet



New Links Dialog

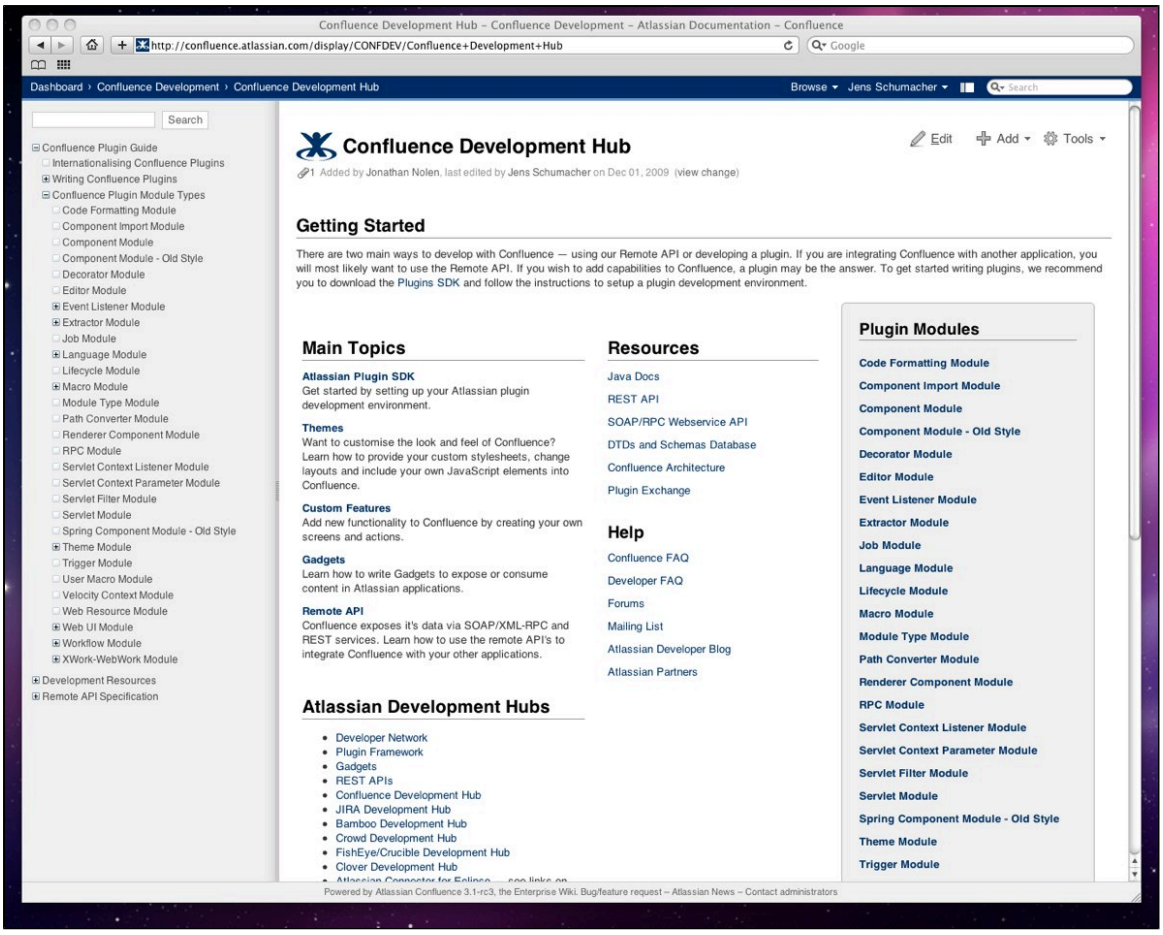
We have rewritten the way you add links in Confluence. The new links dialog is faster, better looking, and more powerful.





New "Documentation Theme"

Confluence now comes bundled with our hugely successful "Documentation Theme". See https://plugins.atlassian.com/plugin/details/16393 for details about the theme.

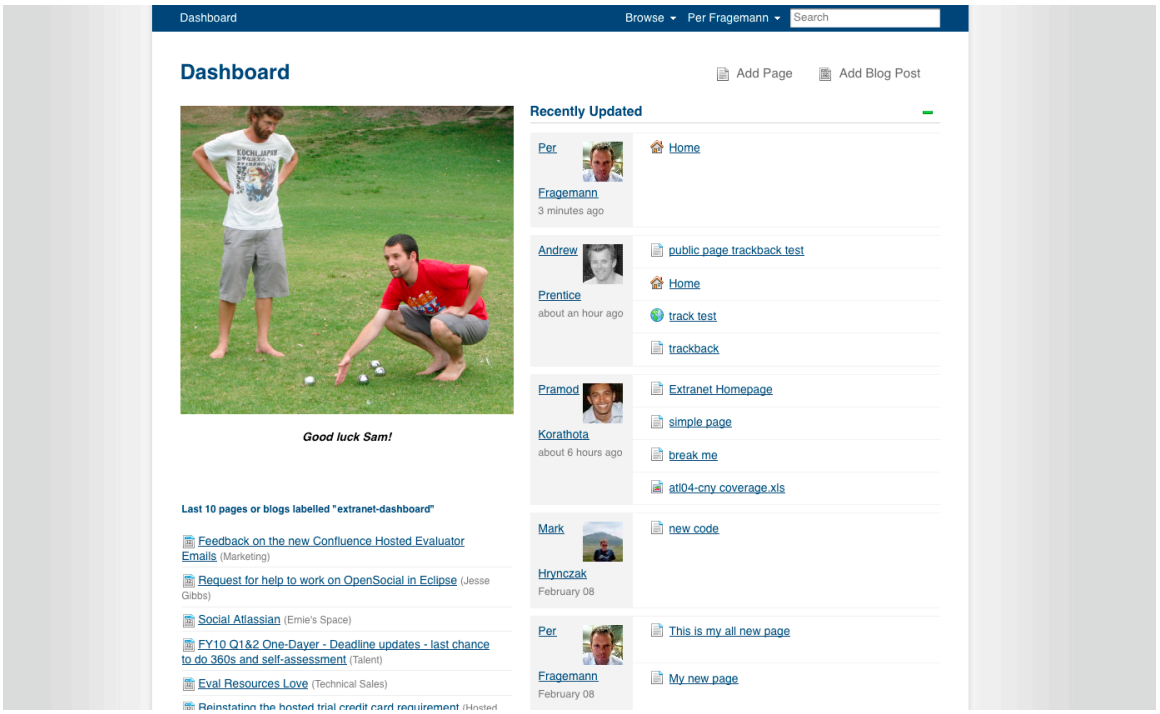


This theme replaces the old "Left Navigation Theme". We are still shipping the Left Nav theme, but it is deprecated and will stop shipping in a later release.



New "Easy Reader" Theme

With today's huge monitors, it can become hard to read large amounts of text that spans the whole screen. We are introducing a new theme which, like most websites, only uses a portion of the screen, to facilitate readability.



This theme replaces the old "Clickr Theme". We are still shipping the Clickr theme, but it is deprecated and will stop shipping in a later release.



Bundled templates and new template marketplace

The Bundled Templates 20% project is actually a framework to allow customers and Atlassian to deploy bundles of templates to a Confluence instance.

3.2 includes the framework plugin and a bundle of default templates[1]; these are available to the administrator to import into their instance as either space templates or global templates.

The templates import framework allows the administrator to preview a template before importing it as well as view the associated wiki-markup.





[1] Current default template bundle will be changed, pending list from Matt and Bill.

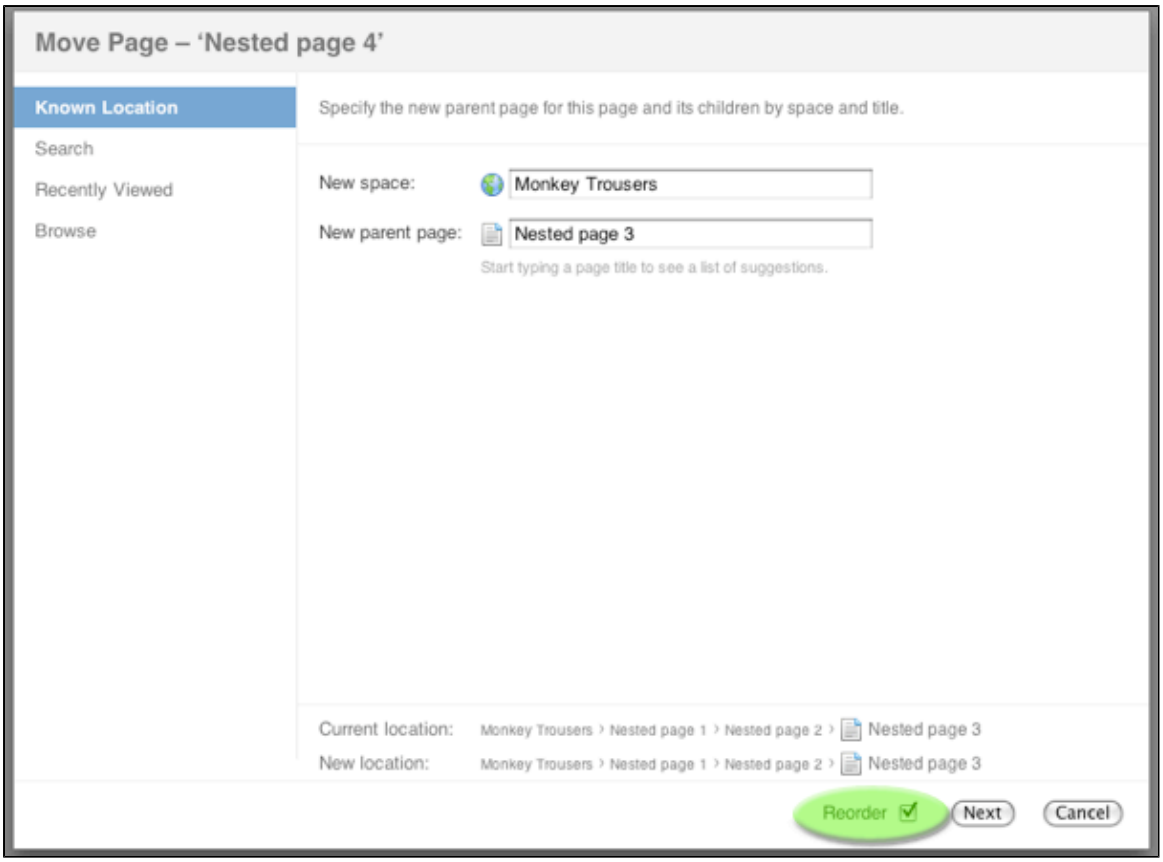
Documentation on how to write your own template plugin can be found here: Creating A Template Bundle.

Please visit PAC for the new category. TODO: add link

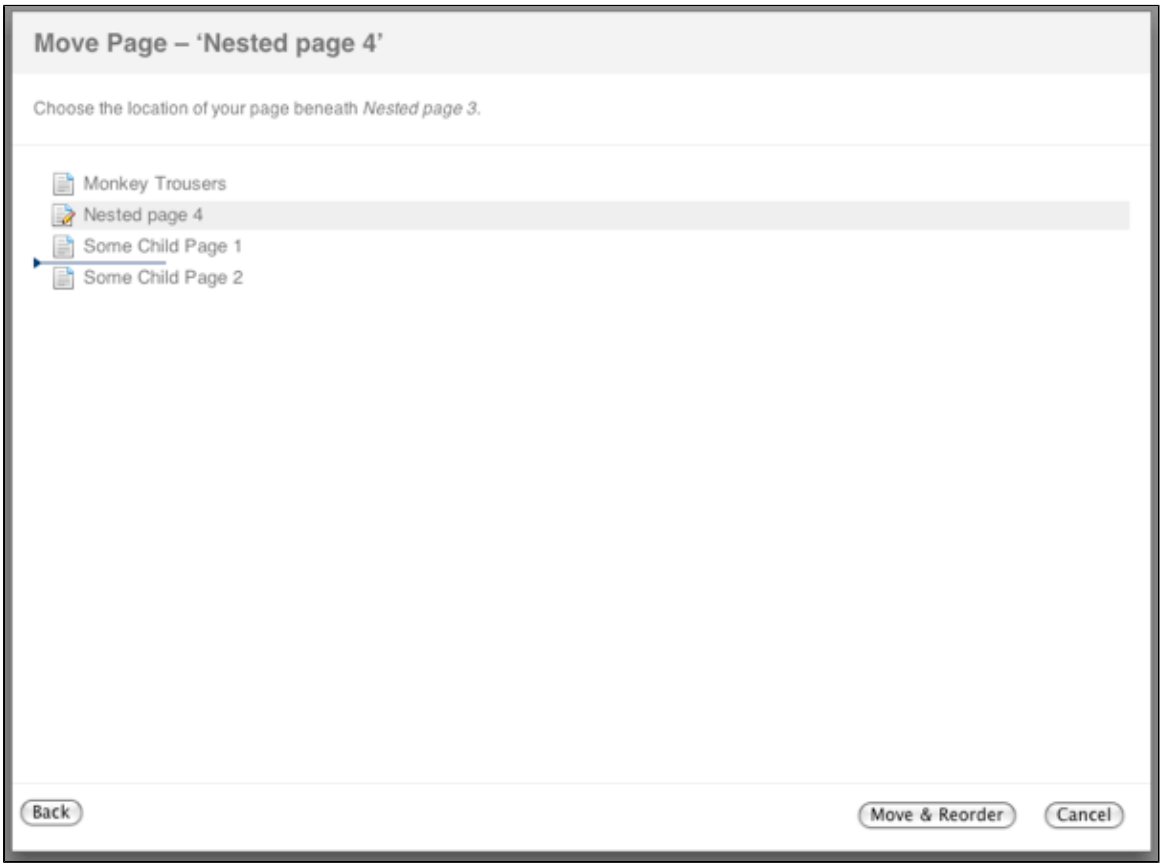


Better Move Page Dialog

We introduced a new page move dialog in Confluence 3.1 but the ability to change the order of pages fell out of scope. This has now been implemented and is accessed as shown -



And looks like -





REST APIs

New REST Resources

- Attachment Resource
  Request: https://confluence.example.com/rest/prototype/1/attachment/161

```
01.<?xml version="1.0" encoding="UTF-8" standalone="yes"?>
02.<attachment niceType="Image" version="1"
03.comment="image of Sydney harbour"
04.fileName="harbour.jpg" type="attachment" id="161">
05.<ownerId>69518087748064278</ownerId>
06.<title>harbour.jpg</title>
07.<lastModifiedDate date="2004-11-21T23:20:18-0600" friendly="Nov 21, 2004"/>
08.<createdDate date="2004-11-21T23:20:18-0600" friendly="Nov 21, 2004"/>
09.</attachment>
```

Some fields removed

More in-depth documentation coming soon.



Useful Improvements and Enhancements

Page History

What used to look like this...



.. now is a bit more understandable:



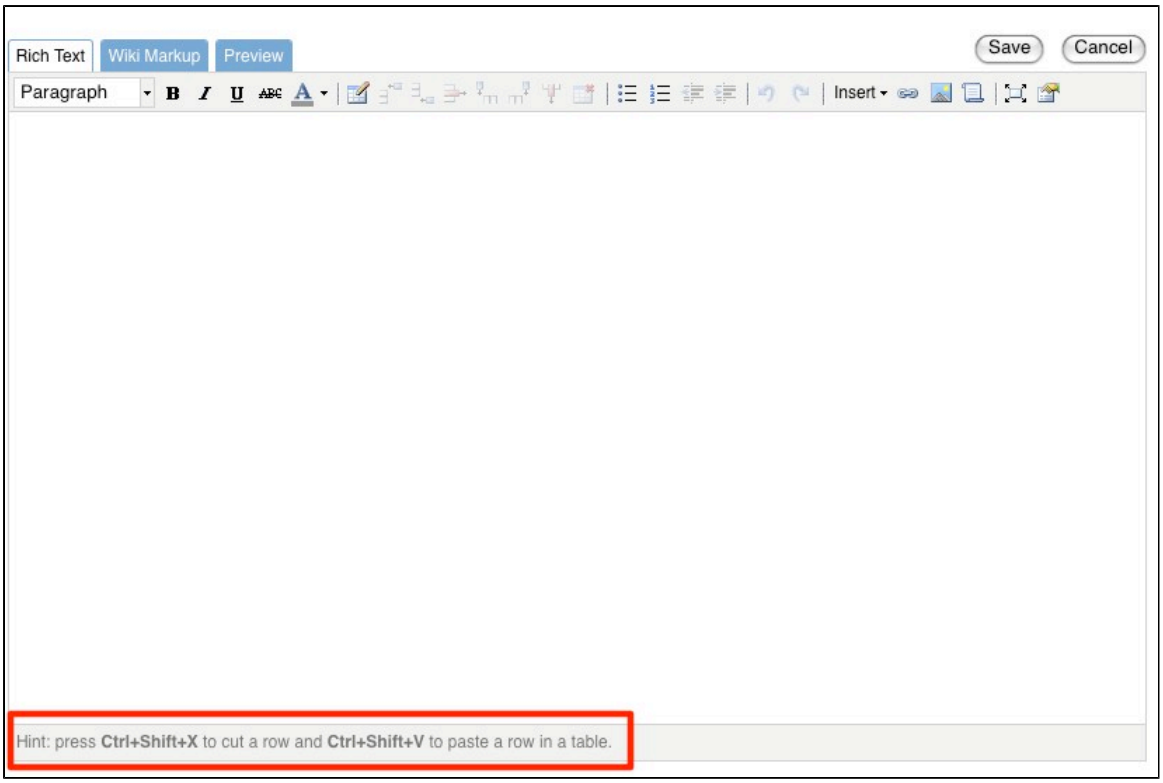
RTE Shortcuts Added

The following RTE shortcuts have been added, please try them out!

- Ctrl+Shift+C Copy table row
- Ctrl+Shift+X Cut table row
- Ctrl+Shift+V Paste table row
- Ctrl+Shift+B Bullet lists
- Ctrl+Shift+N Numbered lists
- Ctrl+Shift+F Full screen mode
- Ctrl+Shift+S Strikethrough
- Ctrl+Shift+T Insert Table
- Ctrl+Shift+A Macro Browser

RTE Hints

At the bottom of the RTE, useful editor hints get displayed. This will help promote the new keyboard shortcuts and the autocomplete feature.

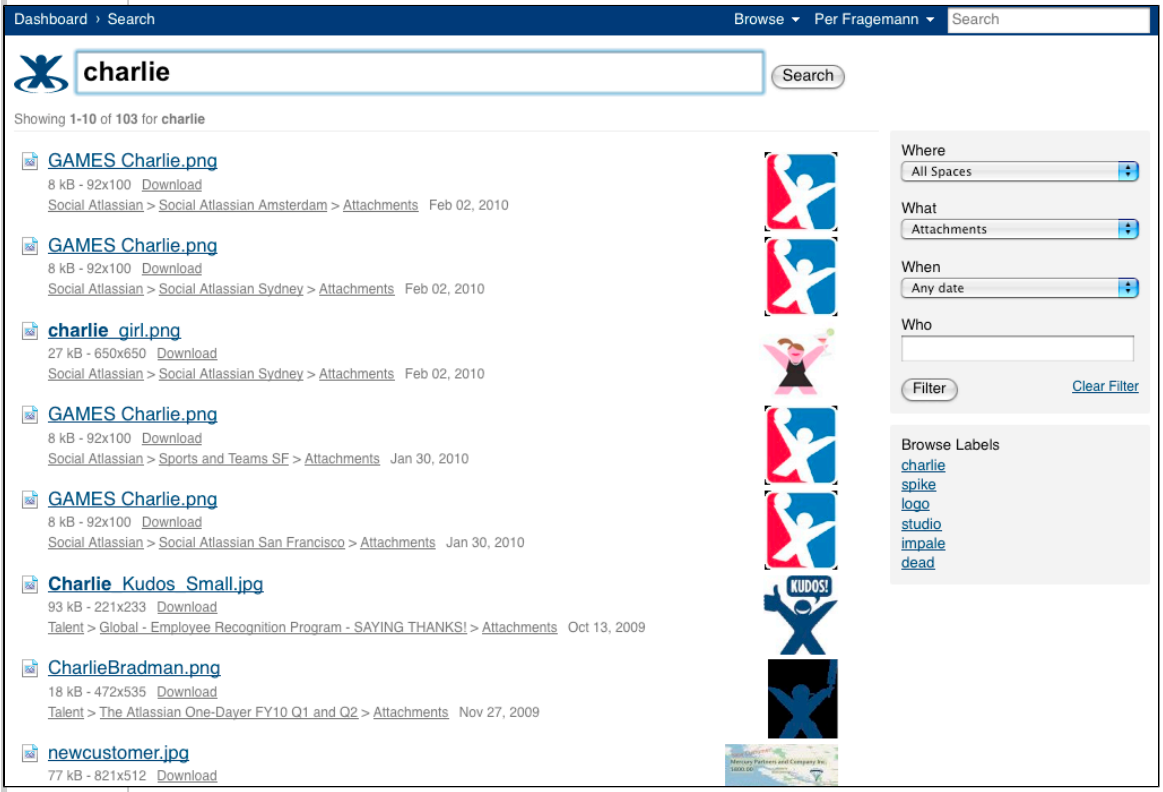


Images in search results

If an image is found, it will be displayed in the search results. The main point here is that we've made it possible to write plugins to spice up search results. You could write custom display for any kind of attachment types, or even change

what the regular pages look like.

A tutorial for how to write your own renderer is located over at the Confluence Dev space



Bundled JDBC drivers

To improve the Evaluator experience, we're now bundling a few of the most common JDBC drivers.

Improved

RSS feed builder screen

The RSS feed builder looks a bit more polished and less scary.



Bug fixes and stability improvements

A host of old bugs have been fixed. Most notable:

- it is now possible to install language packs without restarting the server.
- When adding new pages or posts, there used to be a white-space overlapping the right part of the editor. This is now fixed.
- Removed many instances of hard-coded colours (mostly blue) that prevented colour schemes applying correctly.
- Enhancement for modern browsers: added CSS to handle common forms of over-sized content, so it gets a localised scrollbar instead of making the whole page scroll
- Properly set backgrounds for email and editor popups, so they don't (incorrectly) pick up backgrounds from the main wiki window.

Coherence license changes SEPT 2009 - new Standard and Clustered Confluence Editions

Summary

Oracle Coherence (formerly known as Tangosol Coherence) is the technology that provides clustering and distributed caching in Confluence. It has also been used for caching purposes in non-clustered Confluence deployments.

The Oracle Coherence technology was first incorporated into Confluence version 2.3. Since then, Atlassian has been able to distribute the Coherence technology library files via the following means:

- Included with **all** versions and distributions of Confluence downloadable from our web site since version 2.3, regardless of whether these were intended for clustered or single-server installations.
- From the Atlassian Maven repositories.

However, Atlassian is about to enter a new license agreement with Oracle over the Coherence technology. This means that from late September 2009, Atlassian will only be be permitted to distribute the Coherence library files to customers who have purchased a license for it (that is, a Confluence clustered license).

As a result, the following changes will occur:

- The next version of Confluence (3.0.1) will be released in two editions:
  - **Standard** — Editions of Confluence without the Coherence library files. Ehcache will replace the local caching functionality previously provided by the Coherence technology.
  - **Clustered** — Editions of Confluence containing the Coherence library files.
- Customers who have purchased a non-clustered Confluence license will only be able to download standard editions of Confluence from the Atlassian web site, whereas customers who have purchased a Confluence clustered license will be able to download clustered editions of Confluence.
- **From late September 2009:**
  - **Standard editions of Confluence will be made available for each previous major releases of Confluence back to version 2.6.** These will be available as Confluence versions 2.10.4, 2.9.3, 2.8.3, 2.7.4 and 2.6.3 and will be available to customers with non-clustered licenses.
  - **All other previous versions of Confluence currently available from our download page (from 2.6 to 3.0 inclusively), will be re-released as clustered editions and will only be available to customers with Confluence clustered licenses.**
  - The Coherence library files will no longer be available in any form from the Atlassian Maven repositories.
  - The installation files for all versions of Confluence prior to 2.6 (which are no longer supported) will be removed from the Atlassian web site and will no longer be available for download and installation.



If you are currently running a **clustered installation of Confluence**, please do not upgrade it with a **standard edition** of Confluence.

What are the implications?

I am a Confluence customer with a non-clustered Confluence license, running Confluence 2.3 or later.

The Confluence distribution you are running will continue to function and if it is Confluence version 2.6 or later, be supported by Atlassian in accordance with our standard support policy.

However, if you upgrade to Confluence version 3.0.1 or later or obtain any Confluence version released after late September 2009, you will only be able to download and upgrade to standard editions of Confluence.

I run a customised installation of Confluence 2.3 or later and must build Confluence from source.

Confluence source code downloaded before late September 2009 requires that the Coherence library files are present in either your local or the Atlassian Maven repositories for automated Maven builds to complete successfully.

If you have Confluence source code downloaded before late September 2009 (excluding version 3.0.1) but conduct an automated Maven build of Confluence using this source code after this date, your build will fail if the Coherence library files are not available in your local Maven repository. This is because the Coherence library files will also not be available in the Atlassian Maven repository.

Hence, to build a customised installation of Confluence using this source code, we recommend that you locate the `tangosol-3.3.jar` and `coherence-3.3.jar` from the `WEB-INF/lib` directory of **your own** existing Confluence installation and install them into your local Maven repository using the following commands:

```
1.mvn install:install-file -Dfile=tangosol-3.3.jar -DgroupId=tangosol-coherence \
2.-DartifactId=tangosol -Dversion=3.3 -Dpackaging=jar
3.mvn install:install-file -Dfile=tangosol-3.3.jar -DgroupId=tangosol-coherence \
4.-DartifactId=coherence -Dversion=3.3 -Dpackaging=jar
```



These commands will install the Coherence library files into your *local* Maven repository, which should be available to you only. Please do not upload these files to any public Maven repository nor make them publicly available by any other means. Atlassian's End User License Agreement does not grant permission to redistribute any part of Confluence, which includes these Coherence library files.

Alternatively, you can download the sources for one of the new standard or clustered editions of Confluence and reapply your customisations. These will compile without any additional problems.

I am a plugin developer and wish to compile plugins against old or existing versions of Confluence

This will affect plugin developers in two ways:

1. When building a plugin, Maven will complain about the absence of the Coherence library files in the Atlassian Maven repository.
2. When using the Atlassian Plugin Toolkit to test a plugin, Maven will be unable to download the appropriate Confluence EAR-WAR distribution file (from the Atlassian Maven repository), against which to perform integration tests.

Therefore, we recommend that as soon as possible, plugin developers start compiling their plugins based on the new standard editions of Confluence (without the Coherence library files). Standard editions of Confluence will be binary-compatible with clustered editions and existing clustered instances of Confluence. Hence, plugins developed against standard editions of Confluence will also run on any clustered editions and existing clustered instances of Confluence.

I am a plugin developer concerned about API changes and multiple Confluence editions resulting from these changes

As long as you are using only Confluence APIs to develop plugins, your plugins will be binary compatible with both standard and clustered editions of Confluence. The interfaces of the Confluence `CacheManager`, `Cache` and `ClusterManager` will be the same in both editions of Confluence, although there will be only one important change from previous versions of Confluence.

In existing versions of Confluence, the `ClusterManager` exposes the Coherence `InvocationService` to allow clients to execute code or perform queries across all nodes of the cluster. This API will be unavailable in all standard edition versions of Confluence from version 3.0.1 back to 2.6. If your plugin uses this service (and Atlassian is not aware of any that do) you should instead use Confluence's `RemoteEvent` API to send messages to other cluster nodes and direct them to perform the work.

Also, if for some reason your plugin references some other Coherence classes directly (or imports the Coherence-specific implementations of the `CacheManager`, `Cache` or `ClusterManager`), you will need to rewrite your plugin to use the generic interfaces only.

I am a plugin developer and want to test my plugin against Confluence in a cluster

For testing purposes, you must own a Confluence clustered license and have access to a clustered Confluence installation (either an existing one or one based on the new Confluence clustered edition).

Upgrade Notes Overview

All upgrades between Confluence versions have a lot in common. Please refer to our general information about upgrading Confluence first.

Typically, each major release of Confluence comes with some specific recommendations for upgrading from the previous major version. If you plan to upgrade and skip a few Confluence versions, you must read the upgrade notes for all major versions between your current version and the version to which you are upgrading, to make sure you do not miss something important.



For example:

If you plan to upgrade from Confluence 2.8 to Confluence 3.0, read the upgrade notes for Confluence 2.9 and Confluence 2.10, as well as those for Confluence 3.0.

Also, we strongly recommend that you read the upgrade notes for any minor releases below, since they contain important information that will affect your Confluence upgrade.

Below is a list of upgrade notes for previous major releases of Confluence, as well as the upgrade notes for important minor releases:

- Confluence 3.2 Upgrade Notes
- Confluence 3.1 Upgrade Notes
- Confluence 3.0.1 Upgrade Notes
- Confluence 3.0 Upgrade Notes
- Confluence 2.10 Upgrade Notes
- Confluence 2.9 Upgrade Notes
- Confluence 2.8 Upgrade Notes
- Confluence 2.7 Upgrade Notes
- Confluence 2.6 Upgrade Notes

You will find the upgrade notes attached to the release notes for the relevant version.

RELATED TOPICS

Confluence Release Summary
Release Notes

Confluence Configuration Guide

The pages listed below contain instructions on configuring Confluence. If you cannot find what you are looking for, try the Search box in the left hand navigation.

- Application Server URL encoding
    - Configuring Tomcat's URI encoding
- Database Configuration
    - Migrate to Another Database
    - Database Setup Guides
    - Creating Database Schema Manually
    - Known Issues For Supported Databases
    - Improving Database Performance
    - Troubleshooting External Database Connections
    - Troubleshooting the Embedded Database (hSQL DB)
    - Upgrading From HSQL 1.7.1 to 1.8
- Guide to using Apache Tomcat's Virtual Hosts
- How to dump Active Directory data to a file
- Known Issues with Enterprise or Webhosting environments
- Managing Application Server Memory Settings
- Pull down RSS Feeds or use the Repository plugin through a web proxy
- Running Confluence behind Apache
    - General Apache Configuration Notes
    - Using Apache with mod_jk
    - Using Apache with mod_proxy
    - Using Apache with virtual hosts and mod_proxy
- Setting Up Public Access
- Setup a mail session in standalone version
- Start Confluence automatically on system startup
    - Start Confluence automatically on Linux & Unix
    - Start Confluence automatically on OS X using launchd
    - Start Confluence automatically on Windows as a Service
- The Confluence data model
- Troubleshooting SQL Exceptions
- Webserver Configuration
    - Apache and Apache Connector Tips
    - Configure Web Proxy Support for Confluence

RELATED CONTENT

Documentation Home
Confluence Administrator's Guide
Frequently Asked Questions
Confluence Community
Development Hub
Plugins and Extensions

Application Server URL encoding

Application servers may have different settings for character encodings. We strongly suggest setting this to **UTF-8** where possible.

Information on setting the character encoding is available at:

- Configuring Tomcat's URI encoding

Configuring Tomcat's URI encoding

By default, Tomcat uses ISO-8859-1 character encoding when decoding URLs received from a browser. This can cause problems when Confluence's encoding is UTF-8, and you are using international characters in attachment or page names.

1. Edit `conf/server.xml` and find the line where the `Coyote HTTP Connector` is defined. It will look something like this, possibly with more parameters:

```
1.<Connector port="8080"/>
```

2. Add a `URIEncoding="UTF-8"` property to the connector:

```
1.<Connector port="8080" URIEncoding="UTF-8"/>
```

3. Restart Tomcat

If you are using mod_jk

You should apply the same URIEncoding parameter as above to the AJP connector if you are using mod_jk, and add the following option to your Apache mod_jk configuration:

```
1.<Connector port="8009" protocol="AJP/1.3" URIEncoding="UTF-8"/>
```

```
1.JkOptions +ForwardURICompatUnparsed
```

More information using Apache with Tomcat

For comprehensive examples of how to use Tomcat and Apache with Confluence, see Running Confluence behind Apache.

Database Configuration

This document provides information on configuring an external database.

The Embedded Database for Evaluation Purposes

The Confluence installation includes an embedded HSQLDB database, supplied for the purpose of **evaluating Confluence**.

If you are using the embedded database, the database files are stored in the `\database` folder under your Confluence Home directory. See also Important Directories and Files.



Embedded Database is Not Suitable for Production Instances of Confluence
Production instances of Confluence should use an external database. When using the default HSQLDB database, you run the risk of unrecoverable data loss due to not being transaction safe.

- Corruption is occasionally encountered after sudden power loss and can usually be corrected using this data recovery procedure.
- HSQLDB is still suitable for evaluation purposes, but the risk can only be eliminated by switching databases. External databases may also provide superior speed and scalability.

Selecting an External Database



The XML backup built into Confluence is not well suited for database migration for large data sets (see [Alternative Backup Strategy] for reference). Choose your database wisely; if you need to migrate later, Atlassian support will refer you to a third party database migration tool. Vote for CONF-12599 to add a more robust strategy for large implementation migrations.

Below is more information on selecting and migrating to an external database:

- Migrating to a Different Database
- List Of Supported Databases
- Known Issues For Supported Databases

Database Setup

Setup instructions are shown below. Click a database to review the setup guide and any known issues.

MySQL

PostgreSQL

Oracle

SQL Server

DB2

Other databases should use these instructions.

Optimising Database Performance

To improve database responsiveness:

- Improving Database Performance
- Known Issues For Supported Databases

Database Troubleshooting

For solving database-related problems:

- Troubleshooting External Database Connections
- Troubleshooting the Embedded Database (hSQL DB)
- Interpreting DB2 error codes
- Known Issues For Supported Databases

Obtain technical support from Troubleshooting Problems & Requesting Technical Support.

Migrate to Another Database

Limitations of Database Migration



The XML backup built into Confluence is not well suited for database migration for large data sets (see Production Backup Strategy for reference). If the procedures below do not work, use a commercial database migration tool. Vote for CONF-12599 to add a more robust strategy for large implementation migrations. Atlassian does not support migrating to a new database.

Database Migration

This document outlines how to migrate your data from your existing database to another database. It is designed for migrating from an evaluation to a production database. Large data sets will require third party database migration tools. You should use this method when moving from the embedded DB to an external DB, or from one type of external DB to another (e.g. Oracle to Postgres).



If you are simply moving your DB from one server to another you can just change the JDBC URL in <confluence.home>/confluence.cfg.xml (if you are using a direct JDBC connection) or in the definition of your datasource (if you are connecting via a datasource).

There are two ways you can perform the migration:

1. Method one is the standard procedure.
2. For large installations of Confluence using version 2.2 or later: If the total size of attachments on your installation exceeds 500MB,

use method two.

On this page:

- Method One - Standard Procedure
  - Step One: Backing up your data
  - Step Two: Configuring the Confluence Home Directory
  - Step Three: Setting up the new database
  - Step Four: Setting up Confluence with the new database
- Method Two - For large installations
  - Step One: Backing up your data
  - Step Two: Configuring the Confluence Home Directory
  - Step Three: Moving your attachments
  - Step Four: Setting up new database
  - Step Five: Setting up Confluence with the new database
- A Note about Case Sensitivity in your Database
- Troubleshooting

Method One - Standard Procedure

Step One: Backing up your data

1. Create a backup of your existing data. This is done from the **Administration Console**. Instructions on how to create a backup can be found here.
2. Shut down and backup the Confluence Home Directory.
3. If you are already using an external database, please make a backup of it using the utilities that were installed with it.



Note which plugins are currently installed/enabled, so that you can reinstate them later.

Step Two: Configuring the Confluence Home Directory

1. Create a new Confluence Home Directory. You can place this directory anywhere you like and give it a name of your choice.
2. Open WEB-INF/classes/confluence-init.properties file in your Confluence installation and change the confluence.home property to point to this new Confluence Home Directory.

Step Three: Setting up the new database

Perform the database setup instructions for your database.

Step Four: Setting up Confluence with the new database

If your databases are not already configured for Confluence, refer to Database Configuration to setup your database access.

1. Make sure that the JDBC drivers for your database are available to the application server. If you don't already have the JDBC driver, please download one from here.
2. Make sure that your database is using a **case-sensitive** collation. Please refer to the section on case sensitivity below and see this issue for more details: CONF-7917.
3. If you are running the standalone version of Confluence, copy your JDBC database driver (a .jar file), into the `<confluence-install>/lib` folder.
4. Start up Confluence. You will see the Confluence Setup Wizard.
5. Select '**Custom Install**'.
6. Select your database from the drop down list.
7. Select '**Direct JDBC**' and then enter the details of the new database you want to migrate to.

   Read the documentation on the Setup Wizard for more detailed explanation.
8. When prompted, restore the contents of the backup you made in Step One into the new Confluence site

Your old Confluence data will now be imported to your new database.

Method Two - For large installations

Step One: Backing up your data

1. Before proceeding with these instructions please check that:
   - you are upgrading from at least Confluence version 2.2 *and*
   - your attachments are stored in the file system, and *not in your database.* (To migrate between attachment storage systems, please see Attachment Storage Configuration)

     These instructions will not work if either of the above is not true.
2. From Confluence, go to `Administration -> Backup & Restore` and create a manual backup that **excludes** attachments.
3. Shut down and back up the Confluence Home Directory.
4. If you are already using an external database, please make a backup of it using the utilities that were installed with it.

Step Two: Configuring the Confluence Home Directory

1. Create a new Confluence Home Directory. You can place this directory anywhere you like and give it a name of your choice.
2. Open WEB-INF/classes/confluence-init.properties file in your Confluence installation and change the confluence.home property to point to this new Confluence Home Directory.

Step Three: Moving your attachments

Move the contents of your attachments directory from your old Confluence Home to your new Confluence Home.

Step Four: Setting up new database

Perform the database setup instructions for your database.

Step Five: Setting up Confluence with the new database

If your databases are not already configured for Confluence, refer to Database Configuration to setup your database access.

1. Make sure that the JDBC drivers for your database are available to the application server. If you don't already have the JDBC driver, please download one from here.
2. Make sure that your database is using a **case-sensitive** collation. Please refer to the section on case sensitivity below and see this issue for more details: CONF-7917.
3. If you are running the standalone version of Confluence, copy your JDBC database driver (a .jar file), into the `<confluence-install>/lib` folder.
4. Start up Confluence. You will see the Confluence Setup Wizard.
5. Select '**Custom Install**'.
6. Select your database from the drop down list.
7. Select '**Direct JDBC**' and then enter the details of the new database you want to migrate to.
   Read the documentation on the Confluence Setup Wizard for more detailed explanation.
8. When prompted, restore the contents of the backup you made in Step One into the new Confluence site.

A Note about Case Sensitivity in your Database

'Collation' refers to a set of rules that determine how data is sorted and compared. Case sensitivity is one aspect of collation. Other aspects include sensitivity to kana (Japanese script) and to width (single- versus double-byte characters).

Case-sensitive or case-insensitive collation — how should you create your Confluence database? What about when you are migrating your existing Confluence instance from one database to another?

Setting up a New Confluence Instance

For new Confluence instances, we recommend using case-sensitive collation for your Confluence database, which is the default collation type used by many database systems. The Confluence application itself reduces all usernames into lower-case characters before they are stored in the Confluence database. Therefore, 'joebloggs', 'joeBloggs', 'JoeBloggs', etc. will be treated as the same username on a Confluence installation with case-sensitive database collation.

Migrating an Existing Confluence Instance to a Different Database

The default Confluence Standalone configuration uses case-sensitive database collation. This is often the case with databases on several other systems which were created under default conditions. Therefore, if you are migrating from this type of configuration to a new database, we recommend that the new database uses case-sensitive collation. If you use case-insensitive collation, you may encounter data integrity problems after migration (for example, via an XML import) if data stored within your original Confluence site required case-sensitive distinctions.

Troubleshooting

If you're unable to restore your XML backup, consult our Troubleshooting Guide.

RELATED TOPICS

Migrate to Another Database

Confluence Configuration Guide

Administrators Guide Home Confluence Documentation Home

The Embedded HSQLDB Database

The Confluence installation includes an embedded HSQLDB database, supplied for the purpose of **evaluating Confluence**.

If you are using the embedded database, the database files are stored in the `\database` folder under your Confluence Home directory. See also Important Directories and Files.



Embedded Database is Not Suitable for Production Instances of Confluence
Production instances of Confluence should use an external database. When using the default HSQLDB database, you run the risk of unrecoverable data loss due to not being transaction safe.

- Corruption is occasionally encountered after sudden power loss and can usually be corrected using this data recovery procedure.
- HSQLDB is still suitable for evaluation purposes, but the risk can only be eliminated by switching databases. External databases may also provide superior speed and scalability.

RELATED TOPICS

Important Directories and Files
Database Configuration

Database Setup Guides

- Database Setup For Any External Database
- Database Setup for PostgreSQL
- Database Setup For MySQL
- Database Setup for Oracle
- Database Setup for SQL Server
- Database JDBC drivers

Database Setup For Any External Database

If you are using Confluence in a production environment, data should be stored in an external database. The embedded database is bundled for evaluation purposes and does not offer full transactional integrity in the event of sudden power loss.

This document provides instructions for setting up Confluence for use with a production-ready database. It covers both migration from an evaluation installation of Confluence and installation of an empty database during initial setup. The following specific database guides have additional information:

- PostgreSQL Guide
- MySQL Guide
- Oracle Guide
- SQL Server

Preparation

Install the following on the Confluence server:

- Database administration tool, for example DBVisualizer
- JDBC database drivers
- The database server (unless accessed remotely)

The instructions refer to two particular directories:

- The <*Confluence Installation Directory*> is the directory where you unpacked the Standalone Confluence download.
- The <*Confluence Home Directory*> is the directory where Confluence stores its data, which you set by editing the `confluence-init.properties` file in *Confluence Installation Directory*`/confluence/WEB-INF/classes`.

Database Setup

Create the schema and setup permissions:

1. Visit the Database Configuration page to review any known issues and database setup for your database.
2. Create a new schema using the correct database encoding.
3. Create a user with full read/write access to the Confluence schema, including the ability to create tables.
4. If the database only permits users to log in from approved hosts (e.g. `localhost`), grant database access permission for the Confluence server.
5. If the database is hosted remotely to the Confluence server, set up any firewall permissions.
6. Test the connection by using the database administration tool installed on the Confluence server to log in to the database.

Migration From an Evaluation Instance of Confluence
Continue here if you are migrating from an evaluation instance with the built-in database. If you are installing Confluence for the first time, continue below.

Create Backups

To keep any existing Confluence content:

1. If you are already using an external database, use your database administration tool to create a full database backup.
2. Manually create an XML backup of Confluence under '**Administration**' -> '**Backup & Restore**'. If you have less than 100MB of attachments, check '**Backup attachments**' when creating the backup. If you have over 100MB of attachments, you should not check the 'Backup attachments' and instead you should manually copy the `/attachments` directory in your Confluence home to a backup location. This attachments directory will later be copied into the new home directory.
3. Download the backup file to a backup location.

Database Connection Setup

Set up Confluence's database connection:

1. Stop Confluence if it is already running.
2. The JDBC database drivers for your database must be available to the application server. You can skip this step if the drivers are already loaded.
    a. Copy the database driver JAR file into the `lib` directory. In Confluence Standalone this directory is `/confluence/WEB-INF/lib`. Other application servers will use a different path.
    b. If the application server does not support dynamic library loading, stop your application server.
3. Create a new Confluence home directory.
4. Open the `WEB-INF/classes/confluence-init.properties` file in your Confluence installation and change the `confluence.home` property to point to this new Confluence home directory.
5. Start up Confluence. Refer to the platform-specific installation instructions to learn how. You should be presented with the **Confluence setup wizard**. Enter your license information.
6. Select '**Custom install**'.
7. Select a database from the drop down list.
8. Select `Direct JDBC` and then enter the username, password and database driver of the new database.
9. If you created a Confluence backup earlier and wish to restore it, import it into Confluence now.
10. Once the wizard is complete, if you did not check the 'Backup attachments', copy the backed up `/attachments` directory into the new Confluence home.

RELATED TOPICS

Troubleshooting External Database Connections

Database Setup for PostgreSQL

This document provides instructions for setting up Confluence for use with a PostgreSQL database. Please check the Known Issues for PostgreSQL before you start.

On this page:

- 1. Install PostgreSQL
- 2. Create a User and a Database
- 3. Configure Confluence to use the PostgreSQL Database
- Troubleshooting

1. Install PostgreSQL

To install PostgreSQL,

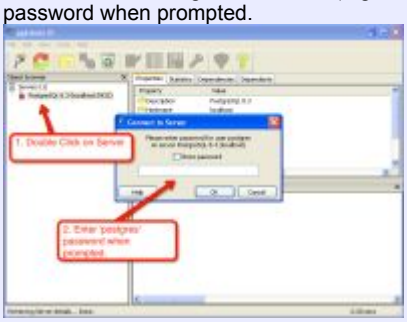
1. Download the database software and installer from the PostgreSQLdownload site and save it to your desktop. Choose the package that matches your operating system. Where available, choose the **One Click Installer**. These instructions assume you will use the One Click Installer. For example:
    - PostgreSQL One Click Installer for Windows.
    - PostgreSQL One Click Installer for Linux.
    - PostgreSQL One Click Installer for Mac OS X.
2. Run the installer. Please note the following information when installing PostgreSQL:
    - The **password** that you are prompted to provide during the installation process is for the 'postgres' account, which is the db root level account.
    - The **default port** for PostgreSQL is 5432. If you decide to change the default port, please ensure that your new port number does not conflict with any services running on that port. You will also need to remember to update all further mentions of db port.
    - Choose the **locale** that best fits your geographic location, when prompted to enter a locale.
    - Do not launch **Stack Builder** at the completion of the installer.
3. PostgreSQL is now installed on your machine.

2. Create a User and a Database

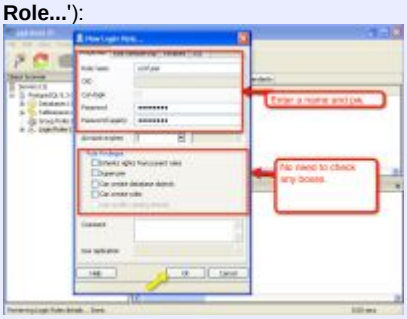
All screenshots below are taken from a PostgreSQL configuration on a Windows machine.

To create a PostgreSQL user and database,

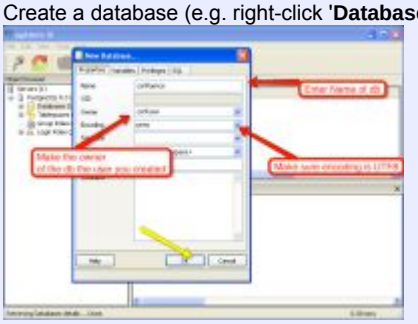
1. Start the '**pgAdmin III**' administration tool on your machine. The pgAdmin III administration console will display. The database user and database that will be used by Confluence are created via the 'pgAdmin III' tool.
2. Connect to the PostgreSQL server (e.g. double-click on the server name in the object browser). Enter a 'postgres' password when prompted.



3. Create a new user, i.e. login role (e.g. right-click click '**Login Roles**' in the object browser and select '**New Login Role...**'):



   - Enter a name and password for the new user.
   - Do not select any role privileges.
4. Create a database (e.g. right-click '**Databases**' and select '**New Database...**'):



   - Enter a name for the new database.
   - Set the owner of the database to the user you created in the previous step.
   - Select '**UTF8**' for '**Encoding**'.

Creating a User and Database via UNIX command-line
If you are on UNIX and do not have the above pgAdmin III administration tool, you can use the command line interface instead. Assuming that you are using the default installation directory of `/opt/PostgreSQL/8.3/bin/`, enter the following commands:

```
sudo -s -H -u postgres
# Create the Confluence user:
/opt/PostgreSQL/8.3/bin/createuser -S -d -r -P -E confuser
# Create the Confluence database:
/opt/PostgreSQL/8.3/bin/createdb -O confuser confluence
exit
```

3. Configure Confluence to use the PostgreSQL Database

Once you have installed and set up PostgreSQL, you will need to configure Confluence to use the PostgreSQL database.

To configure Confluence to use PostgreSQL,

1. Install Confluence, if you haven't done so already. Ensure that you download Confluence Standalone, not the evaluation installer.
2. Ensure that Confluence is stopped (for example, by ensuring that the application server or service which is running Confluence has been stopped or terminated).
3. Install the database drivers, if needed:

   SQL Server JDBC Drivers bundled with Confluence
   The JDBC drivers for this database are bundled with Confluence. You do not have to download or install any JDBC drivers to use this database with Confluence, **if you are using a direct JDBC connection**. If you are connecting via a datasource, you will still need to download and install the drivers manually. See Database JDBC drivers for more information on the bundled JDBC drivers.

   - If you are configuring a datasource to connect to your PostgreSQL database, you will need to place the jar file in **`<confluence install>/WEB-INF/lib`** (for Confluence 2.10 onwards) or **`<confluence install>/common/lib`** (for earlier versions). Information and links to the appropriate database drivers are available on Database JDBC drivers.

     Windows renames .jar extensions to .zip! Just rename it back to .jar. You'll have to set your folder options to view hidden file extensions if you can't rename it without changing the file type (Tools >> Folder Options >> View >> Uncheck "Hide Extensions for known file types."
4. Start Confluence and after entering your license code on the **'Confluence Setup Wizard'** page, click **'Custom Installation'**. The **'Choose a Database Configuration'** page will display.
5. Select '**PostgreSQL**' and click **'External Database'**. The '**Configure Database**' page will display.
6. Choose your desired database connection method (please note that if you choose to connect via datasource, you will need to install the appropriate database drivers as described in the previous step).
7. Enter your PostgreSQL database setup details (as defined in the previous step above):



If the server that is hosting the PostgreSQL database is not the same server as Confluence, then please ensure that the confluence server can contact the database server and also refer to the PostgreSQL documentation on how to set up pg_hba.conf If the pg_hba.conf file is not set properly, remote communication to the PostgresSQL server will fail.

Running SQL Queries
For ongoing maintenance of your server, you can continue to use PGAdmin as your SQL browser.

Troubleshooting

- Known Issues for PostgreSQL contains common issues encountered when setting up your PostgreSQL database to work with Confluence.
- If you are unable to connect to the database from Confluence and they are on different machines, most likely you have a firewall in between the two machines or your pg_hba.conf file is misconfigured. Verify that your firewall is set to allow connections through 5432 or double check your hba configuration
- If Confluence is complaining that it's missing a class file, you might have forgotten to place the jdbc driver in the WEB-INF/lib folder or possibly have placed it in the wrong folder.
- If none of the above describes your issue, please create a support ticket at http://support.atlassian.com and be sure to include your logs (found in confluence-install/logs and confluence-data/logs).

Database Setup For MySQL

This page provides instructions for installing Confluence and the open-source MySQL database on Microsoft Windows, as well as how to set up and configure MySQL to work with Confluence. Additional instructions are also provided for migrating across any existing Confluence database content. Use this guide in conjunction with the more general Database Setup Guide for Any Database. These instructions add some important reference notes specific to MySQL.

We recommend that you use an external database with Confluence for production purposes, as the embedded database does not offer full transactional integrity in the event of sudden power loss. Confluence Standalone with MySQL is well-tested and easily configured for production purposes.

Please note the following points:

- Throughout the instructions below, the `Confluence Installation Directory` refers to the directory where you extracted the Confluence zipped installer.
- The **MySQL Database Setup** procedure below will make modifications to your default MySQL Server settings. These modifications result in:
    1. The default collation (localisation) option being changed to `utf8_bin` (that is, case-sensitive binary UTF8), such that by default, all new database tables will be created with this type of case-sensitive collation.
    2. The default MySQL database storage engine being changed to 'InnoDB'.

On this page:

- 1. Install Confluence
- 2. Install MySQL Server
- 3. Set up your MySQL Database and User
- (Optional) 4. Back Up Confluence Data
- 5. Set Up your Database Connection
- Troubleshooting
- Related Documents

1. Install Confluence

Install Confluence if you have not done so already. Ensure that you download Confluence Standalone, not the evaluation installer.

2. Install MySQL Server

To install MySQL Server,

1. If you do not have an operational MySQL database server instance available, install 'MySQL Community Server' database server (version 5.0).
   The installation package can be downloaded from the MySQL download page or from the version 5.0 download page. Instructions for installing the MySQL 5.0 database server on Windows can found on the 'Installing MySQL on Windows' page of the MySQL web site.
2. Run the 'MySQL Server Instance Config Wizard':
   *If you intend to connect Confluence to an existing, operational MySQL database server instance, we strongly recommend that you reconfigure this database server instance by running through the Config Wizard and initially choosing the **Reconfigure Instance** option.*
   a. At the start of the Config Wizard (or after having chosen **Reconfigure Instance**), choose **Detailed Configuration**.
   b. 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 relevant MySQL documentation for further information.
   c. Choose **'Transactional Database Only'** for the database usage step. This ensures that your MySQL database will use InnoDB as its default storage engine.
      *It is **highly recommended** that you only use the InnoDB storage engine with Confluence. Avoid using the MyISAM storage engine as this can lead to data corruption.*
   d. Set the InnoDB Tablespace settings to your requirements.(The default settings are acceptable).
   e. Set the approximate number of concurrent connections permitted that best suits your Confluence usage requirements. You can use one of the presets if desired or enter a number manually. Refer to the relevant MySQL documentation for further information.
   f. 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.
   g. For the MySQL server instance's default character set option, choose **'Best Support For Multilingualism'** (i.e. UTF-8).
   h. 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 **'Install As Windows Service'** (and **Launch the MySQL Server automatically**) options. Refer to the relevant MySQL documentation for further information.
      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.
   i. Finally, select the **'Modify Security Settings'** option to enter and set your MySQL Server (root) access password.
3. Install the database drivers, if needed:

   SQL Server JDBC Drivers bundled with Confluence
   The JDBC drivers for this database are bundled with Confluence. You do not have to download or install any JDBC drivers to use this database with Confluence, **if you are using a direct JDBC connection**. If you are connecting via a datasource, you will still need to download and install the drivers manually. See Database JDBC drivers for more information on the bundled JDBC drivers.

   - If you are configuring a datasource to connect to your MySQL database, you will need to place the jar file (i.e. `mysql-connector-java-5.x.y-bin.jar` where `x.y` depends on the driver version) in **`<confluence install>/WEB-INF/lib`** (for Confluence 2.10 onwards) or **`<confluence install>/common/lib`** (for earlier versions). Information and links to the appropriate database drivers are available on Database JDBC drivers.

3. Set up your MySQL Database and User

This procedure uses the MySQL Administrator application component of the MySQL GUI Tools package. If MySQL GUI Tools is not already installed or included in your MySQL installation, you will need to download and install this package before proceeding.

To create the database and user privileges,

1. Start the 'MySQL Administrator' application and enter your login credentials when prompted to connect to the MySQL database server.
   *If you ran through the 'MySQL Server Instance Config Wizard' in step 3 of the **Installation and Configuration** (above) and did not change your username, you should be able to connect to the MySQL database server using the default Username of 'root', Server Host of 'localhost' (assuming you are running MySQL Administrator on the same machine as the MySQL Server) and Password which you set in the Config Wizard.*
   If, however, you are running MySQL for the first time and you have not yet run through the Config Wizard, or you do not recall having set your login credentials, you should be able to click OK without a Username and Password.
2. Go to **Service Control** in the left pane and ensure that the MySQL Service is stopped. If not, click '**Stop Service**' to stop the MySQL Service.
3. Go to **Startup Variables** in the left pane and click the '**Advanced**' tab.
4. Ensure that the '**Default Collation**' check box is selected, enter '**utf8_bin**' for the 'Data Collation' value and click '**Apply changes**'.
   *If you did not run through the 'MySQL Server Instance Config Wizard' in step 3 of the **Installation and Configuration** (above), click the '**General Parameters**' tab and ensure that the **Default storage** check box has been selected with its value set to **InnoDB**. Also, click the '**InnoDB Parameters**' tab and ensure that the **Activate InnoDB** check box is selected. If not, implement these settings and click '**Apply changes**' before proceeding.*
5. Go back to **Service Control** in the left pane and click '**Start Service**' to re-start the MySQL Service (with the updated default collation settings).
6. Go to **Catalogs** in the left pane, then right-click in the schema list at the bottom of this pane and select '**Create New Schema**' from the pop-up menu.
7. Call the new Schema name 'confluence' and click '**OK**'.
8. Go to **User Administration** in the left pane and then click '**Add new user**' at the bottom of the window.
9. Enter the MySQL username 'confluenceuser'.
   In fact, you can enter any legitimate MySQL username. However, the remaining steps and procedures on this page assume that you have entered the MySQL username of 'confluenceuser' at this point. If you choose a different MySQL username, substitute 'confluenceuser' in the remaining steps and procedures with your chosen MySQL username.
10. Enter and re-confirm a password for this user and click '**Apply changes**'.
11. Right-click on the username 'confluenceuser' in the left pane and choose '**Add host from which the user can connect**' from the pop-up menu.
    *If prompted to enable the global '**Show hosts in user list**' option, choose '**Yes**'.*
12. Enter the host name 'localhost' and click '**OK**'.
13. Click on '**localhost**' within 'confluenceuser' in the left pane. (You might need to click on the name 'confluenceuser' first).
14. Click the '**Schema Privileges**' tab and then click '**confluence**' in the 'Schemata' list.
15. Make sure that 'confluenceuser@localhost' has all privileges assigned to them by clicking the '**<<**' button to move the privileges from the 'Available' to the 'Assigned' list and click '**Apply changes**'.
    To support international languages in Confluence, you should verify the newly created database is using UTF-8 encoding and re-examine the JDBC URL settings (configured in Stage 3).

(Optional) 4. Back Up Confluence Data

This stage is only required if you have existing Confluence content you wish to transfer.
To back up your Confluence data,

1. Manually create an XML backup of Confluence under Administration -> Backup & Restore. If you have less than 100MB of attachments, check 'Backup attachments' when creating the backup. If you have over 100MB of attachments, you should not check the 'Backup attachments' and instead you should manually copy the /attachments directory in your Confluence home to another location. This attachments directory can then be copied into the new home directory as described later.
2. Download the backup file to a backups folder.

5. Set Up your Database Connection

To set up your Confluence MySQL database connection or to switch to using this external database,

1. Ensure that Confluence is stopped (for example, by ensuring that the application server or service which is running Confluence has been stopped or terminated).
2. Edit `Confluence Installation Directory_/confluence/WEB-INF/classes/confluence-init.properties` and change the `confluence.home` property to point to a new directory. e.g. if you had

   ```
   1.confluence.home=c:/confluencedata
   ```

   You could change it to:

   ```
   1.confluence.home=c:/confluencedata_mysql
   ```

   This is your new `Confluence Home Directory`. (The name doesn't have to end in _mysql – that's just an example)

3. Start Confluence and set up the new configuration.
   a. Enter your license key and click the '**Custom Installation**' button.
   b. Under the 'External Database' heading, select '**MySQL**' from the dropdown list and click '**External Database**'.
   c. On the next page, click '**Direct JDBC**'.
   d. Enter `confluenceuser` in the User Name field, and the password you chose earlier in the Password field
   e. Click the Next button. If you get the error message `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)'` verify that you have properly given the confluenceuser user all the right permissions when connecting from localhost.
   f. *(optional)* If you previously backed up your Confluence data, you can choose to restore it at the 'Load Content' page. Choose 'Restore From Backup', browse for the backup you created and restore it. Otherwise choose either the example or empty site as you wish.

Troubleshooting

- Known Issues for MySQL contains common issues encountered when setting up your MySQL database to work with Confluence.
- If Confluence is complaining that it's missing a class file, you might have forgotten to place the jdbc driver in the `WEB-INF/lib` folder or possibly have placed it in the wrong folder.
- If none of the above describes your issue, please create a support ticket at http://support.atlassian.com and be sure to include your logs (found in confluence-install/logs and confluence-data/logs).

Related Documents

- Configuring Database Character Encoding
- Known Issues for MySQL

Database Setup for Oracle

This guide covers deploying Confluence standalone or WAR distribution with an Oracle database.



This database can only be set up by an Oracle database administrator (DBA)
If you are not a DBA, you should not attempt to set up this database.

Oracle has a history of being extremely difficult to set up. If you do not have access to an experienced Oracle DBA in your organisation, you are recommended to select any free, scalable and easy-to-install alternative rather than proceeding with Oracle. Users evaluating Confluence are recommended to start with an alternative database and only consider migrating to Oracle after approval from their DBA. Atlassian's technical support for Oracle setup difficulties will also reflect the high minimum skill requirements for attempting an Oracle setup.

Database Setup Information

This setup guide must be used in conjunction with the list of Known Issues For Oracle. Please review that page before continuing.

Schema Requirements

Confluence must be deployed to a schema in it's own, separate Oracle instance. This avoids a Hibernate bug triggered by the existence of duplicate tablenames in other schemas, outlined here.

Database Incompatibilities

Orcale 9i and later are supported, please upgrade any older versions prior to installing Confluence.



Database driver update may be required
For all versions of Oracle, you must upgrade to the latest 10g drivers. Check the latest compatibility matrix before proceeding to the JDBC download page. We highly recommend to use the **thin** drivers.

Please check that your version of Oracle does not have any known issues:

Oracle Version	Oracle Driver	Issue	Solution
Any	Pre 10g	Driver incompatibilities	Upgrade to latest 10g drivers if compatible
Pre 9	Any	Earlier DBs not compatible	Upgrade DB to 9.x or later
9.0.1.0	Any	DB not compatible	Upgrade DB to 9.0.2.6 or later
10g	10.1.0.3.0	Drivers not compatible with Tomcat 5.5	Upgrade drivers to 10.1.0.4.0 or later

You might be also interested in a relevant JIRA documentation to check the compatibility of your Oracle server and driver.

Deploying Against Oracle

Complete the instructions for installing Confluence standalone, then **return to this document instead of proceeding** to the Confluence Setup Guide.

Database Preparation

Tailor these instructions to your particular database version:

1. Perform any necessary database or driver upgrades. Download the latest compatible database drivers - click here for Oracle JDBC driver FAQ.
2. Create a separate instance of Oracle and apply any configuration tasks.
3. Create a Confluence user configuration and grant access permissions.

Adding a Datasource to Tomcat

1. Open `<INSTALL>/conf/server.xml` for editing.
2. Locate the section `Host -> Context`

```
1.<Host name="localhost" debug="0" appBase="webapps" unpackWARs="true" autoDeploy="false">
2.          <Context path="" docBase="../confluence" debug="0" reloadable="true">
3.                      <!-- Logger is deprecated in Tomcat 5.5. Logging configuration for
Confluence is specified in confluence/WEB-INF/classes/log4j.properties -->
4.                      <Manager pathname="" />
5.          </Context>
6.</Host>
```

3. Paste in the `Resource` section provided, prior to `Manager` as shown

```
01.<Host name="localhost" debug="0" appBase="webapps" unpackWARs="true" autoDeploy="false">
02.          <Context path="" docBase="../confluence" debug="0" reloadable="true">
03.                      <!-- Logger is deprecated in Tomcat 5.5. Logging configuration for
Confluence is specified in confluence/WEB-INF/classes/log4j.properties -->
04.
05.                      <Resource
06.                      name="jdbc/confluence"
07.                      auth="Container"
08.                      type="javax.sql.DataSource"
09.                      driverClassName="oracle.jdbc.driver.OracleDriver"
10.                      url="jdbc:oracle:thin:@hostname:port:sid"
11.                      username="<username>"
12.                      password="<password>"
13.                      connectionProperties="SetBigStringTryClob=true"
14.                      maxActive="25"
15.                      maxIdle="5"
16.                      maxWait="10000"
17.                      />
18.
19.                      <Manager pathname="" />
20.          </Context>
21.</Host>
```

4. Change the `username` and `password` to match the Oracle login.
5. Change `url` to match `hostname`, `port` and `sid` of the Oracle server. `sid` stands for the Schema ID. For example:

```
1.jdbc:oracle:thin:@example.atlassian.com:1521:confluencedb
```

For connecting to an Oracle RAC cluster, you'll need to edit the connection string directly in `confluence.cfg.xml`, using Oracle's connection syntax like this:

```
01.jdbc:oracle:thin:@(DESCRIPTION=
02.                                                  (SDU=32768)
03.                                                  (enable=broken)
04.                                                  (LOAD_BALANCE=yes)
05.                                                  (FAILOVER=yes)
06.                                                  (ADDRESS=
07.                                                       (PROTOCOL=TCP)
08.                                                       (HOST=dbserver1.example.com)
09.                                                       (PORT=1525))
10.                                                  (ADDRESS=
11.                                                       (PROTOCOL=TCP)
12.                                                       (HOST=dbserver2.example.com)
13.                                                       (PORT=1525))
14.                                                  (CONNECT_DATA=
15.                                                       (SERVICE_NAME=CONFDB)))
```

This example has been broken up over multiple lines for clarity, but it should be compacted into a single line when actually added to `confluence.cfg.xml` on your server.

6. If required, choose different maxActive and maxIdle values. These set how many total database connections will be allowed at one time, and how many will be kept open even when there is no database activity.

Configure Confluence Datasource Access

Confluence must be configured to use this datasource:

1. Edit the file `<INSTALL>/confluence/WEB-INF/web.xml`
2. Go to the end of the file and just before `</web-app>`, insert the following:

```
1.<resource-ref>
2.<description>Connection Pool</description>
3.<res-ref-name>jdbc/confluence</res-ref-name>
4.<res-type>javax.sql.DataSource</res-type>
5.<res-auth>Container</res-auth>
6.</resource-ref>
```

3. Locate the Oracle JDBC database drivers for JDK 1.4, which comes bundled with your database. We recommend using the thin drivers only. Copy `ojdbc14.jar` to the `<INSTALL>/common/lib` directory. This directory path is potentially `<INSTALL>/lib` if Confluence is running off Apache Tomcat version 6 or above.

Confluence Setup Wizard

Now Confluence is ready to attempt to connect to Oracle:

1. Startup Confluence using `<INSTALL>/bin/startup.bat` or `<INSTALL>/bin/startup.sh`
2. Insert your licence and select **External Database**.
3. Select **Datasource Connection** using your Oracle version.
4. Enter `java:comp/env/jdbc/confluence` for the name of the datasource.

Confluence should now deploy using the Oracle database specified. Please read this comment on Oracle database optimisation.

Generic Oracle Configuration Tips

Websphere and Oracle

If you are planning to run Confluence on a Websphere application server and Oracle database, you should read the information on Known Issues for Websphere.

24-hour time format with Oracle 8i

We have received a report from a user that when an Oracle 8i database is configured to use 24-hour time as the default format, an exception like this may occur:

```
01.005-12-06 13:23:20 Loading root WebApplicationContext
02.
03.2005-12-06 13:24:34 StandardContext[]: Exception sending context initialized event to listener
instance
04.of class com.atlassian.confluence.util.ConfluenceContextLoaderListener
05.
06.org.springframework.beans.factory.BeanCreationException: Error creating bean with name
'userAccessor' defined in class path resource [applicationContext.xml]:
07.Can't resolve reference to bean 'userAccessorTarget' while setting property 'target';
08.nested exception is org.springframework.beans.factory.BeanCreationException: Error creating
bean with name 'userAccessorTarget' defined in class path
09.resource [applicationContext.xml]: Can't resolve reference to bean 'spacePermissionManager'
while setting property 'spacePermissionManager';
10.nested exception is org.springframework.beans.factory.BeanCreationException: Error creating
bean with name 'spacePermissionManager' defined in class path resource [securityContext.xml]:
11.Can't resolve reference to bean 'spacePermissionManagerTarget' while setting property 'target';
12.nested exception is org.springframework.beans.factory.BeanCreationException:
13.Error creating bean with name 'spacePermissionManagerTarget' defined in class path resource
[securityContext.xml]: Initialization of bean failed;
14.nested exception is org.springframework.jdbc.UncategorizedSQLException: (Hibernate operation):
encountered SQLException [Cannot create PoolableConnectionFactory];
15.nested exception is org.apache.commons.dbcp.SQLNestedException: Cannot create
PoolableConnectionFactory
16.
17....
18.
19.org.apache.commons.dbcp.SQLNestedException: Cannot create PoolableConnectionFactory, cause:
20.
21.
22.java.sql.SQLException: ORA-00604: error occurred at recursive SQL level 1
23.
24.
25.ORA-12705: invalid or unknown NLS parameter value specified
```

One symptom of this problem is that Confluence may refuse to start after midday.

The workaround is to go to 'General Configuration' and set the default time format to "HH:mm".

RELATED TOPICS

Known Issues For Oracle

Database Setup for SQL Server

Use this guide in conjunction with the more general Database Setup Guide for Any Database. These instructions add some reference notes specific to SQL Server.

1. Review the known issues for SQL Server.
2. Identify which character encoding to use. To do this, check the encoding currently used by your application server and Confluence. All three must use compatible encoding. For example, the default SQL Server encoding of USC-2 is compatible with UTF-8.
3. Create a new database (as an SQL administrator). If you set your application server and Confluence to use an encoding incompatible with USC-2, specify that character encoding for the database.
4. Create a new SQL user account for Confluence (as an SQL administrator). Provide full create, read and write permissions for the table. Please note, Confluence must be able to create its own schema.
5. Install the database drivers, if needed:

   SQL Server JDBC Drivers bundled with Confluence
   The JDBC drivers for this database are bundled with Confluence. You do not have to download or install any JDBC drivers to use this database with Confluence, **if you are using a direct JDBC connection**. If you are connecting via a datasource, you will still need to download and install the drivers manually. See Database JDBC drivers for more information on the bundled JDBC drivers.

   - If you are configuring a datasource to connect to your MS SQL server database, you will need to place the jar file in **`<confluence install>/WEB-INF/lib`** (for Confluence 2.10 onwards) or **`<confluence install>/common/lib`** (for earlier versions). Information and links to the appropriate database drivers are available on Database JDBC drivers. You may also find this page helpful: http://jtds.sourceforge.net/faq.html
6. Start Confluence and visit the home URL (eg http://localhost:8080) to start the Confluence Setup Wizard and select a Custom Install, insert the relevant connection information.
   - When prompted for a **driver class name** in the database setup step enter:

     ```
     1.net.sourceforge.jtds.jdbc.Driver
     ```

   - When prompted for the **jdbc url**, the format to use is:

     ```
     1.jdbc:jtds:sqlserver://<server>:<port>/<database>
     ```

Database JDBC drivers

The JDBC drivers for all databases currently supported for Confluence are linked below. You will need to make the driver available to your application server, as described in the appropriate setup guide.

Please note, we bundle a number of JDBC drivers with Confluence, as shown below. You do not have to download or install the drivers for the relevant databases, **if you are using a direct JDBC connection**. If you are connecting via a datasource, you will still need to download and install the drivers manually.

JDBC Driver Download Links

Database	JDBC Drivers Bundled with Confluence?	JDBC Drivers
PostgreSQL	✓	8.4-701.jdbc3 (note, the JDBC 3 driver will work under the 1.6 JVM)
MySQL	✓	5.1.11
Microsoft SQL Server	✓	jTDS 1.2.2
Oracle	✗	JDBC driver downloads (see Database Setup for Oracle for required JDBC driver versions)
DB2	✗	JDBC drivers should be included with DB2, otherwise they can be downloaded from the IBM website

Creating Database Schema Manually

Database Schema Creation

This document provides information on how to find the SQL for database table creation.

Often, DBAs will require that table creation be done manually. To find the required SQL statements, you can use the script generated from our evaluation version's HSQLDB database.

- Install Confluence with no external database.
- Shut down Confluence.
- Find the file located in `<confluence-home>/database/confluencedb.script`.

The `confluence-home` directory is not your installation directory, its the directory specified in `confluence-init.properties`. For more information, see Confluence Home Directory and Confluence Installation Directory.

To see some example SQL statements, click 'Tools' and select 'Attachments' on this wiki page.

To get Confluence to run against this established schema, configure the database with the normal procedure, then modify the hibernate.connection.url, username and password from `confluence-home/confluence.cfg.xml` or `server.xml`, depending on whether it's a direct jdbc or resoruce connection. This is described in Migrating Confluence Between Servers.

RELATED TOPICS

Database Configuration

Known Issues For Supported Databases

Supported Databases

Please refer to the Supported Platforms topic.

Troubleshooting

- Known Issues for PostgreSQL
- Known Issues For SQL Server
- Known Issues for MySQL
- Configuring Database Character Encoding
- Known Issues for DB2
- Known Issues for Sybase Database

Known Issues For Oracle

Configuring Database Character Encoding

On this page:

- JDBC connection settings
  - MySQL
- Creating a UTF-8 database
  - MySQL
  - PostgreSQL
  - For PostgreSQL running under Windows
  - For PostgreSQL running under Linux
- Updating existing database to UTF-8
  - MySQL database with existing data
- Testing database encoding

The database used with Confluence should be configured to use the same character encoding as Confluence. **The recommended encoding is Unicode UTF-8.**

There are two places where character encoding may need to be configured:

- when creating the database
- when connecting to the database (JDBC connection URL or properties).

The configuration details for each type of database are different. Some examples are below.

JDBC connection settings

MySQL

Append "useUnicode=true to your JDBC URL:

```
jdbc:mysql://hostname:port/database?autoReconnect=true&useUnicode=true&characterEncoding=utf8
```

If you are modifying `confluence.cfg.xml` directly rather than via the Confluence Installation GUI, you'll need to escape ou syntax when the XML is parsed. An effective URL could be similar to:

```
<property
name="hibernate.connection.url">jdbc:mysql://hostname:port/database?autoReconnect=true&amp;u
```

Creating a UTF-8 database

MySQL

1. Create a UTF-8 database with binary UTF-8 collation.
   Binary UTF-8 provides case-sensitive collation.

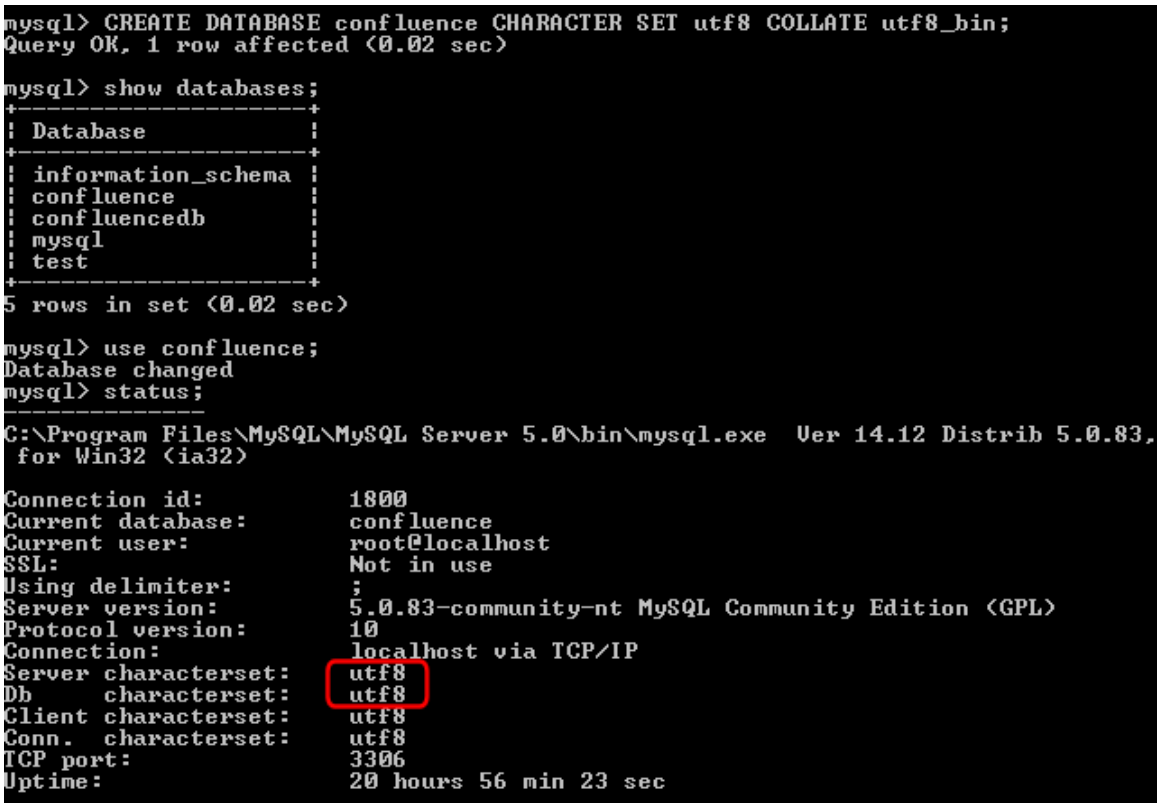
   ```
   1.CREATE DATABASE confluence CHARACTER SET utf8 COLLATE utf8_bin;
   ```

2. You will also need to set the `Server Characterset` to `utf8`. This can be done by adding the following in `my.ini` for Windows or `my.cnf` for other OS. It has to be declared in the Server section, which is the section after `[mysqld]`:

   ```
   1.[mysqld]
   2.default-character-set=utf8
   ```

3. Use the **`status`** command to verify database character encoding information.

   Screenshot: Using the Status Command to Verify Database Character Encoding



4. In some cases, the individual tables collation and character encoding may differ from the one that the database as a whole has been configured to use. Please use the command below to ensure all tables within your Confluence database are correctly configured to use UTF-8 character encoding and binary UTF-8 collation:

```
1.use confluence;
2.show table status;
```

Check for the value listed under the **Collation** column, to ensure it has been set to `utf8_bin` (that is, case-sensitive) collation for all tables.
If not, then this can be changed by the following command, executed for each table in the Confluence database:

```
1.ALTER TABLE tablename CONVERT TO CHARACTER SET utf8 COLLATE utf8_bin;
```

Please substitute the <tablename> above, with each table within the confluence database.

Relevant MySQL manual for more detailed explanation:

- Specifying Character Sets and Collations documentation.
- Connection Character Sets and Collations.
- SHOW TABLE STATUS Syntax.
- ALTER TABLE Syntax.

PostgreSQL

```
1.CREATE DATABASE confluence WITH ENCODING 'UNICODE';
```

Or from the command-line:

```
1.$ createdb -E UNICODE confluence
```

For more information see the PostgreSQL documentation.

For PostgreSQL running under Windows

Please note that international characters sets are only fully supported and functional when using PostgreSQL 8.1 and above under Microsoft Windows.

For PostgreSQL running under Linux

Please make sure you check the following to ensure proper handling of international characters in your database

When PostgreSQL creates an initial database cluster, it sets certain important configuration options based on the host enviroment. The command responsible for creating the PostgreSQL environment `initdb` will check environment variables such as `LC_CTYPE` and

`LC_COLLATE` (or the more general `LC_ALL`) for settings to use as database defaults related to international string handling. As such it is important to make sure that your PostgreSQL environment is configured correctly **before you install Confluence**.

To do this, connect to your PostgreSQL instance using `pgsql` and issue the following command:

```
SHOW LC_CTYPE;
```

If LC_CTYPE is set to either "C" or "POSIX" then certain string functions such as converting to and from upper and lower case will not work correctly with international characters. Correct settings for this value take the form `<LOCALE>.<ENCODING>` (`en_AU.UTF8` for example).

If your LC_CTYPE is incorrect please check the PostgreSQL documentation for information on configuring database localisation. It is not easy to change these settings with a database that already contains data.

Updating existing database to UTF-8

MySQL database with existing data



Before proceeding with the following changes, please **backup your database**.

This example shows how to change your database from latin1 to utf8.

1. Dump the database to a text file using mysqldump tool from the command-line :
   `mysqldump -p --default_character-set=latin1 -u <`*`username`*`> --skip-set-charset confluence > confluence_database.sql`
2. copy `confluence_database.sql` to `confluence_utf8.sql`
3. Open `confluence_utf8.sql` in a text editor and change all character sets from 'latin1' to 'utf8'
4. Encode all the latin1 characters as UTF-8:
   `recode latin1..utf8 confluence_utf8.sql` (the recode utility is described at http://directory.fsf.org/recode.html; it can actually be downloaded from http://recode.progiciels-bpi.ca/, and is available for Ubuntu via `apt-get`)

In MySQL:

1. `DROP DATABASE confluence;`
2. `CREATE DATABASE confluence CHARACTER SET utf8 COLLATE utf8_bin;`

Finally, reimport the UTF-8 text file:

1. `mysql -u <`*`username`*`> -p --default-character-set=utf8 --max_allowed_packet=64M confluence < /home/confluence/confluence_utf8.sql`

To support large imports, the parameter '--max_allowed_packet=64M' used above sets the maximum size of an SQL statement to be very large. In some circumstances, you may need to increase it further, especially if attachments are stored in the database.

Testing database encoding

See Troubleshooting Character Encodings for a number of tests you can run to ensure your database encoding is correct.

RELATED TOPICS:

Character encodings in Confluence
Known Issues for MySQL

Known Issues for DB2

On this page:

- DB2 Does Not Support Unicode Character By Default
- DB2 Does Not Support Transaction Logging Of BLOBs Larger Than 1GB
- Configuring Database Character Encoding
- DB2 Dialect class change (if upgrading from Confluence version 2.2 or older)
- Encoding Test Problems
- Transaction Isolation Level of Read Uncommitted
- Incompatible Data Types on z/OS 8
- SQLCODE: -443, SQLSTATE: 38553, SQLERRMC: SYSIBM.SQLTABLES;TABLES;SYSIBM:CLI:-805

DB2 Does Not Support Unicode Character By Default

To support the Unicode character set, create your database like this:
For 8.2.2 or later:

```
1.db2 create database <name> using codeset utf-8 territory us pagesize 8 k
```

Before 8.2.2

```
1.db2 create database <name> using codeset utf-8 territory us
2.db2 drop tablespace userspace1
3.db2 create bufferpool bufpool8k <number of pages> pagesize 8 k
4.db2 create tablespace userspace1 pagesize 8 k managed by database using \(file \'<location>\'
5.                          <number of pages>\) bufferpool bufpool8k
```

The territory may not be important when using UTF-8, but it must be included.

DB2 Does Not Support Transaction Logging Of BLOBs Larger Than 1GB

http://publib.boulder.ibm.com/infocenter/db2luw/v8/index.jsp?topic=/com.ibm.db2.udb.doc/ad/c0006327.htm

> *"With their potentially large size, LOBs can slow down the performance of your database system significantly when moved into or out of a database. Even though DB2 does not allow logging of a LOB value greater than 1 GB, LOB values with sizes approaching 1 GB can quickly push the database log to near capacity. An error, SQLCODE -355 (SQLSTATE 42993), results from attempting to log a LOB greater than 1 GB in size. The lob-options-clause in the CREATE TABLE and ALTER TABLE statements allows users to turn off logging for a particular LOB column. Although setting the option to NOT LOGGED will improve performance, changes to the LOB values after the most recent backup are lost during roll-forward recovery."*



Although one could set the length for the CustomClobType in BodyContent.hbm.xml field length to larger sizes, say 2^31 to allow storing of BLOBs up to the size of 2 Gb, DB2 maximum length can only be set to 2^30, 1GB.

```
1.<property name="body" type="com.atlassian.confluence.core.persistence.hibernate.CustomClobType"
update="true" insert="true"
2.                                          column="BODY" length="1073741824" />
```

Configuring Database Character Encoding

Refer to Configuring Database Character Encoding.

DB2 Dialect class change (if upgrading from Confluence version 2.2 or older)

If you are using DB2, and are upgrading Confluence from version 2.2 or older, the 'hibernate.dialect' property in the `confluence.cfg.xml` file needs to be changed to 'com.atlassian.hibernate.dialect.DB2Dialect' like so:

```
<property name="hibernate.dialect">com.atlassian.hibernate.dialect.DB2Dialect</property>
```

Encoding Test Problems

At the time of writing this, there is an open bug with the encoding test when using DB2. Details are here: CONF-8588.

Transaction Isolation Level of Read Uncommitted

Due to CONF-9323, DB2 requires a transaction isolation level of "read uncommitted" to avoid occasional deadlocks. The more concurrent usage a Confluence instance experiences, the higher the frequency of deadlocks with a transaction isolation level higher than "read uncommitted".

Set the following property in the `confluence.cfg.xml` file.

```
1.<property name="hibernate.connection.isolation">1</property>
```

Incompatible Data Types on z/OS 8

There are problems when running DB2 on z/OS 8 and 9. This platform is not officially supported but there is an unsupported workaround here.

SQLCODE: -443, SQLSTATE: 38553, SQLERRMC: SYSIBM.SQLTABLES;TABLES;SYSIBM:CLI:-805

```
Caused by: com.ibm.db2.jcc.b.SqlException: DB2 SQL error: SQLCODE: -443, SQLSTATE: 38553,
SQLERRMC: SYSIBM.SQLTABLES;TABLES;SYSIBM:CLI:-805
 at com.ibm.db2.jcc.b.hh.b(hh.java:1369)
 at com.ibm.db2.jcc.b.hh.c(hh.java:1356)
 at com.ibm.db2.jcc.c.db.k(db.java:352)
 at com.ibm.db2.jcc.c.db.e(db.java:96)
 at com.ibm.db2.jcc.c.t.e(t.java:83)
 at com.ibm.db2.jcc.c.sb.h(sb.java:167)
 at com.ibm.db2.jcc.b.hh.q(hh.java:1329)
 at com.ibm.db2.jcc.b.ih.d(ih.java:2529)
 at com.ibm.db2.jcc.b.jh.V(jh.java:146)
 at com.ibm.db2.jcc.b.fb.a(fb.java:7374)
 at com.ibm.db2.jcc.b.fb.a(fb.java:5997)
 at com.ibm.db2.jcc.b.fb.getTables(fb.java:5908)
 at
com.mchange.v2.c3p0.impl.NewProxyDatabaseMetaData.getTables(NewProxyDatabaseMetaData.java:2828)
 at net.sf.hibernate.tool.hbm2ddl.DatabaseMetadata.getTableMetadata(DatabaseMetadata.java:54)
 ... 23 more
```

You might encounter this problem after an upgrade to UDB Version 8.1 FixPak 10 (also known as Version 8.2 FixPak 3) when invoking a DB2 Call Level Interface (CLI) catalog function. The solution to this is to perform a rebind of the `db2schema.bnd` file locally against each database. Refer to the IBM Technote FAQ for more instruction.

RELATED TOPICS

Interpreting DB2 error codes

Interpreting DB2 error codes

When DB2 dies, it'll dump its error codes out in the SQL Exception:

```
DB2 SQL error: SQLCODE: -803, SQLSTATE: 23505
```

You can find the meaning of the error code from the db2 command prompt with "`? sqlxxxx`" where xxxx is the SQLCODE from the error message:

```
[db2inst1@matilda db2inst1]$ db2 ? sql-803

SQL0803N One or more values in the INSERT statement, UPDATE
statement, or foreign key update caused by a DELETE
statement are not valid because the primary key, unique
constraint or unique index identified by
"<index-id>" constrains table "<table-name>" from
having duplicate rows for those columns.

Explanation:

The INSERT or UPDATE object table "<table-name>" is constrained
by one or more UNIQUE indexes to have unique values in certain
columns or groups of columns. Alternatively, a DELETE statement
on a parent table caused the update of a foreign key in a...

(and so on for about three more pages of text...)
```



You may also like to use the SQL Message Finder to find information about iSeries SQL messages. You can search by message ID, SQLCODE, or SQLSTATE value. You can also select an SQLSTATE class code.

There is also an online reference for SQL Messages and Codes which is quite useful.

Known Issues for MySQL

On this page:

- No 'toLower' Capability - Database Case Sensitivity
- Configuring Database Character Encoding
- MySQL Storage Engine
    - Upgrading to Confluence 3.0 or later fails

- MySQL JDBC Drivers
- Access Denied
- Max Allowed Packet Size Exceeded
  - Prior MySQL 4.0, use this syntax instead:
  - From MySQL 4.0, use this syntax
- Duplicate Key Exception During Import
  - Setting the MySQL Collation to be case sensitive
- Database Timeout Issues when creating a Manual Backup
- After a while, database errors are generated and Confluence stops working
- Troubleshooting Tips

This page provides tips on configuring MySQL as Confluence's database.

No 'toLower' Capability - Database Case Sensitivity

Some of the database indices are not available for MySQL. See Creating a Lowercase Page Title Index for more detail on a workaround.

Configuring Database Character Encoding

To prevent problems with character encoding, for consistency, we recommend to use Unicode character encoding UTF-8 among all the entities of your system. See Configuring Database Character Encoding for more details.

When specifying a character encoding as part of your mysql connection url (eg: &characterEncoding=utf8), it is important to ensure that the specified encoding is compatible with the default encoding used by your database. Note: if you do not specify a characterEncoding on the connection url, the connection will default to the server's default character set.

Full details of MySQLs character support is available here: http://dev.mysql.com/doc/mysql/en/charset.html

MySQL Storage Engine

The default storage engine for MySQL is MyISAM. Because this storage engine does not support referential integrity, foreign key constraints or transactions, using it may lead to data corruption. Some known issues caused by using MyISAM include CONF-16070 and CONF-16494. Hence, this storage engine is not recommended for use with Confluence.

For a MySQL command line session, you can set the storage engine by passing the '--default-storage-engine=INNODB' option when starting the session. You can make this the default MySQL Server setting by adding this option to the `my.ini` (or `my.cnf`) configuration file.

For more information, please refer to: http://dev.mysql.com/doc/refman/5.1/en/storage-engines.html.

Upgrading to Confluence 3.0 or later fails

Some customers running Confluence on a MySQL database may find that when they upgrade to Confluence 3.0 or later, their Confluence 3.0 upgrade fails, with the Confluence logs revealing a "Specified key was too long" error. This issue is known to occur when MySQL's MyISAM storage engine and UTF-8 character set is used with Confluence.

If this is the case, please refer to the Upgrade to Confluence 3.0 with MySQL database fails with messages like "specified key was too long" knowledge base article on how to resolve this upgrade issue.

MySQL JDBC Drivers

Ensure that you are using the latest (5.1) MySQL Java Connector. Earlier versions of the MySQL connector have a bug which is triggered by improvements in Confluence 2.2. These earlier connector versions will result in an error being recorded in your logs on upgrade (and will result in unstable operation of Confluence)

```
1.ERROR [hibernate.tool.hbm2ddl.SchemaUpdate] execute could not complete schema update
```

You can download the latest MySQL connector from the MySQL Java Connector 5.1 download page. Please be sure that you remove any older versions of the connector from your application server.

Don't use the debug version of these drivers (the jar file ending in '-g.jar'). This requires extra configuration, see Installing the Driver and Configuring the CLASSPATH

Access Denied

If you get a connection error: **Access denied for user 'confluenceuser'@'localhost.localdomain'** this may be because 127.0.0.1 resolves to 'localhost.localdomain' in your environment. Create a user 'confluenceuser@localhost%' to match any domain starting with localhost.

Max Allowed Packet Size Exceeded

A common problem with MySQL is the max packet size restriction, which can result in an "Object Error" message when installing certain plugins.

If you are using MySQL 4 and prior, you may come across a problem with **`max_allowed_packet`** size.

```
ERROR [sf.hibernate.util.JDBCExceptionReporter] logExceptions Packet for query is too large
(1259485 > 1048576). You can change this value on the server by setting the max_allowed_packet'
variable.
```

To resolve these problems, you need to either specify or increase the value for **`max_allowed_packet`**. Make sure to set the packed size when starting the server, not the client.

Prior MySQL 4.0, use this syntax instead:

```
1.shell> mysqld --set-variable=max_allowed_packet=16M
```



In MySQL 3.23, the largest possible packet is 16MB, due to limits in the client/server protocol. In MySQL 4.0.1 and up, the limit is 1GB.

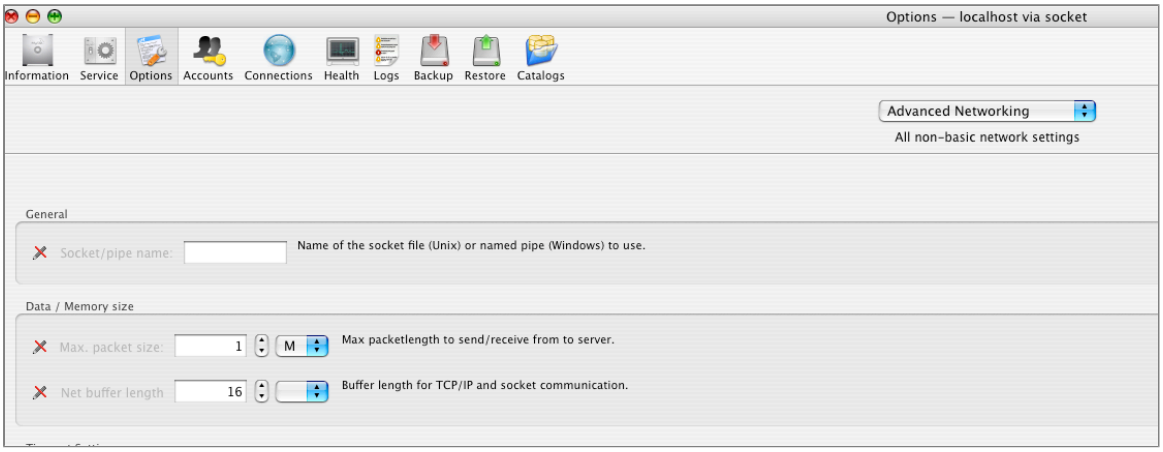
From MySQL 4.0, use this syntax

```
1.shell> mysqld --max_allowed_packet=32M
```

> *For more information, please refer to MySQL manual:*
> *http://dev.mysql.com/doc/refman/5.0/en/packet-too-large.html*

If you use MySQL Administrator, you can set this parameter from the **Startup Variables** section on Windows OS (**Options** on MacOS) > **Advanced Networking** subsection. You can set **`max_allowed_packet`** option to 16 (MB), which in the screenshot below, is currently set at 1 MB.

*Screenshot: Setting the **`max_allowed_packet`** option via the MySQL Administrator*



Duplicate Key Exception During Import

When upgrading to MySQL from another database, such as HSQL, importing the site backup often fails with an error like this:

```
1.Duplicate key or integrity constraint violation message from server:
```

Such errors occur because usually MySQL evaluates unique key constraints and primary key constraints in a case insensitive way. So if you have a space with the key "sp" and another with the key "SP", MySQL will refuse to add the second one.

This problem is avoidable by setting the collation on the database to be case sensitive.

Setting the MySQL Collation to be case sensitive

MySQL uses collations for sorting data and for evaluating uniqueness.

To set the collation to case sensitive when using utf8, use this command:

```
1.CREATE DATABASE confluence CHARACTER SET utf8 COLLATE utf8_bin;
2.
3.Note: The collation must be compatible with the character set.  The name of the database in the
example is confluence.
```

For further information see the MySQL documentation on character sets on collations.

Database Timeout Issues when creating a Manual Backup

A problem that some customers have encountered is their Database connection timing out whilst in the middle of performing a Manual Site

Backup in Confluence.

This issue is indicated in the log files with the following Error Message:

```
com.mysql.jdbc.CommunicationsException: Communications link failure due to underlying exception:
```

According to the MySQL manual:

> *The number of seconds the server waits for activity on a non-interactive connection before closing it. This timeout applies only to TCP/IP and Unix socket file connections, not to connections made via named pipes, or shared memory.*
>
> *On thread startup, the session wait_timeout value is initialized from the global wait_timeout value or from the global interactive_timeout value, depending on the type of client (as defined by the CLIENT_INTERACTIVE connect option to mysql_real_connect()). See also interactive_timeout.*

This problem is resolved by adjusting the **wait_timeout** parameter to a higher value.

If problems persist after making the changes, switch to the Production Backup Strategy.

After a while, database errors are generated and Confluence stops working

MySQL's JDBC drivers usually close a connection that remains idle for a certain amount of time (normally eight hours). Since Confluence uses a connection pool, this means that pooled connections will be terminated if they are not used within a certain time period.

The solution is to append `?autoReconnect=true` to the end of your database's JDBC URL.

- If you are using JNDI data-sources, you will do this in your application server's configuration files.
- If you have configured Confluence to access the database directly, you will need to manually edit the `hibernate.connection.url` property in the `confluence.cfg.xml` file in your `confluence.home` directory. After you have changed the URL in this file, restart Confluence.

Troubleshooting Tips

The following are tips for Troubleshooting MySQL.

- ClientAbortException java.net.SocketException Broken Pipe with Jira User management on MySQL
- Setup Fails Creating MySQL Schema due to Tomcat Incompatibility
- Login after a long inactivity on confluence fails using MySQL
- Characters appear as question marks using MySQL
- Database errors when using MySQL and MyISAM tables
- Error loading plugins or saving content larger than 1Mb
- Upgrade to Confluence 3.0 with MySQL database fails with messages like "specified key was too long"
- Error- Invalid use of group function
- Cannot Create XML Backup due to Corrupt Table
- MySQL Table is Marked as Crashed and Should be Repaired
- Cannot Restore XML Backup due to Data Truncation - MySQL Driver 3.1

Known Issues For Oracle

- Use Oracle with thin Oracle 10g JDBC drivers
- Running Confluence on Websphere with Oracle database
- 24-hour time format with Oracle 8i
- ORA-01430: column being added already exists in table
- Configuring Database Character Encoding
- Knowledge Base Articles

Use Oracle with thin Oracle 10g JDBC drivers

We recommend you to use the following configuration to run Confluence with Oracle:

- Regardless of what version of the Oracle database you are using, you should use the **Oracle 10g** JDBC drivers. (Note: Oracle 10g JDBC drivers will not work with Oracle 8.1.6
  see Oracle FAQ)

- We highly recommend to use the **thin** drivers.

Also see:

Running Confluence on Websphere with Oracle database

If you are planning to run Confluence on a Websphere application server and Oracle database, you should read the information on Known Issues for Websphere.

24-hour time format with Oracle 8i

We have received a report from a user that when an Oracle 8i database is configured to use 24-hour time as the default format, an exception like this may occur:

```
01.005-12-06 13:23:20 Loading root WebApplicationContext
02.
03.
04.2005-12-06 13:24:34 StandardContext[]: Exception sending context initialized event to listener
instance
05.of class com.atlassian.confluence.util.ConfluenceContextLoaderListener
06.
07.org.springframework.beans.factory.BeanCreationException: Error creating bean with name
'userAccessor' defined in class path resource [applicationContext.xml]:
08.Can't resolve reference to bean 'userAccessorTarget' while setting property 'target';
09.nested exception is org.springframework.beans.factory.BeanCreationException: Error creating
bean with name 'userAccessorTarget' defined in class path
10.resource [applicationContext.xml]: Can't resolve reference to bean 'spacePermissionManager'
while setting property 'spacePermissionManager';
11.nested exception is org.springframework.beans.factory.BeanCreationException: Error creating
bean with name 'spacePermissionManager' defined in class path resource [securityContext.xml]:
12.Can't resolve reference to bean 'spacePermissionManagerTarget' while setting property 'target';
13.nested exception is org.springframework.beans.factory.BeanCreationException:
14.Error creating bean with name 'spacePermissionManagerTarget' defined in class path resource
[securityContext.xml]: Initialization of bean failed;
15.nested exception is org.springframework.jdbc.UncategorizedSQLException: (Hibernate operation):
encountered SQLException [Cannot create PoolableConnectionFactory];
16.nested exception is org.apache.commons.dbcp.SQLNestedException: Cannot create
PoolableConnectionFactory
17.
18....
19.
20.org.apache.commons.dbcp.SQLNestedException: Cannot create PoolableConnectionFactory, cause:
21.
22.
23.java.sql.SQLException: ORA-00604: error occurred at recursive SQL level 1
24.
25.
26.ORA-12705: invalid or unknown NLS parameter value specified
```

One symptom of this problem is that Confluence may refuse to start after midday.

The workaround is to go to 'General Configuration' and set the default time format to "HH:mm".

ORA-01430: column being added already exists in table

If any of confluence tables exist in other oracle schemas in the "same database", those tables will not be created, impairing bits of Confluence functionality that depend on those table(s).
This issue has been filed in CONF-3613.

There is a work around however as found by one of our customers:

```
1.CREATE VIEW wiki_dev.all_objects
2.AS
3.SELECT *
4.FROM sys.all_objects
5.WHERE owner = 'WIKI_DEV'
6.;
```

Configuring Database Character Encoding

Refer to Configuring Database Character Encoding.

Knowledge Base Articles

Unable to Install Oracle on Websphere due to Classloader Preferences (Confluence Knowledge Base) Labels: oracle, websphere

Unable to Configure Oracle Data Source in Websphere (Confluence Knowledge Base) Labels: confluence, oracle, websphere

java.sql.SQLException ORA-00942 table or view does not exist (db permissions) (Confluence Knowledge Base) Labels: confluence, installation, oracle, database

Unique Constraint Violation due to Value too Large for Column (Confluence Knowledge Base) Labels: confluence, oracle, restore_backup

RELATED TOPICS

Database Setup for Oracle

Known Issues for PostgreSQL

Database Performance Issue on PostgreSQL 7

There is an issue in versions of PostgreSQL prior to 8.0 that causes the database to do a full table scan rather than an index lookup for many of the database queries performed by Confluence. This will cause Confluence to slow significantly the more data it has stored.

While we fully support Confluence on PostgreSQL 7, we *highly* recommend upgrading to version 8.0 or higher for this reason.

Configuring Database Character Encoding

Refer to Configuring Database Character Encoding.

Incorrect JDBC Driver Used

If you have downloaded and used an incorrect JDBC Driver version, you will see this error when you attempt to setup the Database connection in the Setup Wizard:

```
2008-04-01 01:45:55,371 ERROR [http-8090-Processor4] [[Standalone].[localhost].[/].[action]]
invoke Servlet.service() for servlet action threw exception
java.lang.UnsupportedClassVersionError: Bad version number in .class file
 at java.lang.ClassLoader.defineClass1(Native Method)
 at java.lang.ClassLoader.defineClass(Unknown Source)
 at java.security.SecureClassLoader.defineClass(Unknown Source)
 at org.apache.catalina.loader.WebappClassLoader.findClassInternal(WebappClassLoader.java:1847)
 at org.apache.catalina.loader.WebappClassLoader.findClass(WebappClassLoader.java:873)
 at org.apache.catalina.loader.WebappClassLoader.loadClass(WebappClassLoader.java:1326)
 at org.apache.catalina.loader.WebappClassLoader.loadClass(WebappClassLoader.java:1205)
 at java.lang.ClassLoader.loadClassInternal(Unknown Source)
 at java.lang.Class.forName0(Native Method)
 at java.lang.Class.forName(Unknown Source)
 at
com.atlassian.confluence.setup.actions.AbstractSetupStandardDatabaseAction.checkDriver(AbstractSetupS
at
com.atlassian.confluence.setup.actions.SetupStandardDatabaseAction.execute(SetupStandardDatabaseActio
at com.opensymphony.xwork.DefaultActionInvocation.invoke(DefaultActionInvocation.java:168)
```

You will need to download the correct JDBC Driver based on your **PostgreSQL Database Version** and also the **JDK** you are using. Below is a guide to selecting the correct JDBC version from PostgreSQL's website:

> *Many other versions of the JDBC driver are available. This includes development versions, compatibility with older JDKs, and previous versions of the driver.*
>
> *To determine JDK/JVM compatibility this following list matches up versions of the JVM with the JDBC specification implemented.*
>
> - *JDK 1.1 - JDBC 1. Note that with the 8.0 release JDBC 1 support has been removed, so look to update your JDK when you update your server.*
> - *JDK 1.2, 1.3 - JDBC 2.*
> - *JDK 1.3 + J2EE - JDBC 2 EE. This contains additional support for javax.sql classes.*
> - *JDK 1.4, 1.5 - JDBC 3. This contains support for SSL and javax.sql, but does not require J2EE as it has been added to the J2SE release.*
> - *JDK 1.6 - JDBC4. Support for JDBC4 methods is limited. The driver builds, but the majority of new methods are stubbed out.*

There is also a matrix on the page which lists all supported versions for the respective PostgreSQL database versions.

Download PostgreSQL JDBC Drivers

Known Issues For SQL Server

Due to numerous reported performance issues with SQL Server 2000, it is strongly recommended that you use SQL Server 2005 instead. SQL Server 2000 is not supported. Confluence with SQL Server 2005 is known to not work very well under high load environments. Please see the deadlock issues below

- Database Deadlock Issues
- Unicode Characters Not Supported By Default
- SQL Server 2005 and Case-Sensitive Schema Objects.
- Case-Sensitive Collation
    - Setting up a New Confluence Instance
    - Migrating an Existing Confluence Instance to a Different Database
- Use jTDS Driver 1.0.3 or Later for JDBC Connection - Reason 1
- Use jTDS Driver 1.0.3 or Later for JDBC Connection - Reason 2
- Configuring Database Character Encoding
- Setting up Authentication Mode of the SQL Server
- Moving from One SQL Server to Another
    - Known Issues

Database Deadlock Issues

Problem

When creating a page, you may encounter deadlocks and get the following error:

```
01.2008-11-11 17:01:21,355 ERROR [http-8080-Processor6] [sf.hibernate.util.JDBCExceptionReporter]
logExceptions Transaction (Process ID 55) was deadlocked on lock resources with another process
and has been chosen as the deadlock victim. Rerun the transaction.
02.   -- url: /confluence/pages/docreatepage.action | userName: tstcreator3 | action: docreatepage
03.2008-11-11 17:01:21,355 ERROR [http-8080-Processor6] [sf.hibernate.util.JDBCExceptionReporter]
logExceptions Transaction (Process ID 55) was deadlocked on lock resources with another process
and has been chosen as the deadlock victim. Rerun the transaction.
04.   -- url: /confluence/pages/docreatepage.action | userName: tstcreator3 | action: docreatepage
05.2008-11-11 17:01:21,357 ERROR [http-8080-Processor6] [sf.hibernate.impl.SessionImpl] execute
Could not synchronize database state with session
06.   -- url: /confluence/pages/docreatepage.action | userName: tstcreator3 | action: docreatepage
07.2008-11-11 17:01:21,357 ERROR [http-8080-Processor6] [sf.hibernate.impl.SessionImpl] execute
Could not synchronize database state with session
08.   -- url: /confluence/pages/docreatepage.action | userName: tstcreator3 | action: docreatepage
09.2008-11-11 17:01:21,359 ERROR [http-8080-Processor6]
[atlassian.xwork.interceptors.XWorkTransactionInterceptor] onThrowable Invoking rollback for
transaction on action '/pages/docreatepage.action (CreatePageAction.doAdd())' due to throwable:
org.springframework.dao.ConcurrencyFailureException: Hibernate operation: could not delete:
[com.atlassian.confluence.pages.Draft#14319839]; SQL []; Transaction (Process ID 55) was
deadlocked on lock resources with another process and has been chosen as the deadlock victim.
Rerun the transaction.; nested exception is java.sql.SQLException: Transaction (Process ID 55) was
deadlocked on lock resources with another process and has been chosen as the deadlock victim.
Rerun the transaction.
10.   -- url: /confluence/pages/docreatepage.action | userName: tstcreator3
11.org.springframework.dao.ConcurrencyFailureException: Hibernate operation: could not delete:
[com.atlassian.confluence.pages.Draft#14319839]; SQL []; Transaction (Process ID 55) was
deadlocked on lock resources with another process and has been chosen as the deadlock victim.
Rerun the transaction.; nested exception is java.sql.SQLException: Transaction (Process ID 55) was
deadlocked on lock resources with another process and has been chosen as the deadlock victim.
Rerun the transaction.
12.Caused by: java.sql.SQLException: Transaction (Process ID 55) was deadlocked on lock resources
with another process and has been chosen as the deadlock victim. Rerun the transaction.
13.          at net.sourceforge.jtds.jdbc.SQLDiagnostic.addDiagnostic(SQLDiagnostic.java:367)
14.          at net.sourceforge.jtds.jdbc.TdsCore.tdsErrorToken(TdsCore.java:2816)
15.          at net.sourceforge.jtds.jdbc.TdsCore.nextToken(TdsCore.java:2254)
16.          at net.sourceforge.jtds.jdbc.TdsCore.getMoreResults(TdsCore.java:631)
17.          at net.sourceforge.jtds.jdbc.JtdsStatement.processResults(JtdsStatement.java:584)
18.          at net.sourceforge.jtds.jdbc.JtdsStatement.executeSQL(JtdsStatement.java:546)
19.          at
net.sourceforge.jtds.jdbc.JtdsPreparedStatement.executeUpdate(JtdsPreparedStatement.java:505)
20.          at
com.mchange.v2.c3p0.impl.NewProxyPreparedStatement.executeUpdate(NewProxyPreparedStatement.java:
105)
21.          at net.sf.hibernate.impl.NonBatchingBatcher.addToBatch(NonBatchingBatcher.java:22)
22.          at net.sf.hibernate.persister.EntityPersister.delete(EntityPersister.java:581)
23.          at net.sf.hibernate.impl.ScheduledDeletion.execute(ScheduledDeletion.java:29)
24.          at net.sf.hibernate.impl.SessionImpl.execute(SessionImpl.java:2449)
25.          at net.sf.hibernate.impl.SessionImpl.executeAll(SessionImpl.java:2435)
26.          at net.sf.hibernate.impl.SessionImpl.execute(SessionImpl.java:2397)
27.          at net.sf.hibernate.impl.SessionImpl.flush(SessionImpl.java:2261)
28.          at net.sf.hibernate.transaction.JDBCTransaction.commit(JDBCTransaction.java:61)
```

This problem will be more prominent under high load environments when there are more concurrent page creates/edits.

Reason

SQL Server is escalating row locks (in this case deleting a draft from the CONTENT table) to table locks and stopping other transactions from accessing the table.

Solution

Configure your database to use the isolation level, Read Committed with Row Versioning. You can do this by executing the following query:

```
1.ALTER DATABASE <database name>
2.        SET READ_COMMITTED_SNAPSHOT ON
3.        WITH ROLLBACK IMMEDIATE;
```

More information on SQL Server isolation levels can be found in the Microsoft documentation.

From version 2.10.x onwards, indexes will be automatically created upon performing an upgrade.
For customers using a previous version, please execute the following DDL against your confluence database:

```
1.CREATE NONCLUSTERED INDEX [c_pageid_idx] ON [dbo].[CONTENT]
2.(
3.[PAGEID] ASC
4.)WITH (PAD_INDEX = OFF, STATISTICS_NORECOMPUTE = OFF, SORT_IN_TEMPDB = OFF, IGNORE_DUP_KEY =
OFF, DROP_EXISTING = OFF, ONLINE = OFF, ALLOW_ROW_LOCKS = ON, ALLOW_PAGE_LOCKS = ON) ON [PRIMARY]
```

Unicode Characters Not Supported By Default

Problem
Non-ASCII characters will not be displayed by Confluence.

Reason
The default SQL Server dialect uses column types that do not support Unicode, specifically the **char**, **varchar** and **text** column types. See CONF-4786 for details.

Solution
To add Unicode support, use the Unicode SQL Server dialect which uses **nchar**, **nvarchar** and **ntext** column types. Unicode SQL Server dialect has the downside of halving the maximum length of each column from 8000 characters to 4000, as every char is stored in two bytes.

Enable Unicode SQL Server dialect on a new setup, perform these steps prior to 'Stage 3 - Database Connection Setup' of the Database Setup For Any External Database:

1. Open the **<confluence installation folder>\confluence\WEB-INF\classes\database-defaults\mssql.properties** file within your Confluence installation folder. In version 2.6 onwards, this file is located in the `/confluence/WEB-INF/lib/confluence.x.x.x.jar` file. Jar files are similar to zip formats, so you may need to extract the `mssql.properties` file and place it in `\confluence\WEB-INF\classes\database-defaults`. You can create the `database-defaults` directory if it does not exist.
2. Comment the line: *dialect=net.sf.hibernate.dialect.SQLServerDialect*
3. Uncomment the line: *#dialect=net.sf.hibernate.dialect.SQLServerIntlDialect*
4. Start the Confluence Setup Wizard

For existing SQL Server instances wishing to enable Unicode support:

1. From Confluence, create an XML site backup
2. From your DBA tool, create a full backup of the Confluence schema contents
3. Stop Confluence
4. Move your home directory
5. Drop all tables from the Confluence schema
6. Change the dialect to Unicode using the above instructions
7. Follow the 'Stage 3 - Database Connection Setup' of Database Setup For Any External Database to setup the connection again and import the XML backup
8. Once the Setup Wizard is complete, stop Confluence
9. Copy any customised content and plugins from your old home directory into the new home directory
10. Start Confluence

This will not restore previously stored Unicode characters.

SQL Server 2005 and Case-Sensitive Schema Objects.

A default installation of SQL Server 2005 on a Windows system configured for English will be case sensitive for schema objects. That is to say that a table called 'CONTENT' is not the same as a table called 'content'.

Confluence will not run correctly in this case. During installation of SQL Server, be sure to choose a **case-insensitive schema**.

Case-Sensitive Collation

'Collation' refers to a set of rules that determine how data is sorted and compared. Case sensitivity is one aspect of collation. Other aspects include sensitivity to kana (Japanese script) and to width (single- versus double-byte characters).

Case-sensitive or case-insensitive collation — how should you create your Confluence database? What about when you are migrating your existing Confluence instance from one database to another?

Setting up a New Confluence Instance

For new Confluence instances, we recommend using case-sensitive collation for your Confluence database, which is the default collation type used by many database systems. The Confluence application itself reduces all usernames into lower-case characters before they are stored in the Confluence database. Therefore, 'joebloggs', 'joeBloggs', 'JoeBloggs', etc. will be treated as the same username on a Confluence installation with case-sensitive database collation.

Migrating an Existing Confluence Instance to a Different Database

The default Confluence Standalone configuration uses case-sensitive database collation. This is often the case with databases on several other systems which were created under default conditions. Therefore, 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.

Information on different collation options in SQL Server 2005 can be found in the Microsoft documentation.

Use jTDS Driver 1.0.3 or Later for JDBC Connection - Reason 1

Problem
When using jTDS driver 1.0.2 or earlier, Confluence may freeze when performing certain functions, and you see a warning like the one displayed below:

```
1.\[ERROR\] ActionSupport - \-An error occured while storing the requested page\!
<org.springframework.jdbc.UncategorizedSQLException: (Hibernate operation): encountered
SQLException
2.\[The amount of data read from the stream is not = length.\]; nested exception is
java.sql.SQLException: The amount of data read from the stream is not = length.>
```

This can occur with jTDS driver 1.0.2 but is fixed in 1.0.3 - see the jTDS homepage . It will prevent backups from suceeding and lock access to the database when viewing certain pages.

Solution
Upgrade to jTDS driver 1.0.3 or later from the jTDS download page.

Use jTDS Driver 1.0.3 or Later for JDBC Connection - Reason 2

Problem
When using MS SQL Server without the jTDS drivers, you may receive an error similar to:

```
net.sf.hibernate.LazyInitializationException: Exception initializing proxy:
\[com.atlassian.confluence.spaces.SpaceDescription#5\]
 at net.sf.hibernate.proxy.LazyInitializer.initializeWrapExceptions(LazyInitializer.java:64)
 at net.sf.hibernate.proxy.LazyInitializer.getImplementation(LazyInitializer.java:164)
...
Caused by: java.sql.SQLException: \[Microsoft\]\[SQLServer 2000 Driver for JDBC\]ResultSet can not
re-read row data for column 2.
 at com.microsoft.jdbc.base.BaseExceptions.createException(Unknown Source)
 at com.microsoft.jdbc.base.BaseExceptions.getException(Unknown Source)
 at com.microsoft.jdbc.base.BaseResultSet.validateColumnIndex(Unknown Source)
 at com.microsoft.jdbc.base.BaseResultSet.getLong(Unknown Source)
 at com.microsoft.jdbc.base.BaseResultSet.getLong(Unknown Source)
 at org.jboss.resource.adapter.jdbc.WrappedResultSet.getLong(WrappedResultSet.java:338)
 at net.sf.hibernate.type.LongType.get(LongType.java:18)
 at net.sf.hibernate.type.NullableType.nullSafeGet(NullableType.java:62)
 at net.sf.hibernate.type.NullableType.nullSafeGet(NullableType.java:53)
 at net.sf.hibernate.type.ManyToOneType.hydrate(ManyToOneType.java:61)
 at net.sf.hibernate.loader.Loader.hydrate(Loader.java:686)
 at net.sf.hibernate.loader.Loader.loadFromResultSet(Loader.java:627)
 at net.sf.hibernate.loader.Loader.instanceNotYetLoaded(Loader.java:586)
 at net.sf.hibernate.loader.Loader.getRow(Loader.java:501)
 at net.sf.hibernate.loader.Loader.getRowFromResultSet(Loader.java:213)
 at net.sf.hibernate.loader.Loader.doQuery(Loader.java:281)
 at net.sf.hibernate.loader.Loader.doQueryAndInitializeNonLazyCollections(Loader.java:133)
 at net.sf.hibernate.loader.Loader.loadEntity(Loader.java:911)
 at net.sf.hibernate.loader.Loader.loadEntity(Loader.java:931)
 at net.sf.hibernate.loader.EntityLoader.load(EntityLoader.java:59)
 at net.sf.hibernate.loader.EntityLoader.load(EntityLoader.java:51)
 at net.sf.hibernate.persister.EntityPersister.load(EntityPersister.java:419)
```

We use a component in Confluence called Hibernate. According to Hibernate Documentation there may be issues with the Microsoft JDBC drivers resulting in the error you are seeing.

Solution
Upgrade to jTDS driver 1.0.3 or later from the jTDS download page. You may also consider alternative SQL Server drivers listed on the

Hibernate page.

Configuring Database Character Encoding

When creating the database schema, the database character encoding must be compatible with the application and Confluence character encoding as described in Configuring Database Character Encoding. If setting MS SQL to use UTF-8 is not an option, you can create a schema that uses UCS-2 encoding, and have the application and Confluence use UTF-8 encoding.

Setting up Authentication Mode of the SQL Server

During a login process to SQL server or while setting up your DB connection through the Setup Wizard, the following error may appear:

```
Login failed for user 'username'. The user is not associated with a trusted SQL Server connection.
(Microsoft SQL Server, Error: 18452).
```

The cause of this error is that the SQL server has been configured to operate in 'Windows Authentication Mode (Windows Authentication)' and doesn't allow
the use of SQL accounts.

In order to resolve this problem, change the Authentication Mode of the SQL server from 'Windows Authentication Mode (Windows Authentication)'
to 'Mixed Mode (Windows Authentication and SQL Server Authentication)'.

Please refer to this Microsoft document for more details.

Moving from One SQL Server to Another

In Confluence, tables are created using the database login name for the schema instead of the dbo owner. Thus when moving data from one SQL server to another, you need to ensure that you **create the database user first before exporting and importing the data**, otherwise that user cannot be created.

To illustrate with an example, say the Confluence database login was 'confuser' — the table names would be 'confuser.table1', 'confuser.table2', etc.

When performing a database backup and restore into the new server, the tables will still be 'confuser.table1' and 'confuser.table2'.

However, if the database user 'confuser' has not been created first on the new server then you can no longer access the tables with a login of 'confuser', because although the database was copied across, the login object was not. You cannot create the login with the same name at this stage as it will complain that there is already an object in the database with the same name (the user object).

Known Issues

- Unicode Characters Not Supported By Default
- Invalid object name hibernate_unique_key due to Invalid Table Name
- Invalid Username or Password when Delegating User Management to use JIRA Logins
- Use jTDS Driver 1.0.3 or Later for JDBC Connection
- Database Deadlock on Microsoft SQL Server

Known Issues for Sybase Database

Find below a list of tips relating to using Sybase with Confluence

Confluence and JIRA shared user base management fails due to wrong casing of database columns and names.

- See Override properties in JIRA to Confluence Bridge

Configuring Database Character Encoding

Refer to Configuring Database Character Encoding

Case sensitivity issues

All versions of Confluence 2.3 require a patch in order to work with Sybase. This is downloadable from http://jira.atlassian.com/browse/CONF-7925.

Improving Database Performance

Diagnosis

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

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

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

Upgrade your Database and Drivers

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

Upgrade your hardware

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

Ensure you have the Latest Database Indices

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

Confluence 2.10 or Manual .ddl Indices

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

Additional Indices not Included in 2.10

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

Tuning the Database Connection Pool

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

Configure a Database Query Timeout

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

Related Articles

Troubleshooting Database Issues.

Creating a Lowercase Page Title Index

Diagnosis

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

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

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

> *select ... from CONTENT page0_, SPACES space1_*
> *where page0_.CONTENTTYPE='PAGE'*

> *and ((**lower(space1_.SPACEKEY)= @P0** and page0_.SPACEID=space1_.SPACEID)*
> *and(**lower(page0_.TITLE)= @P1**)*
> *and(page0_.PREVVER is null)and(page0_.CONTENT_STATUS='current'))*

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

Generic solution

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

```
1.create index CONFTITLE_LOWER on CONTENT(lower(TITLE));
```

Sources:

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

SQL Server

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

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

Sources:

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

MySQL

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

Source:

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

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

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

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

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

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

```
01.<query name="confluence.page_findLatestBySpaceKeyTitle"><![CDATA[
02.          from Page page
03.          where lower(page.space.key) = :spaceKey and
04.                    lower(page.title) = :pageTitle and
05.                    page.originalVersion is null and
06.                    page.contentStatus = 'current'
07.]]></query>
08.
09.<query name="confluence.page_findLatestBySpaceKeyTitleOptimisedForComments"><![CDATA[
10.          from Page page
11.          left join fetch page.comments as theComments
12.          left join fetch theComments.children
13.          where lower(page.space.key) = :spaceKey and
14.                    lower(page.title) = :pageTitle and
15.                    page.originalVersion is null and
16.                    page.contentStatus = 'current'
17.]]></query>
```

DB2 for Linux, Unix or Windows

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

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

Related pages

- Improving Database Performance
- CONF-10030: Queries that use 'lower' do not use index because of case sensitivity

Troubleshooting External Database Connections

The Atlassian Database Check Utility

A common administration issue when configuring Confluence is identifying database connectivity problems. The helper JSP can isolate database connectivity issues. It simply checks whether you can connect to a database with your application server. If your application server cannot connect to the database, Confluence certainly will not.

Purpose:

- Check that your application server can successfully query your database (either via immediate JDBC connectivity or a Datasource in the context of your application server).
- Pinpoint problems in your configuration which may occur if the above is failing.

Functionality:

- Check that a JDBC Driver can be loaded into memory and view what is already loaded.
- Connect to a JDBC URL and do a 'select 1' from the database.
- Find a DataSource in the JNDI environment and do the above.
- View the System classpath (to ensure that the JDBC jar file is there).

Using the Utility

If you have already setup Confluence completely

1. Download testdatabase.jsp to 'MY-CONF-INSTALL-PATH\confluence'
2. Restart Confluence
3. Go to http://MY-CONF-SERVER:MY-CONF-PORT/testdatabase.jsp
4. Check that your database driver is loaded into memory. If not, check the system classpath for the JDBC driver file, and that the driver is in the `<confluence-install>/lib` folder (for Confluence version 2.10 onwards) or `<confluence-install>/common/lib` (for earlier versions) as per these instructions.
5. Enter the DB settings Confleunce is using and test the database. If an error appears, check that the db service is running, the location matches, and that any users specified actually exist with the right login and permissions. You may be able to find a workaround by googling the error.

If you cannot setup Confluence because of an error in 'Configuring Database'

1. Record the DB settings you are using for your direct JDBC or datasource connection in the 'Configure Database' step of your setup
2. Download testdatabase.jsp to 'MY-CONF-INSTALL-PATH\confluence'
3. Rename 'MY-CONF-INSTALL-PATH\confluence\WEB-INF\web.xml' to 'backup web.xml' (This disables redirection)
4. Restart Confluence
5. Go to http://MY-CONF-SERVER:MY-CONF-PORT/testdatabase.jsp
6. Check that your database driver is loaded into memory. If not, check the system classpath for the JDBC driver file, and that the driver is in the confluence-install/common/lib folder as per the instructions
7. Enter the DB settings you recorded and test the database. If an error appears, check that the db service is running, the location matches, and that any users specified actually exist with the right login and permissions. You may be able to find a workaround by googling the error.
8. After correcting the error, rename 'MY-CONF-INSTALL-PATH\confluence\WEB-INF\backup web.xml' back to 'web.xml'

Notes

This JSP is bundled in Confluence 1.2.4 and onwards. It can be used with earlier versions quite easily.
If you use this utility, please let us know ways in which we could improve it or leave helpful hints for others here.

For a comprehensive set of database manual that might be helpful for troubleshooting, please refer to the following links:

- PostgreSQL
- MySQL

Request Technical Support

If you're still stuck after attempting the suggestions above, lodge a free technical support request with information on your database setup.

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:

```
1.<bean id="transactionManager" class=
"org.springframework.orm.hibernate.HibernateTransactionManager">
2.          <property name="sessionFactory">
3.                    <ref bean="sessionFactory"/>
4.          </property>
5.          <property name="defaultTimeout" value="120"/>
6.</bean>
```

The timeout is measured in seconds and will forcibly abort queries that take longer than this. In some cases, these errors are not handled gracefully by Confluence and will result in the user seeing the Confluence error page.

4. Start Confluence.

Once the timeout is working properly in your test environment, migration the configuration change to Confluence.

 You will need to reapply these changes when upgrading Confluence, as the original `databaseSubsystemContext.xml` file changes from version to version.

Troubleshooting the Embedded Database (hSQL DB)

java.sql.SQLException: User not found: SA

Also see http://hsqldb.sourceforge.net/doc/guide/ch01.html#N101C2 .

HSQLDB periodically must update its files to represent changes made in the database. In doing so, it must delete the current confluencedb.data file on the filesystem (beneath conf.home/database) and replace it with a new one.

If an administrator issues a shutdown on Confluence in this period, data can be lost, and is typically noticed by the error message, when starting Confluence up again, of 'User not found: SA' .

Users encountering this problem should seek to restore backups, contained in the backup directory beneath confluence.home. If daily backups have been turned off, and no other copy of data remains, the data is lost.

 HSQLDB should not be used as a production database. It is included for evaluation purposes only.

Hibernate logging

It can be useful to enable detailed Hibernate logging when debugging problems with HSQL.

Connecting to the Embedded Database

Connecting to the Embedded HSQL Database can be quite difficult. You may need to connect to the database to retrieve information, or for troubleshooting purposes.

Please follow the instructions on how to you can connect to the embedded HSQL Database using the free Database Administration Tool DBVisualizer.

Connecting to HSQLDB using DBVisualizer

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

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

Prerequisites

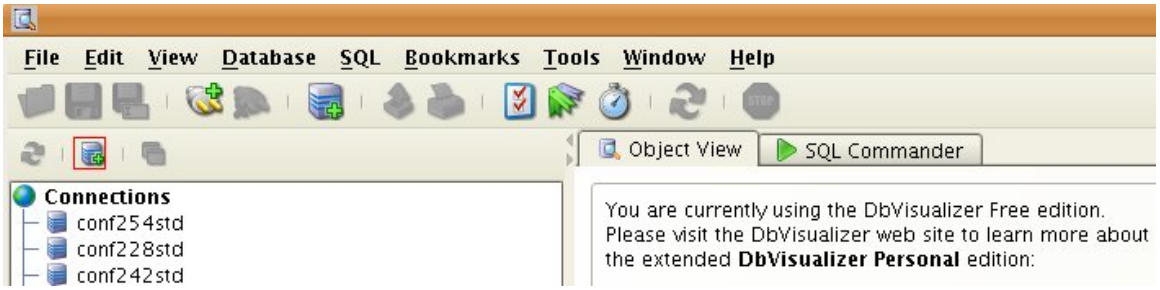
1. Download and install the latest copy of DBVisualizer.
2. You will also need to download a copy (preferably the latest version) of HSQLDB
3. Extract the contents of the HSQLDB archive
4. Ensure that Confluence is not running.

Connection Procedure

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

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

1. Enter Connection Name



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

2. Select JDBC Driver

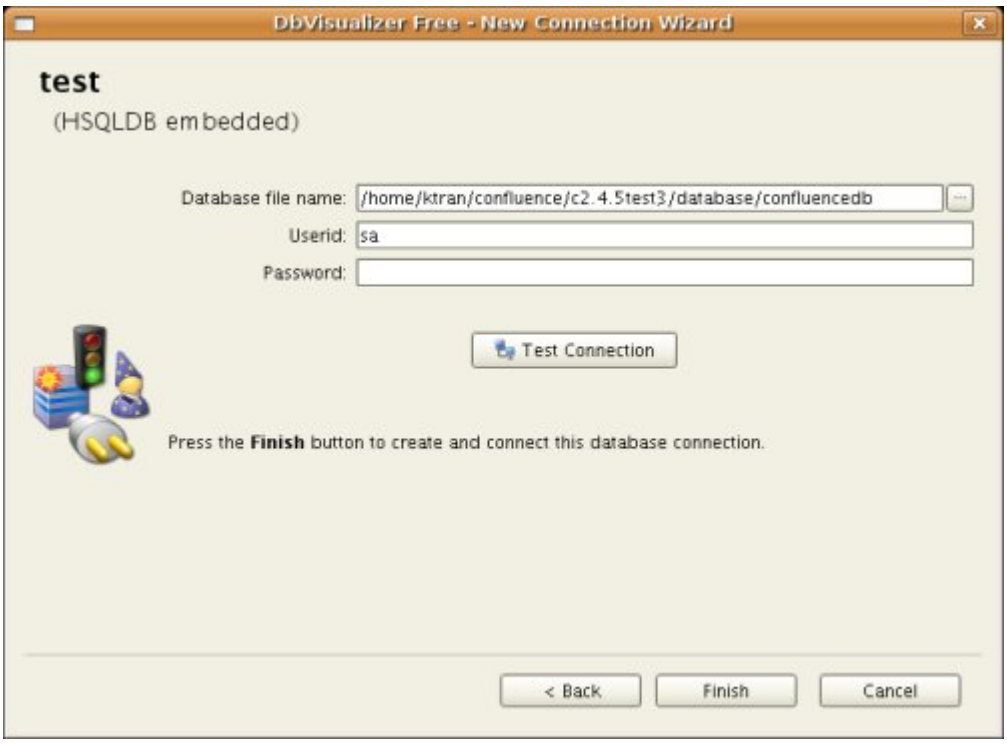


1. From the drop down list select **HSQLDB Embedded**
2. Click on Load Driver Files
3. Browse to directory where the `HSQLDB.jar` file is located

3. Select Database Path

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

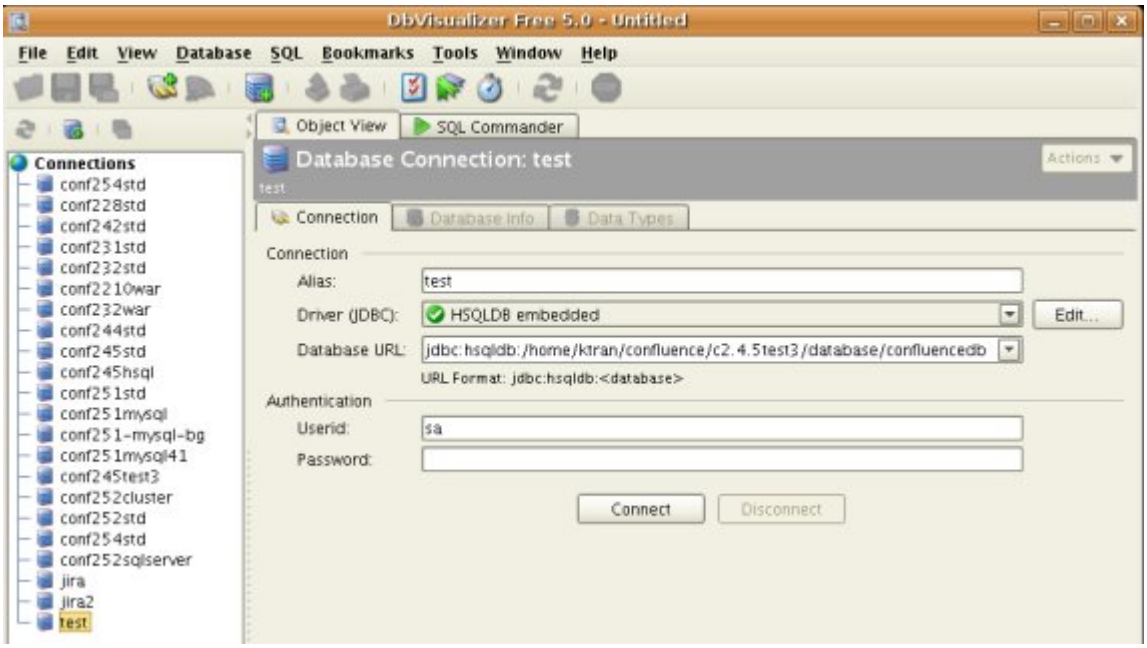
4. Enter Connection Details



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

refer to the example screenshot above if you are unsure

5. Connect to embedded Database



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

HSQL database manager

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

Related Topics

Universal SQL client Squirrel
HSQL
Enable Hibernate Logging
Database Tables Reference
Confluence data model

Database Tables Reference

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

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



Right Click and Select *Save Link As* here to download this image.

Upgrading From HSQL 1.7.1 to 1.8

If you have received an error message while upgrading Confluence which said "HSQL Database needs to be cleaned. Please run HSQL17To18PreUpgradeCleaner.", and referred you to this page, you need to do the following:

1. Shut down Confluence
2. Open a command window.
3. Change your directory to the directory you set as your confluence.home directory.
4. Execute the following command: `java -jar <confluence install directory>/hsqlcleaner/hsqlcleaner.jar`
5. Start Confluence -- the upgrade will now work.

> ⚠ Don't run this unless asked to by Confluence - it won't work

Guide to using Apache Tomcat's Virtual Hosts

Introduction

Sometimes it is necessary to have Tomcat serve different applications on the same context path, but different host names. Most commonly, this is when trying to use a simple mod_proxy configuration with Apache.

Tomcat configuration

Tomcat allows name-based virtual hosting, where the hostname of the request determines which application processes it. The following configuration shows how two virtual hosts can be configured for Jira and Confluence on the same Tomcat instance:

Confluence application server URL	http://confluence-app-server.internal.example.com:8080/
JIRA application server URL	http://jira-app-server.internal.example.com:8080/

Below is a minimal configuration of Tomcat's server.xml which configures separate hosts for JIRA and Confluence on the URLs above.

```
<Server port="8005" shutdown="SHUTDOWN">
    <Service name="Catalina">
        <Connector port="8080" />
        <Engine name="Catalina" defaultHost="confluence-app-server.internal.example.com">

            <Host name="confluence-app-server.internal.example.com" appBase="webapps">
        <Context path="" docBase="/opt/webapps/confluence-2.2/confluence"/>
                <Logger className="org.apache.catalina.logger.FileLogger"/>
            </Host>

            <Host name="jira-app-server.internal.example.com" appBase="webapps">
        <Context path="" docBase="/opt/webapps/jira-3.6.1/jira.war"/>
                <Logger className="org.apache.catalina.logger.FileLogger"/>
            </Host>

        </Engine>
    </Service>
</Server>
```

Points to note:

- the HTTP connector is accessible on port 8080, as per the URLs above
- the AJP connector is not included in this minimal configuration. If you want to mod_jk with this configuration, you need to ensure you add it.
- in the Engine, the Confluence host is configured as the default host. The default host is used if the request doesn't contain a Host header, or is accessed by a name the server doesn't recognise.
- because the above configuration uses name-based virtual hosting, you need to have entries in your DNS server for "confluence-app-server" and "jira-app-server" that point to the application server

More information

After doing this, you can configure Apache virtual hosts to map subdomains like jira.example.com and confluence.example.com to your application.

How to dump Active Directory data to a file

You can extract all the data present in your Active Directory onto a file using the following command (please ensure you are logged in with sufficient rights to do this)

```
csvde -f test.csv
```

This command will perform a CSV dump of every entry in your Active Directory server. You should be able to see the full DN's of users and groups.

Known Issues with Enterprise or Webhosting environments

When you attempt to run Confluence, you may get the following error:

```
java.security.AccessControlException: access denied (java.lang.RuntimePermission
accessDeclaredMembers)
                    at
java.security.AccessControlContext.checkPermission(AccessControlContext.java(Compiled Code))
                    at
java.security.AccessController.checkPermission(AccessController.java(Compiled Code))
                    at java.lang.SecurityManager.checkPermission(SecurityManager.java(Compiled
Code))
```

Some of the libraries Confluence relies on to function make use of features of the Java language that may be restricted by Java security policies. This does not normally cause any problems: the default security configuration of most application servers will happily run Confluence. However, in some shared-hosting or enterprise environments, security settings may be such that Confluence can not function.

The permissions required by Confluence to run are detailed in the sample policy file below. You may need to give this information to your systems administrator so that they can be deployed with the Confluence application.

```
grant codeBase "file:${catalina.home}/webapps/confluence/-" {
permission java.security.AllPermission;
};

grant {
permission java.lang.RuntimePermission "accessDeclaredMembers";
permission java.lang.reflect.ReflectPermission "suppressAccessChecks";
permission java.lang.RuntimePermission "defineCGLIBClassInJavaPackage";
};
```

Managing Application Server Memory Settings

The minimum and maximum JVM heap space allocated to the application server affects performance. Confluence administrators may wish to modify this value from the defaults depending on their server load. This document only provides guidelines rather than rules, so administrators optimising for performance should use this document as a starting point only.



For a comprehensive overview of memory management, and memory tuning in Confluence under Sun JRE, please read Garbage Collector Performance Issues

Testing For Optimum Memory Settings

In the general case, both JIRA & Confluence users will benefit from setting the minimum and maximum values identical. In larger installations, there is benefit to memory tuning, if there is a perceived performance issue. If you are experiencing Out of Memory Heap errors, try doubling the -Xmx and -Xms values for your installation to see if this resolves or helps resolve your issue. If not, please lodge a support ticket as there may be other factors contributing.

Memory usage is most likely to be maximised under peak load, and when creating a site XML backup. In many cases, the backup can be the cause of the OOM, so increase -Xmx values and verify if a backup was occurring at the time of OOM. A quick rule of thumb for gauging the success of a memory adjustment is using simple anecdotal evidence from users. Is it snappier? The same? How does it handle while a backup is occurring?



Atlassian recommends in normal use, to disable the XML backup and use a Production Backup Strategy.

- If you normally perform manual XML site backups on your server, test your maximum memory requirements by performing a site XML backup while the server is under maximum load
- If you do not create manual XML site backups, simply monitor the server while under maximum load

Applying Memory Settings

Users of the EAR/WAR distribution set their memory in their application server. For Confluence Standalone, the heap space is set in the Confluence install directory.

- Windows users edit `TOMCAT_HOME/bin/setenv.bat`
- Other OS users edit `TOMCAT_HOME/bin/setenv.sh`



If you are running Tomcat as a Windows service, you may need to run a `regedt32` to increase memory as in the article on increasing JIRA memory. Although the article discusses JIRA, it should also be applicable to Confluence.

The relevant parameter is

```
JAVA_OPTS="-Xms256m -Xmx1024m"
```

Where `-Xms` is the minimum and `-Xmx` is the maximum memory available to Apache Tomcat. In this sample the maximum memory has been set to 1024 megabytes, but you will have to figure out the best setting based on your particular environment. In regedit, it will be listed as JvmMs and JvmMx.

Note: In newer Tomcat distributions the file may not exist. Feel free to create it in the ./bin directory.

Related Topics

- Garbage Collector Performance Issues
- 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
- Websphere- Tuning JVM

Pull down RSS Feeds or use the Repository plugin through a web proxy

You will need to make Confluence aware of your proxy.

Running Confluence behind Apache

Introduction

Running Confluence behind a web server should be done for performance reasons in high-load environments. In general, web server caching and thread management is far superior to that provided by your application server's HTTP interface.

To run Confluence behind the Apache httpd web server, there are two main configuration options: **mod_jk** or **mod_proxy**.

Connection type	Features
mod_proxy (also known as reverse proxy)	• **recommended connection method** • simple HTTP proxy to application server • works with all application servers • if application paths are consistent, there is minimal load on the web server
mod_jk (also known as AJP)	• uses the AJP binary protocol • provides failover (and load balancing, which Confluence supports only with a clustered license) • only works with some application servers (typically Tomcat) • if application paths are consistent, there is some load on the web server to translate requests to AJP

Features common to both mod_proxy and mod_jk
• application paths must be consistent to avoid complex and slow URL rewriting • works with name-based virtual hosting, both on web server and app server • web server keeps a pool of connections to application server

Mod_proxy documentation

- Using Apache with mod_proxy is the main documentation for this configuration.
- If you want to set up the common configuration of JIRA and Confluence virtual hosts, you can use Apache's virtual hosts with separate application servers, then Tomcat's virtual hosts to run both applications on a single instance of Tomcat.

Mod_jk documentation

- Using Apache with mod_jk is the main documentation for this configuration.
- You can follow a similar method to the mod_proxy documentation above for setting up virtual hosts in Apache and Tomcat, if required.

Mod_jk2 not supported

The misleadingly-named mod_jk2 is an older method of connecting to Tomcat from Apache. Since mod_jk2 is no longer supported by the Apache Foundation, we do not support this configuration, and are not updating our mod_jk2 documentation. Mod_jk2 also has unresolved problems with Unicode URLs; you need to use either mod_proxy or mod_jk for international characters to work correctly in Confluence.

Other related documentation

- Apache and Apache Connector Tips
- Using the (older) mod_jk2 connector
- Configuring Tomcat's URI encoding
- Adding SSL for Secure Logins and Page Security

General Apache Configuration Notes

On this page:

- Prefer Apache mod_deflate to Confluence's built-in gzip implementation
- Ensure keepalive is enabled
- Enable keepalive for recent MSIE user agents

Prefer Apache mod_deflate to Confluence's built-in gzip implementation

1. Disable gzip in confluence. See Compressing an HTTP Response within Confluence.
2. Enable gzip compression in Apache. For RedHat distributions this can be achieved by adding the following lines:

```
AddOutputFilterByType DEFLATE text/html text/plain text/xml text/css
application/x-javascript

# ensure sensible defaults
DeflateBufferSize 8192
DeflateCompressionLevel 4
DeflateMemLevel 9
DeflateWindowSize 15
```

Ensure keepalive is enabled

```
KeepAlive On
```

Enable keepalive for recent MSIE user agents

The standard Apache SSL configuration is very conservative when it comes to MSIE and SSL. By default all keepalives are disabled when using HTTPS with MSIE. While MSIE will always be *special*, the issues with SSL and MSIE have been solved since Service Pack 2 for Windows XP, released over 4 years go. For anyone using an XP machine SP2 or above, it is safe to allow keepalive for MSIE 6 and above.

Remove the following lines:

```
SetEnvIf User-Agent ".*MSIE.*" \
        nokeepalive ssl-unclean-shutdown \
        downgrade-1.0 force-response-1.0
```

Add these in their place:

```
BrowserMatch "MSIE [1-5]" nokeepalive ssl-unclean-shutdown downgrade-1.0 force-response-1.0
BrowserMatch "MSIE [6-9]" ssl-unclean-shutdown
```

RELATED TOPICS

Running Confluence behind Apache
Apache and Apache Connector Tips
Configuring Tomcat's URI encoding
Adding SSL for Secure Logins and Page Security

Using Apache with mod_jk

Introduction

The Apache web server is often used in front of an application server to improve performance in high-load environments. Mod_jk allows request forwarding to an application via a protocol called AJP. Configuration of this involves enabling mod_jk in Apache, configuring a AJP connector in your application server, and directing Apache to forward certain paths to the application server via mod_jk.

Mod_jk is sometimes preferred to mod_proxy because AJP is a binary protocol, and because some site administrators are more familiar with it than with mod_proxy..

The scope of this documentation is limited to configuring the AJP connector in Tomcat 5.x. Other application servers may support AJP connectors; please consult your application server documentation for instructions on how to configure it.

The configuration below assumes your Confluence instance is accessible on the same path on the application server and the web server. For example:

Externally accessible (web server) URL	http://www.example.com/confluence/
Application server URL (HTTP)	http://app-server.internal.example.com:8080/confluence/

The AJP connection of the application server is set to: app-server.internal.example.com:8009.

Configuring mod_jk in Apache

The standard distribution of Apache does not include mod_jk. You need to download it from the JK homepage and put the mod_jk.so file in your Apache modules directory.

Next, add the following in **httpd.conf** directly or included from another file:

```
# Put this after the other LoadModule directives
LoadModule jk_module modules/mod_jk.so

# Put this in the main section of your configuration (or desired virtual host, if using Apache
virtual hosts)
JkWorkersFile conf/workers.properties
JkLogFile logs/mod_jk.log
JkLogLevel info

JkMount /confluence worker1
JkMount /confluence/* worker1
```

Configuring workers.properties

Create a new file called 'workers.properties', and put it in your Apache conf directory. (The path for workers.properties was one of the configuration settings above.)

```
worker.list=worker1

worker.worker1.host=app-server.internal.example.com
worker.worker1.port=8009
worker.worker1.type=ajp13
```

Tomcat 5.x configuration

In Tomcat 5, the AJP connector is enabled by default on port 8009. An absolutely minimal Tomcat server.xml is below for comparison. The relevant line is the Connector with port 8009 – make sure this is uncommented in your server.xml.

```
<Server port="8005" shutdown="SHUTDOWN">
    <Service name="Catalina">

        <!-- Define a HTTP/1.1 Connector on port 8080 -->
        <Connector port="8080" />

        <!-- Define an AJP 1.3 Connector on port 8009 -->
        <Connector port="8009" protocol="AJP/1.3" />

        <Engine name="Catalina" defaultHost="localhost">
            <Host name="localhost" appBase="webapps">
        <Context path="/confluence" docBase="/opt/webapps/confluence-2.2/confluence"/>
                <Logger className="org.apache.catalina.logger.FileLogger"/>
            </Host>
        </Engine>
    </Service>
</Server>
```

Points to note:

- the Connector on port 8009 has protocol of "AJP/1.3". This is critical.
- the Context path of the Confluence application is "/confluence". This must match the path used to access Confluence on the web server.
- we recommend keeping your application Contexts outside the server.xml in Tomcat 5.x. The above example includes them for demonstration only.

Improving the performance of the mod_jk connector

The most important setting in high-load environments is the number of processor threads used by the Tomcat AJP connector. By default, this is 200, but you should increase it to match Apache's maxThreads setting (256 by default):

```
<Connector port="8009" minSpareThreads="5" maxThreads="256" protocol="AJP/1.3" />
```

All the configuration parameters for the AJP connector are covered in the Tomcat documentation.

Ensuring UTF-8 compatibility

If you have problems downloading attachments with non-ASCII characters in the filename, add the following to your Apache configuration:

```
JkOptions +ForwardURICompatUnparsed
```

And specify UTF-8 as the URIEncoding in the AJP connector configuration:

```
<Connector port="8009" protocol="AJP/1.3" URIEncoding="UTF-8" />
```

These settings are discussed further on Configuring Tomcat's URI encoding.

More information

The Tomcat JK website has complete documentation on workers.properties and Apache configuration. You can also find information there on how to use mod_jk with IIS.

Alternatives

If you're not happy with mod_jk, or find it too difficult to configure, you can:

- use mod_proxy, which works with any application server, and together with mod_proxy_html allows complex URL rewriting to deal with different application paths on the web server and the application server.

Using Apache with mod_proxy

This page describes how to integrate Confluence into an Apache website, using mod_proxy.

On this page:

- Simple configuration
  - Set the context path
  - Configure mod_proxy
  - Set the URL for redirection
- Complex configuration
- Adding SSL
  - More information
  - Alternatives



The content on this page relates to platforms which are not supported for Confluence. Consequently, Atlassian can not 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.

There are some common situations where you might do this:

- You have an existing Apache-based website, and want to add Confluence to the mix (eg. http://www.example.com/confluence).
- You have two or more Java applications, each running in their own application server on different ports, eg. http://localhost:8080/confluence and http://localhost:8081/jira. By setting up Apache with mod_proxy, you can have both available on the regular HTTP port (80), eg. at http://www.example.com/confluence and http://www.example.com/jira. If you are running JIRA and Confluence, we recommend this setup. It allows each app to be restarted, managed and debugged separately.

This page describes how to configure mod_proxy. We describe two options:

- If you want a URL like http://www.example.com/confluence/, go to the simple configuration.
- If you want a URL like http://confluence.example.com/, go to the complex configuration.

Simple configuration

Set the context path

First, set your Confluence application path (the part after hostname and port) correctly. Say you want Confluence available at http://www.example.com/confluence/, and you currently have it running at http://localhost:8080/. The first step is to get Confluence available at http://localhost:8080/confluence/.

To do this in Tomcat (bundled with Confluence), edit `conf/server.xml`, locate the "Context" definition:

```
1.<Context path="" docBase="../confluence" debug="0" reloadable="true">
```

and change it to:

```
1.<Context path="/confluence" docBase="../confluence" debug="0" reloadable="true">
```

Then restart Confluence, and ensure you can access it at http://localhost:8080/confluence/

Configure mod_proxy

Now enable mod_proxy in Apache, and proxy requests to the application server by adding the example below to your Apache httpd.conf (note: the files may be different on your system; the JIRA docs describe the process for Ubuntu/Debian layout):

```
01.# Put this after the other LoadModule directives
02.LoadModule proxy_module /usr/lib/apache2/modules/mod_proxy.so
03.LoadModule proxy_http_module /usr/lib/apache2/modules/mod_proxy_http.so
04.
05.# Put this in the main section of your configuration (or desired virtual host, if using Apache
virtual hosts)
06.ProxyRequests Off
07.ProxyPreserveHost On
08.
09.<Proxy *>
10.          Order deny,allow
11.          Allow from all
12.</Proxy>
13.
14.ProxyPass /confluence http://localhost:8080/confluence
15.ProxyPassReverse /confluence http://localhost:8080/confluence
16.<Location /confluence>
17.          Order allow,deny
18.          Allow from all
19.</Location>
```

Note to Windows Users

It is recommended that you specify the absolute path to the `mod_proxy.so` and `mod_proxy_http.so` files.

Set the URL for redirection

You will need to modify the `server.xml` file in your tomcat's *conf* directory and set the URL for redirection.

Locate this code segment

```
1.<Connector port="8080" maxHttpHeaderSize="8192"
2.                           maxThreads="150" minSpareThreads="25" maxSpareThreads="75"
3.                           enableLookups="false" redirectPort="8443" acceptCount="100"
4.                           connectionTimeout="20000" disableUploadTimeout="true" />
```

And append the following segment:

```
1.<Connector port="8080" maxHttpHeaderSize="8192"
2.                           maxThreads="150" minSpareThreads="25" maxSpareThreads="75"
3.                           enableLookups="false" redirectPort="8443" acceptCount="100"
4.                           connectionTimeout="20000" disableUploadTimeout="true"
5.                           proxyName="www.example.com" proxyPort="80"/>
```

Replace **www.example.com** with the URL you wish to be redirected to.

Complex configuration

A complex configuration involves using the mod_proxy_html filter to modify the proxied content en-route. This is required if the Confluence path differs between Apache and the application server. For example:

Externally accessible (Apache) URL	http://confluence.example.com/
Application server URL	http://app-server.internal.example.com:8080/confluence/

Notice that the application path in the URL is different in each. On Apache, the path is /, and on the application server the path is /confluence.

For this configuration, you need to install the mod_proxy_html module, which is not included in the standard Apache distribution.

Alternative solutions are discussed below.

```
# Put this after the other LoadModule directives
LoadModule proxy_module modules/mod_proxy.so
LoadModule proxy_http_module modules/mod_proxy_http.so
LoadModule proxy_html_module modules/mod_proxy_html.so

<VirtualHost *>
        ServerName confluence.example.com

        # Put this in the main section of your configuration (or desired virtual host, if
using Apache virtual hosts)
        ProxyRequests Off
        ProxyPreserveHost On

        <Proxy *>
                Order deny,allow
                Allow from all
        </Proxy>

        ProxyPass / http://app-server.internal.example.com:8080/confluence
        ProxyPassReverse / http://app-server.internal.example.com:8080/confluence

        ProxyHTMLURLMap / /confluence/

        <Location />
                Order allow,deny
                Allow from all
        </Location>
</VirtualHost>
```

The ProxyHTMLURLMap configuration can become more complex if you have multiple applications running under this configuration. The mapping should also be placed in a Location block if the web server URL is a subdirectory and not on a virtual host. The Apache Week tutorial has more information how to do this.

Adding SSL

If you're running Apache in front of Tomcat, it's a good idea to terminate your SSL configuration at Apache, then forward the requests to Tomcat over HTTP. You can set up Apache to terminate the SSL connection and use the ProxyPass and ProxyPassReverse directives to pass the connection through to Tomcat (or the appropriate application server) which is running Confluence.

1. Create a new SSL host by creating a virtual host on 443
2. The standard http connection on apache could be used to redirect to https if you want or it could just be firewalled.
3. Within the VirtualHost definition:
    a. define the SSL options (SSLEngin and SSLCertificateFile)
    b. define the ProxyPass and ProxyPassReverse directives to pass through to Tomcat.

Because of how the ProxyPass and ProxyPassReverse directives work, you should not need to modify the tomcat installation at all.

Most of the relevant Apache Config:

```
Listen 443

NameVirtualHost *:443
<VirtualHost *:443>
        SSLEngine On
        SSLCertificateFile /etc/apache2/ssl/apache.pem
        ProxyPass / http://localhost:8080/
        ProxyPassReverse / http://localhost:8080/
</VirtualHost>
```

More information

- The mod_proxy_html site has documentation and examples on the use of this module in the complex configuration.
- Apache Week has a tutorial that deals with a complex situation involving two applications and ProxyHTMLURLMap.
- Using Apache with virtual hosts and mod_proxy shows how to configure the special case where you want JIRA and Confluence running on separate application servers on virtual host subdomains.

Alternatives

If Tomcat is your application server, you have two options:

- use mod_jk to send the requests to Tomcat
- use Tomcat's virtual hosts to make your Confluence application directory the same on the app server and the web server, removing the need for the URL mapping.

If your application server has an AJP connector, you can:

- use mod_jk to send the requests to your application server.

Using Apache with virtual hosts and mod_proxy

Introduction

The Apache web server is often used in front of an application server to improve performance in high-load environments. Mod_proxy simply redirects requests for certain URLs to another web server, so it typically requires no additional configuration on the application server.

This page documents a very common configuration request: configuring JIRA and Confluence on two Apache virtual hosts, running on different application servers. This is just a special case of mod_proxy configuration.

You can use virtual hosts in your application server if you want to run JIRA and Confluence on the same application server. There is a sample configuration for Tomcat you can use after configuring Apache.

Apache configuration

For this configuration to work properly, the application paths must be the same on both the application servers and the web server. For both JIRA and Confluence below, this is /.

JIRA external URL	http://jira.example.com/
JIRA application server URL	http://jira-app-server.internal.example.com:8080/
Confluence external URL	http://confluence.example.com/
Confluence application server URL	http://confluence-app-server.internal.example.com:8080/

Add the following to your Apache httpd.conf:

```
01.# Put this after the other LoadModule directives
02.LoadModule proxy_module /usr/lib/apache2/modules/mod_proxy.so
03.LoadModule proxy_http_module /usr/lib/apache2/modules/mod_proxy_http.so
04.
05.# Put this with your other VirtualHosts, or at the bottom of the file
06.NameVirtualHost *
07.<VirtualHost *>
08.          ServerName confluence.example.com
09.
10.          ProxyRequests Off
11.          <Proxy *>
12.                    Order deny,allow
13.                    Allow from all
14.          </Proxy>
15.
16.          ProxyPass / http://confluence-app-server.internal.example.com:8080/
17.          ProxyPassReverse / http://confluence-app-server.internal.example.com:8080/
18.          <Location />
19.                    Order allow,deny
20.                    Allow from all
21.          </Location>
22.</VirtualHost>
23.<VirtualHost *>
24.          ServerName jira.example.com
25.
26.          ProxyRequests Off
27.          <Proxy *>
28.                    Order deny,allow
29.                    Allow from all
30.          </Proxy>
31.
32.          ProxyPass / http://jira-app-server.internal.example.com:8080/
33.          ProxyPassReverse / http://jira-app-server.internal.example.com:8080/
34.          <Location />
35.                    Order allow,deny
36.                    Allow from all
37.          </Location>
38.</VirtualHost>
```

Points to note:

- ProxyPass and ProxyPassReverse directives send traffic from the web server to your application server.
- The application path is the same on the application server and on the web server (both are /).
- Because the above configuration uses name-based virtual hosting, you must configure your DNS server to point both names (jira.example.com, confluence.example.com) to your web server.

More information

For different ways to configure mod_proxy, see Using Apache with mod_proxy.

If you use Tomcat, mod_jk provides a different way of connecting Apache via AJP. You can also use the above configuration with just one application server if you use Tomcat's virtual hosts.

Setting Up Public Access

Granting of permissions to use Confluence can be done on the basis of membership of a group, to a particular user, or to the 'Anonymous' user. There is not an actual user named 'Anonymous', it is just a name for a category of granted permissions.

In the security administration of Confluence, the 'Anonymous' user includes all logged-in users, *and* anonymous users who have not logged in. That is, if you allow the 'Anonymous' user to do something, you are allowing all users to do it.

To enable public access to a confluence space, you must give 'Anonymous' the following permissions:

- The global '**Use Confluence**' permission. This permission determines whether a user can access the Confluence installation at all, and is set by the site administrator in **Administration** -> **Global Permissions**.
- The relevant space permissions. The 'Anonymous' user must have at least the 'View Space' permission for a space to be publicly accessible. You set these permissions via '**Browse Space**' -> '**Space Admin**' -> '**Permissions**'.

While these two permissions are the bare minimum necessary for public access to a space, you may wish to give 'Anonymous' additional permissions if you want a space to allow anonymous comments, or anonymous editing.

We **severely warn against** giving the 'Anonymous' user any administrative privileges, either within a space, or especially globally over the Confluence instance. Giving administrative privileges to untrusted users may lead to a serious security compromise of your site.

Related

Enabling or Disabling Public Signup

Setup a mail session in standalone version

1. Add the following to your web.xml file (insert it just before </web-app>):

```
1.<resource-ref>
2.          <description>Test description</description>
3.          <res-ref-name>mail/Session</res-ref-name>
4.          <res-type>javax.mail.Session</res-type>
5.          <res-auth>Container</res-auth>
6.</resource-ref>
```

2. Add the following to your server.xml file (anywhere inside the <context ...> tags)

For Tomcat 5.0.x

```
01.<Resource name="mail/Session" auth="Container" type="javax.mail.Session"/>
02.<ResourceParams name="mail/Session">
03.          <parameter>
04.               <name>mail.smtp.host</name>
05.               <value>mail.example.com</value>
06.          </parameter>
07.          <parameter>
08.               <name>mail.smtp.port</name>
09.               <value>25</value>
10.          </parameter>
11.</ResourceParams>
```

For Tomcat 5.5.x and Tomcat 6.x

```
1.<Resource name="mail/Session" auth="Container" type="javax.mail.Session" mail.smtp.host=
"mail.example.com" mail.smtp.port="25" />
```

To use the TLS protocol, add the property mail.smtp.starttls.enable="true" to this configuration.

You only need to specify the `mail.smtp.port` parameter if you are using a non-standard port for SMTP.

3. Restart the server. Go to create SMTP mail server screen. For JNDI location, enter:

```
java:comp/env/mail/Session
```

Start Confluence automatically on system startup

Confluence can be configured to start automatically on system startup, allowing it to recover automatically after a reboot.

Start Confluence automatically on OS X using launchd

Start Confluence automatically on Linux & Unix

Start Confluence automatically on Windows as a Service

Start Confluence automatically on Linux & Unix

On Unix/Linux, 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 Unix/Linux:

1. Create a `confluence` user for instance, using the following command:

```
sudo useradd --create-home -c "Confluence role account" confluence
```

2. Create a directory to install Confluence into:

```
sudo mkdir /usr/local/confluence
sudo chown confluence: /usr/local/confluence
```

3. Log in as the `confluence` user to install Confluence:

```
sudo su - confluence
cd /usr/local/confluence/
tar zxvf /tmp/confluence-3.0.1-std.tar.gz
ln -s confluence-3.0.1-std/ current
```

4. Edit current/confluence/WEB-INF/classes/confluence-init.properties, and set confluence.home=/usr/local/confluence/home
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.

```
01.#!/bin/sh -e
02.# Confluence startup script
03.#chkconfig: 2345 80 05
04.#description: Confluence
05.
06.# Define some variables
07.# Name of app ( JIRA, Confluence, etc )
08.APP=confluence
09.# Name of the user to run as
10.USER=confluence
11.# Location of application's bin directory
12.CATALINA_HOME=/usr/local/confluence/current
13.# Location of Java JDK
14.export JAVA_HOME=/usr/lib/jvm/java-6-sun
15.
16.case "$1" in
17.     # Start command
18.     start)
19.          echo "Starting $APP"
20.          /bin/su -m $USER -c "$CATALINA_HOME/bin/startup.sh &> /dev/null"
21.          ;;
22.     # Stop command
23.     stop)
24.          echo "Stopping $APP"
25.          /bin/su -m $USER -c "$CATALINA_HOME/bin/shutdown.sh &> /dev/null"
26.          echo "$APP stopped successfully"
27.          ;;
28.        # Restart command
29.        restart)
30.                     $0 stop
31.                     sleep 5
32.                     $0 start
33.                     ;;
34.     *)
35.          echo "Usage: /etc/init.d/$APP {start|restart|stop}"
36.          exit 1
37.          ;;
38.esac
39.
40.exit 0
```

6. Make this file executable:

```
sudo chmod +x /etc/init.d/confluence
```

7. Set this file to run at the appropriate runleve. 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 Standalone, please see the Installation guide
at http://confluence.atlassian.com/display/DOC/Confluence+Installation+Guide
Using CATALINA_BASE:   /usr/local/confluence/current
Using CATALINA_HOME:   /usr/local/confluence/current
Using CATALINA_TMPDIR: /usr/local/confluence/current/temp
Using JRE_HOME:       /usr/lib/jvm/java-1.5.0-sun
done.
```

You should then see this running at http://<server>:8080/.

Adding Confluence as a service for Ubuntu Jaunty (or later)

To continue configuring Confluence to start automatically as a service on Ubuntu Jaunty (or later):

1. After logging in as the `confluence` user to install Confluence, create `start` and `stop` scripts in `/usr/local/confluence`:

Example `start` script:

```
#!/bin/bash
export JAVA_HOME=/usr/lib/jvm/java-6-sun-1.6.0.16/
export JDK_HOME=/usr/lib/jvm/java-6-sun-1.6.0.16/
cd /usr/local/confluence/current/bin
./startup.sh
```

Example `stop` script:

```
#!/bin/bash
export JAVA_HOME=/usr/lib/jvm/java-6-sun-1.6.0.16/
export JDK_HOME=/usr/lib/jvm/java-6-sun-1.6.0.16/
cd /usr/local/confluence/current/bin
./shutdown.sh
```

2. Make both of these scripts executable. For example, by issuing the command: `sudo chmod a+x /usr/local/confluence/start /usr/local/confluence/stop`.
3. Karmic and later: Create two text files in `/etc/init/` called `confluence-up.conf` and `confluence-down.conf`:

`confluence-up`:

```
start on runlevel [2345]

script

 date >> /tmp/confluence-startup.out
 exec sudo -u confluence /usr/local/confluence/start >> /tmp/confluence-startup.out 2>&1

end script
```

`confluence-down`:

```
start on runlevel [16]

expect fork
respawn

exec sudo -u confluence /usr/local/confluence/stop >> /tmp/confluence-shutdown.out 2>&1
```

... and make them readable to all users:
`sudo chmod a+r /etc/init/confluence-up.conf /etc/init/confluence-down.conf`

1. Jaunty, Intrepid: Create two text files in `/etc/event.d/` called `confluence-up` and `confluence-down`:

`confluence-up`:

```
start on runlevel 2
start on runlevel 3
start on runlevel 4
start on runlevel 5

exec sudo -u confluence /usr/local/confluence/start >> /tmp/confluence-startup.out 2>&1
```

`confluence-down`:

```
start on runlevel 1
start on runlevel 6

exec sudo -u confluence /usr/local/confluence/stop >> /tmp/confluence-shutdown.out 2>&1
```

... and make them readable to all users:
`sudo chmod a+r /etc/event.d/confluence-up /etc/event.d/confluence-down`

RELATED TOPICS

Start Confluence automatically on system startup

Start Confluence automatically on OS X using launchd

For long-term use, you should configure Confluence to restart automatically when the operating system restarts. On Mac OS X, the system startup program called launchd manages long running processes – daemons or services.

Apple provides an introduction to launchd. Below we tell you how to use launchd to start Confluence automatically on Mac OS X when running Tomcat.

On this page:

- Using launchd with Tomcat
  - Step 1. Add a Wrapper Shell Script
  - Step 2. Add a launchd Property List
  - Starting and Stopping Confluence Manually
  - Troubleshooting

Using launchd with Tomcat

The Confluence standalone distribution ships with Tomcat. There is a mismatch between how launchd expects a daemon to behave, and how the default startup scripts for Tomcat operate:

- OS X's launchd expects the process it starts to run forever, but '`catalina.sh start`' starts the JVM to run Tomcat and then exits.
- Tomcat provides '`catalina.sh stop`' to shut down Tomcat cleanly by connecting to a socket which Tomcat listens on, but launchd stops daemons by sending them a signal that kills the process immediately if no specific handling is included.

You will need a wrapper shell script and properties list to make launchd work with Tomcat.

Step 1. Add a Wrapper Shell Script

Add the following wrapper shell script to `$CATALINA_HOME/bin`:

launchd_wrapper.sh

```
01.#!/bin/bash
02.
03.function shutdown()
04.{
05.                    date
06.                    echo "Shutting down Confluence"
07.                    $CATALINA_HOME/bin/catalina.sh stop
08.}
09.
10.date
11.echo "Starting Confluence"
12.export CATALINA_PID=/tmp/$$
13.
14.# Uncomment to increase Tomcat's maximum heap allocation
15.# export JAVA_OPTS=-Xmx512M $JAVA_OPTS
16.
17.. $CATALINA_HOME/bin/catalina.sh start
18.
19.# Allow any signal which would kill a process to stop Tomcat
20.trap shutdown HUP INT QUIT ABRT KILL ALRM TERM TSTP
21.
22.echo "Waiting for `cat $CATALINA_PID`"
23.wait `cat $CATALINA_PID`
```

The above shell script starts Tomcat and then waits for the process to complete, so launchd is happy that Tomcat is still running. The script also installs a signal handler, which calls the `shutdown()` function to cleanly shut down Tomcat when launchd signals the script.

You can try this script manually: Start the script, watch Confluence start, and then type **ctrl-C** and see Confluence shut down cleanly. (Note that it will **not** shut down cleanly if Tomcat has not started yet. It takes a few seconds for Tomcat to start listening on the shutdown socket.)

Step 2. Add a launchd Property List

The launchd property list (`.plist`) tells launchd how to start Tomcat.

Add the following plist file to `/Library/LaunchDaemons`, which is the location for system-wide services which are not part of base OS X:

confluence.plist

```
01.<?xml version="1.0" encoding="UTF-8"?>
02.<!DOCTYPE plist PUBLIC "-//Apple Computer//DTD PLIST 1.0//EN" "
http://www.apple.com/DTDs/PropertyList-1.0.dtd">
03.<plist version="1.0">
04.<dict>
05.                     <key>Disabled</key>
06.                     <false/>
07.                     <key>EnvironmentVariables</key>
08.                     <dict>
09.                                         <key>CATALINA_HOME</key>
10.                                         <string>/Users/myname/conf/confluence-2.1.3-std</string>
11.                                         <key>JAVA_HOME</key>
12.                                         <string>/Library/Java/Home</string>
13.                     </dict>
14.                     <key>Label</key>
15.                     <string>com.atlassian.confluence</string>
16.                     <key>OnDemand</key>
17.                     <false/>
18.                     <key>ProgramArguments</key>
19.                     <array>
20.                                         <string
>/Users/myname/conf/confluence-2.1.3-std/bin/launchd_wrapper.sh</string>
21.                     </array>
22.                     <key>RunAtLoad</key>
23.                     <true/>
24.                     <key>ServiceDescription</key>
25.                     <string>Confluence</string>
26.                     <key>StandardErrorPath</key>
27.                     <string>/Users/myname/conf/confluence-2.1.3-std/logs/launchd.stderr</string>
28.                     <key>StandardOutPath</key>
29.                     <string>/Users/myname/conf/confluence-2.1.3-std/logs/launchd.stdout</string>
30.                     <key>UserName</key>
31.                     <string>root</string>
32.</dict>
33.</plist>
```

Notes:

1. Replace '`/Users/myname/conf/confluence-2.1.3-std`' with the path to your Confluence installation. The string occurs four times in the above script.
2. JAVA_HOME is set to use the default JDK. On OS X version 10.4.4, the default JDK is 1.4.2. You will need to change this value if you want to use a different version of Java. For example, if you want to use JDK 1.5, you will need to change JAVA_HOME to `/System/Library/Frameworks/JavaVM.framework/Versions/1.5`.
3. In the above script, we have specified '`root`' as the `UserName`. If necessary, change the UserName to the user you want Tomcat to run as.

Starting and Stopping Confluence Manually

To start and stop Confluence manually, use the following commands:

- Start:
  ```
  cd /Library/LaunchDaemons
  sudo launchctl load -w confluence.plist
  ```
- Stop:
  ```
  cd /Library/LaunchDaemons
  sudo launchctl unload -w confluence.plist
  ```

Troubleshooting

- Make sure both files `launch_wrapper.sh` and `confluence.plist` have the necessary file privileges.
- Check the console logging and log file for any abnormalities.

RELATED TOPICS

Start Confluence automatically on system startup

Start Confluence automatically on Windows as a Service

For long-term use, we recommend that you configure Confluence to start automatically when the operating system restarts. For Windows servers, this means configuring Confluence to run as a Windows service.

There are two ways to install Confluence Standalone as a service: via the Confluence installer or manually as described below.

On this page:

- Reasons for Starting Confluence as a Service
- Manually Installing Confluence Standalone as a Service
- Managing Confluence as a Service
- Upgrading Confluence
- Troubleshooting Confluence while Running as a Windows Service
- Requesting Support

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.

Manually Installing Confluence Standalone as a Service

From your Windows-based server:

1. Open a command prompt in the `<CONFLUENCE-INSTALL>/bin` directory.
2. Confirm that the JAVA_HOME variable is set to the JDK base directory with the command:

   ```
   1.echo %JAVA_HOME%
   ```

   Note that any directory in the path with spaces (eg. `C:\Program Files` must be converted to its eight-character equivalent (e.g. `C:\Progra~1`).

3. If you are installing Confluence on a Windows 2008 server, be sure to run the command prompt using 'run as administrator'. (Otherwise running 'service.bat', as described in the next step, will fail.)
4. Use the following command to install the service with default settings:

   ```
   1.service.bat install Confluence
   ```

5. Now, to have the service start automatically when the server starts, run:

   ```
   tomcat6 //US//Confluence --Startup auto
   ```

6. If you have a less than a 512 megabytes of memory, skip this step. For users with large Confluence installations, you can increase the maximum memory Confluence can use. (The default is 256MB). For example, you can set the maximum memory to 512 megs using:

   ```
   tomcat6 //US//Confluence --JvmMx 512
   ```

7. If you do not have any JVM parameters you pass to your standalone distribution of Confluence, you can skip this step. If you do, add them to the service using:

   ```
   tomcat6 //US//Confluence ++JvmOptions="-Djust.an.example=True"
   ```

8. For further configuration options, please refer to the Tomcat Windows Service How-To guide
9. Go to your Windows Control Panel -> Administrative Tools -> Services -> Apache Tomcat Confluence and right-click on Properties to verify the settings are correct.
10. 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 `%CATALINA_HOME%`, index and database directories. Note that any network drives must be specified by UNC and not letter mappings (eg. `\\backupserver\confluence` not `z:\confluence`)

    Confluence is now installed as a service, but will not automatically start up until the next server reboot

11. Start the Confluence service with the command:

```
1.net start Confluence
```

Managing Confluence as a Service

You can manage the Confluence service from the command prompt.

- Stop Confluence with:

```
1.net stop Confluence
```

- Uninstall the Confluence service with:

```
1.service.bat remove Confluence
```

Upgrading Confluence

After upgrading Confluence, you can either uninstall and reinstall the Windows service or change the StartPath parameter to your new folder. Refer to the Tomcat documentation for help.

Troubleshooting Confluence while Running as a Windows Service

- Check the Knowledge Base articles:
  - Confluence Standalone Unable to Run as Windows Service in Windows 64-Bit
  - Confluence does not Start due to Windows Firewall
  - Unable to Install Service on Windows Vista
  - Unable to configure Confluence to run as a service on Tomcat
  - Could not Start Confluence as a Service After Allocating JVM Memory
- If none of the above solves your problem, please refer to the complete list of known issues in our Knowledge Base.
- When investigating memory issues or bugs, it may be useful to view information from Confluence's garbage collection. To turn on the verbose garbage collection, use the command:

```
tomcat6 //US//Confluence ++JvmOptions="-Xloggc:<CONFLUENCE-INSTALL>\logs\atlassian-gc.log"
```

- The Confluence 2.9 installer does not work when installed as service, due to a missing semi-colon in `service.bat`. Please refer to reported issue CONF-12785.
- You can use a Sysinternals tool called **`Procmon.exe`** from the The Microsoft Windows Sysinternals Team, to check that the error occurred at the specific time when the Confluence service started. You need to match the time when Tomcat failed, as captured by this tool, against the time in the Windows Event Viewer.

> **Note**
> We do not recommend that you run this tool for too long as it may disrupt other Atlassian applications. Once you have captured the required information you will need to press `Ctrl + E` to stop capturing.

Requesting Support

If, after following the troubleshooting guide above, you still cannot make Confluence run as a Windows Service or if there is an error when setting the JVM configuration for the service, you can create a support request at http://support.atlassian.com.

Please provide the following information when creating your support request, because we will need it to assist you:

- Are you running a 32 bit or 64 bit Windows?
- Give us the result of running **`java -version`** from Windows command line console.
- A screen shot of your Windows Registry setting for Tomcat.
- If you have modified **service.bat**, please give us a copy of this file for review.
- What application server are you using? eg. Are you using the Confluence Standalone distribution?
- Your atlassian-confluence.log file.

RELATED TOPICS

Start Confluence automatically on system startup
Fix Out of Memory errors by Increasing Available Memory
Editing the Windows Registry

The Confluence data model

On this page:

- General Database Diagram
- Authentication
  - Atlassian-user
  - OpenSymphony
- Content
- Clustering
  - System information
  - Spaces
- Appearance
- Miscellaneous



The Hibernate mapping files are the authoritative reference. These are the *.hbm.xml files which have been bundled into the main Confluence .jar file in recent releases.

This document is little more than the Confluence schema with added comments, but the priority was placed on making the information available.

General Database Diagram



Authentication

Atlassian-user

This is the "new" authentication system, which is more flexible and extensible than `OpenSymphony`.

```
Table "groups"
  Column   |          Type          | Modifiers
-----------+------------------------+-----------
 id        | bigint                 | not null
 groupname | character varying(255) | not null
Indexes:
    "groups_pkey" PRIMARY KEY, btree (id)
```

```
Table "users"
  Column  |            Type             | Modifiers
----------+-----------------------------+-----------
 id       | bigint                      | not null
 name     | character varying(255)      | not null
 password | character varying(255)      |
 email    | character varying(255)      |
 created  | timestamp without time zone |
 fullname | character varying(255)      |
Indexes:
    "users_pkey" PRIMARY KEY, btree (id)
    "users_name_key" UNIQUE, btree (name)
```

`local_members`: establishes many-to-many association between users and groups.

```
Table "local_members"
 Column  |  Type  | Modifiers
---------+--------+-----------
 userid  | bigint | not null
 groupid | bigint | not null
Indexes:
    "local_members_pkey" PRIMARY KEY, btree (groupid, userid)
Foreign-key constraints:
    "fk6b8fb445117d5fda" FOREIGN KEY (groupid) REFERENCES groups(id)
    "fk6b8fb445ce2b3226" FOREIGN KEY (userid) REFERENCES users(id)
```

`external_entities`: Maps users from LDAP (or any other external authentication system) to IDs in Confluence DB

```
Table "external_entities"
 Column |          Type          | Modifiers
--------+------------------------+-----------
 id     | bigint                 | not null
 name   | character varying(255) |
 type   | character varying(255) | not null
Indexes:
    "external_entities_pkey" PRIMARY KEY, btree (id)
```

`external_members`: associates LDAP (or other external) users with local groups.

```
Table "external_members"
   Column    |  Type  | Modifiers
-------------+--------+-----------
 extentityid | bigint | not null
 groupid     | bigint | not null
Indexes:
    "external_members_pkey" PRIMARY KEY, btree (groupid, extentityid)
Foreign-key constraints:
    "fkd8c8d8a5117d5fda" FOREIGN KEY (groupid) REFERENCES groups(id)
    "fkd8c8d8a5f25e5d5f" FOREIGN KEY (extentityid) REFERENCES external_entities(id)
```

OpenSymphony

The "old" authentication system, which was the default prior to 2.7.

```
Table "os_group"
  Column   |          Type          | Modifiers
-----------+------------------------+-----------
 id        | bigint                 | not null
 groupname | character varying(255) | not null
Indexes:
    "os_group_pkey" PRIMARY KEY, btree (id)
    "os_group_groupname_key" UNIQUE, btree (groupname)
```

```
Table "os_user"
  Column  |          Type          | Modifiers
----------+------------------------+-----------
 id       | bigint                 | not null
 username | character varying(255) | not null
 passwd   | character varying(255) |
Indexes:
    "os_user_pkey" PRIMARY KEY, btree (id)
    "os_user_username_key" UNIQUE, btree (username)
```

```
Table "os_user_group"
  Column   |  Type   | Modifiers
----------+--------+-----------
 group_id | bigint | not null
 user_id  | bigint | not null
Indexes:
    "os_user_group_pkey" PRIMARY KEY, btree (user_id, group_id)
Foreign-key constraints:
    "fk932472461e2e76db" FOREIGN KEY (group_id) REFERENCES os_group(id)
    "fk93247246f73aee0f" FOREIGN KEY (user_id) REFERENCES os_user(id)
```

Content

The actual information that users are storing and sharing.

`attachmentdata`: stores the binary data for attached files.
Only used when Confluence is configured to store attachments in the database; otherwise, attachments are stored in the local filesystem.

```
Table "attachmentdata"
      Column      |  Type   | Modifiers
------------------+---------+-----------
 attachmentdataid | bigint  | not null
 attversion       | integer | not null
 data             | bytea   |
 attachmentid     | bigint  |
Indexes:
    "attachmentdata_pkey" PRIMARY KEY, btree (attachmentdataid)
    "attch_data_idx" btree (attachmentid)
Foreign-key constraints:
    "fk9dc3e34d34a4917e" FOREIGN KEY (attachmentid) REFERENCES attachments(attachmentid)
```

`attachments`: metadata for attachments.

```
Table "attachments"
       Column       |            Type             | Modifiers
--------------------+-----------------------------+-----------
 attachmentid       | bigint                      | not null
 title              | character varying(255)      | not null
 contenttype        | character varying(255)      | not null
 pageid             | bigint                      | not null
 creator            | character varying(255)      |
 creationdate       | timestamp without time zone |
 lastmodifier       | character varying(255)      |
 lastmoddate        | timestamp without time zone |
 filesize           | bigint                      |
 attachment_comment | character varying(255)      |
 attversion         | integer                     |
 prevver            | bigint                      |
Indexes:
    "attachments_pkey" PRIMARY KEY, btree (attachmentid)
    "att_pageid_idx" btree (pageid)
    "att_prevver_idx" btree (prevver)
Foreign-key constraints:
    "fk54475f9017d4a070" FOREIGN KEY (prevver) REFERENCES attachments(attachmentid)
    "fk54475f908c38fbea" FOREIGN KEY (pageid) REFERENCES content(contentid)
```

`bodycontent`: stores the actual content of Confluence pages. No versioning information or other metadata is stored here, though; that's all in the `content` table.

```
Table "bodycontent"
    Column     |  Type  | Modifiers
---------------+--------+-----------
 bodycontentid | bigint | not null
 body          | text   |
 contentid     | bigint |
Indexes:
    "bodycontent_pkey" PRIMARY KEY, btree (bodycontentid)
    "body_content_idx" btree (contentid)
Foreign-key constraints:
    "fka898d4778dd41734" FOREIGN KEY (contentid) REFERENCES content(contentid)
```

`content`: a persistence table for the ContentEntityObject class of objects. The subclass is indicated by the `contenttype` column.

```
Table "content"
      Column      |            Type             | Modifiers
------------------+-----------------------------+-----------
 contentid        | bigint                      | not null
 contenttype      | character varying(255)      | not null
 title            | character varying(255)      |
 version          | integer                     |
 creator          | character varying(255)      |
 creationdate     | timestamp without time zone |
 lastmodifier     | character varying(255)      |
 lastmoddate      | timestamp without time zone |
 versioncomment   | text                        |
 prevver          | bigint                      |
 content_status   | character varying(255)      |
 spaceid          | bigint                      |
 parentid         | bigint                      |
 messageid        | character varying(255)      |
 draftpageid      | character varying(255)      |
 draftspacekey    | character varying(255)      |
 drafttype        | character varying(255)      |
 draftpageversion | integer                     |
 pageid           | bigint                      |
 parentcommentid  | bigint                      |
 username         | character varying(255)      |
Indexes:
    "content_pkey" PRIMARY KEY, btree (contentid)
    "c_draftpageid_idx" btree (draftpageid)
    "c_draftspacekey_idx" btree (draftspacekey)
    "c_drafttype_idx" btree (drafttype)
    "c_messageid_idx" btree (messageid)
    "c_parentcommid_idx" btree (parentcommentid)
    "c_parentid_idx" btree (parentid)
    "c_prevver_idx" btree (prevver)
    "c_spaceid_idx" btree (spaceid)
    "c_title_idx" btree (title)
    "c_username_idx" btree (username)
Foreign-key constraints:
    "fk6382c05917d4a070" FOREIGN KEY (prevver) REFERENCES content(contentid)
    "fk6382c05974b18345" FOREIGN KEY (parentid) REFERENCES content(contentid)
    "fk6382c0598c38fbea" FOREIGN KEY (pageid) REFERENCES content(contentid)
    "fk6382c059b2dc6081" FOREIGN KEY (spaceid) REFERENCES spaces(spaceid)
    "fk6382c059b97e9230" FOREIGN KEY (parentcommentid) REFERENCES content(contentid)
```

`content_label`: Arbitrary text labels for content.

```
Table "content_label"
    Column    |            Type             | Modifiers
--------------+-----------------------------+-----------
 id           | bigint                      | not null
 labelid      | bigint                      | not null
 contentid    | bigint                      | not null
 spacekey     | character varying(255)      |
 owner        | character varying(255)      |
 creationdate | timestamp without time zone |
 lastmoddate  | timestamp without time zone |
Indexes:
    "content_label_pkey" PRIMARY KEY, btree (id)
    "cl_contentid_idx" btree (contentid)
    "cl_labelid_idx" btree (labelid)
    "cl_lastmoddate_idx" btree (lastmoddate)
    "cl_spacekey_idx" btree (spacekey)
Foreign-key constraints:
    "fkf0e7436e27072aef" FOREIGN KEY (labelid) REFERENCES label(labelid)
    "fkf0e7436e8dd41734" FOREIGN KEY (contentid) REFERENCES content(contentid)
```

`label`: the other half of the `content_label` system.

```
Table "label"
    Column    |            Type             | Modifiers
--------------+-----------------------------+-----------
 labelid      | bigint                      | not null
 name         | character varying(255)      |
 owner        | character varying(255)      |
 namespace    | character varying(255)      |
 creationdate | timestamp without time zone |
 lastmoddate  | timestamp without time zone |
Indexes:
    "label_pkey" PRIMARY KEY, btree (labelid)
    "l_name_idx" btree (name)
    "l_namespace_idx" btree (namespace)
    "l_owner_idx" btree ("owner")
```

`content_perm`: content-level permissions objects.

```
Table "content_perm"
    Column    |            Type             | Modifiers
--------------+-----------------------------+-----------
 id           | bigint                      | not null
 cp_type      | character varying(10)       | not null
 username     | character varying(255)      |
 groupname    | character varying(255)      |
 cps_id       | bigint                      |
 creator      | character varying(255)      |
 creationdate | timestamp without time zone |
 lastmodifier | character varying(255)      |
 lastmoddate  | timestamp without time zone |
Indexes:
    "content_perm_pkey" PRIMARY KEY, btree (id)
    "cp_gn_idx" btree (groupname)
    "cp_os_idx" btree (cps_id)
    "cp_un_idx" btree (username)
Foreign-key constraints:
    "fkbd74b31676e33274" FOREIGN KEY (cps_id) REFERENCES content_perm_set(id)
```

`content_perm_set`: one-to-many mapping for content items and their permissions, with added metadata.

```
Table "content_perm_set"
     Column     |            Type             | Modifiers
----------------+-----------------------------+-----------
 id             | bigint                      | not null
 cont_perm_type | character varying(10)       | not null
 content_id     | bigint                      |
 creationdate   | timestamp without time zone |
 lastmoddate    | timestamp without time zone |
Indexes:
    "content_perm_set_pkey" PRIMARY KEY, btree (id)
    "cps_content_idx" btree (content_id)
Foreign-key constraints:
    "fkbf45a7992caf22c1" FOREIGN KEY (content_id) REFERENCES content(contentid)
```

Clustering

`clustersafety`: normally, this table only contains one row. The value of the `safetynumber` is what Confluence uses to find out whether another instance is sharing its database without being part of the cluster.

```
Table "clustersafety"
     Column      |  Type   | Modifiers
-----------------+---------+-----------
 clustersafetyid | bigint  | not null
 safetynumber    | integer |
Indexes:
    "clustersafety_pkey" PRIMARY KEY, btree (clustersafetyid)
```

System information

`confversion` used by the upgrade system to determine what to expect from the database, so as to negotiate upgrades.

```
Table "confversion"
    Column     |            Type             | Modifiers
---------------+-----------------------------+-----------
 confversionid | bigint                      | not null
 buildnumber   | integer                     | not null
 installdate   | timestamp without time zone |
 versiontag    | character varying(255)      |
 creationdate  | timestamp without time zone |
 lastmoddate   | timestamp without time zone |
Indexes:
    "confversion_pkey" PRIMARY KEY, btree (confversionid)
    "confversion_buildnumber_key" UNIQUE, btree (buildnumber)
```

`plugindata`: records which plugins have been installed, and when.
`data` is a blob of the actual plugin .jar file. This is principally cluster-related.

```
Table "plugindata"
    Column    |            Type             | Modifiers
--------------+-----------------------------+-----------
 plugindataid | bigint                      | not null
 pluginkey    | character varying(255)      | not null
 filename     | character varying(255)      | not null
 lastmoddate  | timestamp without time zone |
 data         | bytea                       |
Indexes:
    "plugindata_pkey" PRIMARY KEY, btree (plugindataid)
    "plugindata_filename_key" UNIQUE, btree (filename)
    "plugindata_pluginkey_key" UNIQUE, btree (pluginkey)
```

Spaces

`spacegroups`: this table is only used by the hosted environment.

```
Table "spacegroups"
     Column     |            Type             | Modifiers
----------------+-----------------------------+-----------
 spacegroupid   | bigint                      | not null
 spacegroupname | character varying(255)      |
 spacegroupkey  | character varying(255)      | not null
 licensekey     | text                        |
 creator        | character varying(255)      |
 creationdate   | timestamp without time zone |
 lastmodifier   | character varying(255)      |
 lastmoddate    | timestamp without time zone |
Indexes:
    "spacegroups_pkey" PRIMARY KEY, btree (spacegroupid)
    "spacegroups_spacegroupkey_key" UNIQUE, btree (spacegroupkey)
```

```
Table "spacepermissions"
    Column     |            Type             | Modifiers
---------------+-----------------------------+-----------
 permid        | bigint                      | not null
 spaceid       | bigint                      |
 permtype      | character varying(255)      | not null
 permgroupname | character varying(255)      |
 permusername  | character varying(255)      |
 creator       | character varying(255)      |
 creationdate  | timestamp without time zone |
 lastmodifier  | character varying(255)      |
 lastmoddate   | timestamp without time zone |
Indexes:
    "spacepermissions_pkey" PRIMARY KEY, btree (permid)
    "sp_permtype_idx" btree (permtype)
    "sp_pgname_idx" btree (permgroupname)
    "sp_puname_idx" btree (permusername)
    "sp_spaceid_idx" btree (spaceid)
Foreign-key constraints:
    "fkd33f23beb2dc6081" FOREIGN KEY (spaceid) REFERENCES spaces(spaceid)
```

`spaces`: information about the spaces themselves: key, human-friendly name and numeric ID.

```
Table "spaces"
    Column    |            Type             | Modifiers
--------------+-----------------------------+-----------
 spaceid      | bigint                      | not null
 spacename    | character varying(255)      |
 spacekey     | character varying(255)      | not null
 spacedescid  | bigint                      |
 homepage     | bigint                      |
 creator      | character varying(255)      |
 creationdate | timestamp without time zone |
 lastmodifier | character varying(255)      |
 lastmoddate  | timestamp without time zone |
 spacetype    | character varying(255)      |
 spacegroupid | bigint                      |
Indexes:
    "spaces_pkey" PRIMARY KEY, btree (spaceid)
    "spaces_spacekey_key" UNIQUE, btree (spacekey)
    "s_homepage_idx" btree (homepage)
    "s_spacedescid_idx" btree (spacedescid)
    "s_spacegroupid_idx" btree (spacegroupid)
Foreign-key constraints:
    "fk9228242d11b7bfee" FOREIGN KEY (homepage) REFERENCES content(contentid)
    "fk9228242d16994414" FOREIGN KEY (spacegroupid) REFERENCES spacegroups(spacegroupid)
    "fk9228242d2c72d3d2" FOREIGN KEY (spacedescid) REFERENCES content(contentid)
```

Appearance

`decorator`: storage of custom display templates, for customising layouts.

```
Table "decorator"
    Column     |            Type             | Modifiers
---------------+-----------------------------+-----------
 decoratorid   | bigint                      | not null
 spacekey      | character varying(255)      |
 decoratorname | character varying(255)      |
 body          | text                        |
 lastmoddate   | timestamp without time zone |
Indexes:
    "decorator_pkey" PRIMARY KEY, btree (decoratorid)
    "dec_key_idx" btree (spacekey)
    "dec_name_idx" btree (decoratorname)
```

Miscellaneous

`os_propertyentry`: for arbitrary association of entities and properties.

```
Table "os_propertyentry"
   Column    |            Type             | Modifiers
-------------+-----------------------------+-----------
 entity_name | character varying(125)      | not null
 entity_id   | bigint                      | not null
 entity_key  | character varying(200)      | not null
 key_type    | integer                     |
 boolean_val | boolean                     |
 double_val  | double precision            |
 string_val  | character varying(255)      |
 text_val    | text                        |
 long_val    | bigint                      |
 int_val     | integer                     |
 date_val    | timestamp without time zone |
Indexes:
    "os_propertyentry_pkey" PRIMARY KEY, btree (entity_name, entity_id, entity_key)
```

`bandana`: a catch-all persistence layer. It 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.

```
Table "bandana"
     Column     |          Type          | Modifiers
----------------+------------------------+-----------
 bandanaid      | bigint                 | not null
 bandanacontext | character varying(255) |
 bandanakey     | character varying(100) |
 bandanavalue   | text                   |
Indexes:
    "bandana_pkey" PRIMARY KEY, btree (bandanaid)
    "band_context_idx" btree (bandanacontext)
    "band_key_idx" btree (bandanakey)
```

`extrnlnks`: storage of referral links.

```
Table "extrnlnks"
    Column    |            Type             | Modifiers
--------------+-----------------------------+-----------
 linkid       | bigint                      | not null
 contenttype  | character varying(255)      | not null
 viewcount    | integer                     | not null
 url          | character varying(255)      | not null
 contentid    | bigint                      | not null
 creator      | character varying(255)      |
 creationdate | timestamp without time zone |
 lastmodifier | character varying(255)      |
 lastmoddate  | timestamp without time zone |
Indexes:
    "extrnlnks_pkey" PRIMARY KEY, btree (linkid)
    "el_contentid_idx" btree (contentid)
Foreign-key constraints:
    "fk97c10fe78dd41734" FOREIGN KEY (contentid) REFERENCES content(contentid)
```

`hibernate_unique_key`: used by the high/low ID generator - the subsystem which generates our primary keys.
Mess with this at the cost of being able to create objects.

```
Table "hibernate_unique_key"
 Column  |  Type   | Modifiers
---------+---------+-----------
 next_hi | integer |
```

`indexqueueentries`: arbitrates full-content indexing across the system.
This table generally contains the last 12 hours or so of updates, to allow re-syncing of cluster nodes after restarts.

```
Table "indexqueueentries"
    Column    |            Type             | Modifiers
--------------+-----------------------------+-----------
 entryid      | bigint                      | not null
 creationdate | timestamp without time zone |
 type         | integer                     |
 handle       | character varying(255)      |
Indexes:
    "indexqueueentries_pkey" PRIMARY KEY, btree (entryid)
```

`keystore`: used by the trusted apps framework to store the server's private key, and other servers' public keys.

```
Table "keystore"
  Column   |          Type          | Modifiers
-----------+------------------------+-----------
 keyid     | bigint                 | not null
 alias     | character varying(255) | not null
 type      | character varying(32)  | not null
 algorithm | character varying(32)  | not null
 keyspec   | text                   | not null
Indexes:
    "keystore_pkey" PRIMARY KEY, btree (keyid)
```

`links`: tracks links within the server (i.e. across and within spaces).

```
Table "links"
    Column     |            Type             | Modifiers
---------------+-----------------------------+-----------
 linkid        | bigint                      | not null
 destpagetitle | character varying(255)      |
 destspacekey  | character varying(255)      | not null
 contentid     | bigint                      | not null
 creator       | character varying(255)      |
 creationdate  | timestamp without time zone |
 lastmodifier  | character varying(255)      |
 lastmoddate   | timestamp without time zone |
Indexes:
    "links_pkey" PRIMARY KEY, btree (linkid)
    "l_contentid_idx" btree (contentid)
    "l_destspacekey_idx" btree (destspacekey)
Foreign-key constraints:
    "fk45157998dd41734" FOREIGN KEY (contentid) REFERENCES content(contentid)
```

`notifications`: storage of page- and space-level watches.

```
Table "notifications"
     Column     |            Type             | Modifiers
----------------+-----------------------------+-----------
 notificationid | bigint                      | not null
 pageid         | bigint                      |
 spaceid        | bigint                      |
 username       | character varying(255)      | not null
 creator        | character varying(255)      |
 creationdate   | timestamp without time zone |
 lastmodifier   | character varying(255)      |
 lastmoddate    | timestamp without time zone |
Indexes:
    "notifications_pkey" PRIMARY KEY, btree (notificationid)
    "n_pageid_idx" btree (pageid)
    "n_spaceid_idx" btree (spaceid)
Foreign-key constraints:
    "fk594acc88c38fbea" FOREIGN KEY (pageid) REFERENCES content(contentid)
    "fk594acc8b2dc6081" FOREIGN KEY (spaceid) REFERENCES spaces(spaceid)
```

`pagetemplates`: acts as the back-end of the templates feature.

```
Table "pagetemplates"
    Column    |            Type             | Modifiers
--------------+-----------------------------+-----------
 templateid   | bigint                      | not null
 templatename | character varying(255)      | not null
 templatedesc | character varying(255)      |
 labels       | character varying(255)      |
 content      | text                        |
 spaceid      | bigint                      |
 prevver      | bigint                      |
 version      | integer                     | not null
 creator      | character varying(255)      |
 creationdate | timestamp without time zone |
 lastmodifier | character varying(255)      |
 lastmoddate  | timestamp without time zone |
Indexes:
    "pagetemplates_pkey" PRIMARY KEY, btree (templateid)
    "pt_prevver_idx" btree (prevver)
    "pt_spaceid_idx" btree (spaceid)
Foreign-key constraints:
    "fkbc7ce96a17d4a070" FOREIGN KEY (prevver) REFERENCES pagetemplates(templateid)
    "fkbc7ce96ab2dc6081" FOREIGN KEY (spaceid) REFERENCES spaces(spaceid)
```

```
Table "trackbacklinks"
    Column    |            Type             | Modifiers
--------------+-----------------------------+-----------
 linkid       | bigint                      | not null
 contenttype  | character varying(255)      | not null
 viewcount    | integer                     | not null
 url          | character varying(255)      | not null
 title        | character varying(255)      |
 blogname     | character varying(255)      |
 excerpt      | character varying(255)      |
 contentid    | bigint                      | not null
 creator      | character varying(255)      |
 creationdate | timestamp without time zone |
 lastmodifier | character varying(255)      |
 lastmoddate  | timestamp without time zone |
Indexes:
    "trackbacklinks_pkey" PRIMARY KEY, btree (linkid)
    "tbl_contentid_idx" btree (contentid)
Foreign-key constraints:
    "fkf6977a478dd41734" FOREIGN KEY (contentid) REFERENCES content(contentid)
```

`confancestors`: used to speed up permissions checks, by allowing quick lookup of all a page's ancestors.

```
Table "confancestors"
      Column      |  Type   | Modifiers
------------------+---------+-----------
 descendentid     | bigint  | not null
 ancestorid       | bigint  | not null
 ancestorposition | integer | not null
Indexes:
    "confancestors_pkey" PRIMARY KEY, btree (descendentid, ancestorposition)
Foreign-key constraints:
    "fk9494e23c37e35a2e" FOREIGN KEY (ancestorid) REFERENCES content(contentid)
    "fk9494e23cc45e94dc" FOREIGN KEY (descendentid) REFERENCES content(contentid)
```

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:

```
1.org.springframework.dao.DataIntegrityViolationException: (HibernateTemplate): data integrity
violated by SQL ''; nested exception is java.sql.BatchUpdateException: Duplicate entry '1234' for
key 1
2.   at
org.springframework.jdbc.support.SQLStateSQLExceptionTranslator.translate(SQLStateSQLExceptionTransla
88)
3.caused by: java.sql.BatchUpdateException: Duplicate entry '1234' for key 1
4.   at com.mysql.jdbc.ServerPreparedStatement.executeBatch(ServerPreparedStatement.java:647)
```

or

```
1.(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.

   ```
   1.## log hibernate prepared statements/SQL queries (equivalent to setting
   'hibernate.show_sql' to 'true')
   2.#log4j.logger.net.sf.hibernate.SQL=DEBUG, confluencelog
   3.#log4j.additivity.net.sf.hibernate.SQL=false
   4.
   5.## log hibernate prepared statement parameter values
   6.#log4j.logger.net.sf.hibernate.type=DEBUG, confluencelog
   7.#log4j.additivity.net.sf.hibernate.type=false
   ```

   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 adminstrator 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

Webserver Configuration

- Apache and Apache Connector Tips

- Configure Web Proxy Support for Confluence

Apache and Apache Connector Tips

The speed of downloading attachments is extremely slow. We are experiencing the following speeds

Large file served directly through Apache: 15000 KB/sec
Large file served directly from Tomcat HTTP connector: 14500 KB/sec
Large file served from Confluence (using Apache/mod_jk/Tomcat): 84 KB/sec

You can see that the file served from Confluence is ~176 times slower!

Solution

We upgraded mod_jk from version 1.2.8 to 1.2.10 and the download speed improved significantly to ~12000 KB/sec.

Configure Web Proxy Support for Confluence

Some of Confluence's macros, such as {rss} and {jiraissues} need to make web requests to remote servers in order to retrieve data. If Confluence is deployed within a data-centre or DMZ, it may not be able to access the Internet directly to make these requests. If you find that the {rss} macro does not work, ask your network administrator if Confluence needs to access the Internet through a web proxy.

Configuring an outbound HTTP proxy in Confluence

Proxy support is configured by passing certain system properties to the Java Virtual Machine on startup. These properties follow the conventions defined by Sun:

- http.proxyHost
- http.proxyPort (default: 80)
- http.nonProxyHosts (default: <none>)

At a minimum, you need to define `http.proxyHost` to configure an HTTP proxy. System property configuration is described on the Configuring System Properties page.

Properties `http.proxyHost` and `http.proxyPort` indicate the proxy server and port that the http protocol handler will use.

```
-Dhttp.proxyHost=proxy.example.org -Dhttp.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 |, and in addition a wildcard character (*) can be used for matching. For example:

```
-Dhttp.nonProxyHosts=*.foo.com|localhost
```

Note: You may need to escape the pipe character (|) in some command-line environments.

If the `http.nonProxyHosts` property is not configured, all web requests will be sent to the proxy.

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

Authentication has a few more options in Confluence 2.10 and later, as documented below.

HTTP proxy (Microsoft ISA) NTLM authentication (2.10 and later)

Confluence 2.10 and later supports NTLM authentication for outbound HTTP proxies when Confluence is running on a Windows server.

To clarify, this means the {rss} and {jiraissues} macro will be able to contact external websites if requests have to go through a proxy that requires Windows authentication. This support is not related to logging in Confluence users automatically with NTLM, for which there is a user-contributed authenticator available.

To configure NTLM authentication for your HTTP proxy, you need to define a domain system property, `http.auth.ntlm.domain`, in addition to the properties for host, port and username mentioned above:

```
-Dhttp.auth.ntlm.domain=MYDOMAIN
```

Configuring authentication order (2.10 and later)

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
```

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 yet. See CONF-16941.

Confluence FAQ

Confluence FAQ

Solutions to frequently asked questions and queries about Confluence and commonly encountered issues with the product:

Administration FAQ

- Add Spell Checking To Confluence
- Can I run multiple instances of Confluence & connect them to a central database?
- Capturing HTTP traffic using Wireshark or Fiddler
- Configuring Confluence to send email notifications
- Copy Or Rename A Space
- Customising Confluence Icons
- Disable public account signups
- Displaying System Properties
- Editing your database password
- Enable public anonymous access
- Fix Out of Memory errors by Increasing Available Memory
- Getting a License for a Staging Environment
- How can I retrieve a recently deleted space or page?
- How do I adjust the session timeout
- How do I change the space key?
- How do I check which spaces have email accounts
- How do I configure the Plugin Repository to update its plugins information offline?
- How do I disable indexing of attachments
- How Do I Find My License from the File System?
- How Do I Get More Statistics From Confluence?
- How Do I Identify Inactive Users in Confluence
- How do I know what Confluence version I am running?
- How do I make commenters added as page watchers automatically?
- How do I prevent personal spaces from being shown on the dashboard
- How do I Remove a User who has Content Created
- How do I Remove the Last Updated and Created By Text?
- How do I supress cluster warning message in confluence?
- How to Disable Emoticons
- How to display a banner like the Confluence Documentation space
- How to Force External Links to Open in a New Window
- How to get a Java Heap Dump
- How to Hide the Referrer
- How to Restore Deactivated Users
- How to run a SQL script on your database
- How to Search Confluence for Uses of a Macro
- How to turn on Debugging for indexing
- List page- and space-related details for an attachment using the attachment's name
- Migrate Confluence from one database to another
- Need to Change Server ID for Test Install
- Page Restrictions Performance Considerations
- Passing username and password via URL
- Preventing and Cleaning up SPAM
- Rebuild the Content Indices from scratch
- Redirect users to a page on login
- Redirect users to a site-wide home page after a successful login
- Restrict Attachments Based On File Type
- Search for User Properties in the Database
- Using Firebug Lite in Internet Explorer when browsing a Confluence page
- What are the IP Address Ranges for Atlassian's Servers?
- Where are the files that used to be in my Confluence installation directory?
- Where are user macros stored?

Backup FAQ

- Are there any scripts for backup creation and restore?
- Backup will not import
- Can Confluence be restored from a backup minus attachments?
- Can XML backups be deleted automatically?
- Does running a daily XML backup slow performance?
- How can I reduce the space taken up by automatic backups?
- How to Change the Version of a Space Backup
- How to Find Attachments in Attachments Folder
- Is it Possible to Store the Confluence Home Directory on a Network Share?
- Providing MySQL database with Content Anonymised

Configuration FAQ

- Running Tomcat on a Different Port
- Change default font, color, or spacing in Confluence
- How do I Change the Time of Daily Report Updates
- How to audit Confluence - enabling user access logging
- Share users and groups between Confluence and JIRA
- How do I Modify the Frequency of Content Indexing
- How do I change the default polling time for email in Confluence?
- How do I Configure an Automatic Refresh of the Recently Updated List
- How to Revert from Clustering to Single Node
- Adding a Site-Wide Banner
- Customise Confluence Page Exports
- Editing the Footer
- How do i completely remove the "Space Details" page from Confluence exports?
- Where does Confluence store all its data?
- Running Confluence Behind a Caching Proxy Server
- I am trying to install Confluence but the demo-site.zip file is missing
- How do I Disable Automatic Mail Polling?
- Disabling Profile Pictures on the Recently Updated Dashboard
- How to Disable Profile Pictures from the Recently Updated Section of the Dashboard
- Remove Version from Footer
- Configuring a MySQL Datasource in Apache Tomcat

Installation FAQ

- The Confluence window closes immediately when started
- How do I re-trigger the setup wizard
- Confluence starts but a problem prevents me from accessing the dashboard
- I receive a BUILD FAILED message when trying to create an EAR file in Confluence 2.6 or 2.7
- How much disk space does Confluence need?
- How Do I Make Confluence Accessible from the Root Context with a Tomcat EAR WAR configuration
- How To Run Confluence Standalone and Apache on Port 80 (Different IP Addresses)

JIRA Integration FAQ

- The JIRA Issues Macro behaves problematically
- The JIRA Issues Macro generates an error
- When setting up JIRA user management in Confluence, the Confluence login page loads but login fails
- When setting up JIRA user management in Confluence, the Confluence login page loads with an error

LDAP FAQ

- Are all users in LDAP visible in Confluence administration and can they can be assigned permissions or to groups?
- Can we use LDAP and Confluence groups simultaneously, as a 'mixed mode', where some groups are kept in Confluence and others in LDAP?
- Confluence Domino LDAP FAQs
- Confluence integration with LDAP and Active Directory FAQs
- How are LDAP or Active Directory users counted toward my license limit?
- How can I assign an LDAP user a Confluence account?
- How can I enable LDAP?
- How does Confluence handle user deletions from LDAP? Is the user's assignment to one or more groups still visible?
- I am having a problem with Confluence LDAP integration
- I enabled LDAP and some users are now returned twice under the user browser
- If a Confluence user had a lowercase username, but an LDAP user has the same username in UPPERCASE, does it matter which one I use?
- If a user already exists in Confluence and an LDAP user with the same username is added, which account's password gets used?

Mail Archiving FAQ

- Can Confluence replace my regular mail client?
- How do I get mail into Confluence?
- How do I use the mail archive?
- Okay, I've imported the mail, but where is it?

New User FAQ

- Can I use CamelCaseLinks like they do on WardsWiki?
- Can Users Edit Individual Sections Within a Page?
- How does Confluence differ from a wiki?

RSS Feeds FAQ

- Create an RSS feed for mail from only specified mail accounts
- How do I fix a "Could not download (Feed URL) - Connection timed out (errno238)" error?
- How do I fix a "Could not retrieve (Feed URL) - Not Permitted" error?
- How do I fix an "Error formatting 'macro rss java.lang.NullPointerException'" error?
- How do I fix an "Unable to retrieve (Feed URL) - Connection refused - connect" error?
- How do I force authentication for public feeds?
- Is it possible to delete a feed?
- I want to remove RSS Feeds completely

Upgrade FAQ

- I cannot find the "Rich Text" editor. Is the editor part of Confluence 1.4.3?
- Server ID FAQ
- Upgrade My Trial To A Commercial Version

Usage FAQ

- Add many files to a page at once
- Create a page by passing parameters to a template
- How do I obtain content that hasn't been modified in a certain period of time
- How to Add a Quick Search for Firefox
- How to disable PDF Export
- How to Find Pages with no Label
- How to Make Confluence Open a New Tab when Clicking on the Attachments Link
- How to Reset a Custom Layout
- Redirect to a specific page (home page) within the site after logging in
- Setup email notifications of page updates

Unsupported Platform Information

- Setting up Confluence with IIS
- Using the IBM 64bit J9 JDK

Support Policies

- Bug Fixing Policy
- How to Report a Security Issue
- New Features Policy
- Patch Policy
- Security Advisory Publishing Policy
- Security Patch Policy
- Severity Levels for Security Issues

RELATED TOPICS

Plugin Development
Fix 'Not supported by BasicDataSource' Setup or Startup Error
Troubleshooting HTTPS or SSL-related problems

Administration FAQ

This section contains solutions for common issues or queries associated with administering Confluence.

This section focuses on providing instructions to either perform administration-level tasks or customise Confluence's functionality via its Administration Console.

View one of the following issues or queries for more information:

- Add Spell Checking To Confluence
- Can I run multiple instances of Confluence & connect them to a central database?
- Capturing HTTP traffic using Wireshark or Fiddler
- Configuring Confluence to send email notifications
- Copy Or Rename A Space
- Customising Confluence Icons
- Disable public account signups
- Displaying System Properties
- Editing your database password
- Enable public anonymous access
- Fix Out of Memory errors by Increasing Available Memory
    - Editing the Windows Registry
- Getting a License for a Staging Environment
- How can I retrieve a recently deleted space or page?
- How do I adjust the session timeout

- How do I change the space key?
- How do I check which spaces have email accounts
- How do I configure the Plugin Repository to update its plugins information offline?
- How do I disable indexing of attachments
- How Do I Find My License from the File System?
- How Do I Get More Statistics From Confluence?
- How Do I Identify Inactive Users in Confluence
- How do I know what Confluence version I am running?
- How do I make commenters added as page watchers automatically?
- How do I prevent personal spaces from being shown on the dashboard
- How do I Remove a User who has Content Created
- How do I Remove the Last Updated and Created By Text?
- How do I supress cluster warning message in confluence?
- How to Disable Emoticons
- How to display a banner like the Confluence Documentation space
- How to Force External Links to Open in a New Window
- How to get a Java Heap Dump
- How to Hide the Referrer
- How to Restore Deactivated Users
- How to run a SQL script on your database
- How to Search Confluence for Uses of a Macro
- How to turn on Debugging for indexing
- List page- and space-related details for an attachment using the attachment's name
- Migrate Confluence from one database to another
- Need to Change Server ID for Test Install
- Page Restrictions Performance Considerations
- Passing username and password via URL
- Preventing and Cleaning up SPAM
- Rebuild the Content Indices from scratch
- Redirect users to a page on login
- Redirect users to a site-wide home page after a successful login
- Restrict Attachments Based On File Type
- Search for User Properties in the Database
- Using Firebug Lite in Internet Explorer when browsing a Confluence page
- What are the IP Address Ranges for Atlassian's Servers?
- Where are the files that used to be in my Confluence installation directory?
- Where are user macros stored?

Add Spell Checking To Confluence

Confluence has no inbuilt support for spell checking. You may wish to vote for Confluence to add it's own spell checking or add spell checking to your browser instead:

- Add spell checking to Internet Explorer
- Install the Firefox browser with inbuilt spell checking

Can I run multiple instances of Confluence & connect them to a central database?

Confluence can be clustered.

If running as a single node, you may only have one instance of Confluence connecting to a single database.

There are a couple of reasons for this, but it all comes down to the fact that the Confluence application maintains a lot of state (caches, search indexes) separate to the database, and multiple front-ends will quickly see that state get out of sync, with disastrous effect.

Because of this, Confluence periodically checks to make sure it's the only application accessing the database, and if it finds a conflict it will shut down rather than risk corrupting your data.

Capturing HTTP traffic using Wireshark or Fiddler

This is a quick guide to help you start capturing HTTP traffic when requested by support. This can be helpful either for network traffic issues or for understanding issues with page content loading.



When submitting the captured result to support...
Don't forget to mention the IP Address of the servers involved so Support can go through the TCP dump. Also please mention the time when you performed the operation requested by support.



If you are using HTTPS, try to capture on a HTTP channel if possible. Alternatively, you can try capturing on Windows using Fiddler after enabling the decrypt HTTPS functionality.

Linux/Unix: Wireshark

1. Install Wireshark. (Mirror here)
2. Open your Internet browser.
3. Clear your browser cache.
4. Open Wireshark
5. Click on "Capture > Interfaces". A pop up window will show up.
6. You probably want to capture traffic that goes through your Ethernet Driver. Click on the `Start` button to start capturing traffic via this interface.
7. Visit the URL that you wanted to capture the traffic from.
8. Go back to your Wireshark screen and click the fourth button or press Ctrl + E.
9. After the traffic capture is stopped, please save the captured traffic into a file (in *.pcap format) and attach it to your support ticket.

Wireshark cannot sniff traffic within the same machine (localhost) on Windows. If you would like to sniff local traffic on Windows, try Fiddler.

Windows: Fiddler

1. Download Fiddler.
2. Open it.
3. Clear your browser cache.
4. Browse to your site. Visit the pages that are problematic and a contrasting non-problematic page if appropriate, for contrast.

> Fiddler can capture local traffic by using the machine's name as the host name rather than 'localhost'.

5. Click `File > Save > All Sessions...`.
6. Attach the resulting file for Support.

Configuring Confluence to send email notifications

Configuring Confluence to send email notifications requires a Confluence Administrator to set this up through the **Administration Console**.

To configure confluence to send notifications and test that it is working, follow these instructions:

1. Set up a mail server at Administration -> Mail servers. See Configuring a Server for Outgoing Mail.
    - Make sure you use "Send Test Email" to check that the server is working. Check that you get the test email in your inbox.
2. Go to your profile (using the Preferences link) and click Edit Profile -> Email Preferences, then enable Notify On My Action. (Otherwise you'll have to use multiple users. By default confluence does not send you notifications for your own changes.) See Managing Watches
3. While editing your profile, make sure you have an email address configured. See Configuring a Server for Outgoing Mail
4. Go to a page you wish to get notifications about.
5. Click the Envelope icon in the top right corner to "watch" that page. See Watching a Page.
6. Edit the Page.
7. Either wait a while or: Go to Administration -> Mail Queue. Click "Flush Mail Queue". See The Mail Queue.
8. Check your email.

For instructions on configuring user-level email notifications, see Setup email notifications of page updates.

RELATED TOPICS

Setup email notifications of page updates

Copy Or Rename A Space

Currently Confluence does not support renaming or copying spaces through the user interface. You may wish to vote towards these feature requests:

- Feature request to clone an entire space - this will enable duplication of every page, news item and comment plus space themes and colour schemes.
- Feature request to copy a page hierarchy between existing spaces.

Use The Copy Space Plugin

The best option is to install the Copy Space Plugin.

Note that this plugin is still in **beta** release, and is not officially supported by Atlassian.

Note that copying a space can take a long time and may *appear* to time out when using this plugin (even when the copied space is created correctly). If this occurs, please ensure that your space has not been created before attempting to copy the space again.

Read the developer's notes in the Atlassian blogs.

Manually Clone Or Rename A Space

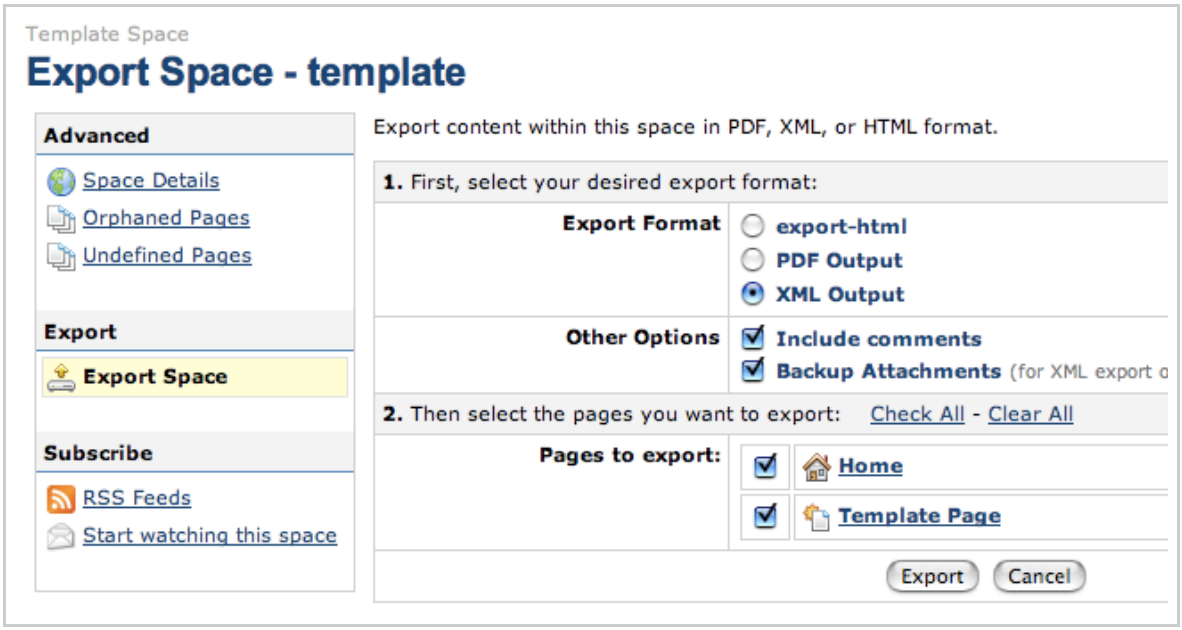
Alternatively, it is possible to manually clone or rename a space by modifying an XML backup of the target space.

Notes

- May require manually updating external links to that space.
- The instructions have been confirmed for Confluence 2.2 onwards. Users running older versions are recommended to upgrade Confluence before continuing.
- Where examples are given, they involve changing `oldkey:Old Space Name` into `newkey:New Space Name`. You need to substitute your own keys for `oldkey` and `newkey`, and your own space names for `Old Space Name` and `New Space Name`.
- When importing a space export for a space that already exists, the previous space content will be overwritten.
- Read the process in full before beginning.

Stage 1: Rename Space

1. Select a new, unique space key and name for the second space. Space keys may only consist of ASCII letters or numbers (A-Z, a-z, 0-9) and no empty spaces are allowed in the key.
2. Clone your production Confluence instance to a test server on another computer now. (For instructions, see Migrating Confluence Between Servers.) You should now have a production server and test server both containing the same data and can avoid the risk of corrupting your production Confluence instance.
3. On the *test server*, login as an administrator.
4. Go to Browse Space -> Advanced -> Export Space and export the target space as `XML` including attachments.



5. Save the space backup.
6. Open the space backup file with a zip file editor and find the file `entities.xml`.
7. Edit `entities.xml` in a text editor.
8. Do a 'Search & Replace' on the old space name as shown below. Swap out `Old Space Name` and `New Space Name` for the actual names.

Search For	Replace With
[CDATA[Old Space Name]	[CDATA[New Space Name]

9. Do a 'Search' for any occurrences of the old space name that occur in user content. You may wish to replace *some or all* of these references with the new space name. Replace `Old Space Name` and `New Space Name` with the actual names.

Search For
Old Space Name

10. Do four 'Search & Replaces' on each of the references to the old space key as shown below. Swap out `oldkey` and `newkey` for your actual keys.

Search For	Replace With

[oldkey]	[newkey]
spaceKey=oldkey	spaceKey=newkey
[oldkey:	[newkey:
key=oldkey]	key=newkey]

For example:

Space name : Test Space
Space key : **test**
New space name : Test Space 2
New space key : **test2**

```
<property name="name"><![CDATA[Test Space]]></property>
<property name="key"><![CDATA[test]]></property>
```

The above search and replace ensures that you will change the **test** `oldkey` to **test2**, and change the "Test Space* `Old Space Name` to **Test Space 2**.

11. Save the modified `entities.xml`.
12. Overwrite the original `entities.xml` in the space backup with the modified version.
13. Login to the *test instance* as a Confluence administrator.
14. Go to Administration -> Backup & Restore. Under 'Upload a zipped backup to Confluence', browse to select to the modified space backup. Check the 'Build index' option and select 'Upload & Restore'.
15. Once the restore process has completed, access the new space and test that you can access pages, embedded links and attachments. Any error in this step indicates that your search and replace was performed incorrectly and you should to retry from step 2.

Stage 2: Restore On Production

1. Backup your production instance now.
2. Login to the *production instance* as a Confluence administrator.
3. Go to Administration -> Backup & Restore. Under 'Upload a zipped backup to Confluence', browse to select to the modified space backup. Check the 'Build index' option and select 'Upload & Restore' If any data loss occurs as a result of using this workaround, immediately revert to the backup.
4. Once the restore process has completed, access the new space and test that you can access pages, embedded links and attachments. If any error occurs during this step, revert to the site backup.
5. If renaming a space, you can delete the old space by going to Browse Space -> Space Admin -> Remove Space. Click OK to remove the old space.

Stage 3: Rename Space References

Links in other spaces to the old space will remain unchanged. If you are renaming a space, you will need to change these links to point to the new space. Users who are copying a space can leave the links pointing to the original space by skipping this stage.

Changing these links depends on if you want to change every link, or only some. If not all links must be changed or you are unwilling to stop your production instance, this must be done by editing each page individually. If all links must be changed, follow the instructions below.

1. Create a site XML backup including attachments from Administration -> Backup & Restore.
2. Save the site XML backup file.
3. Stop the production instance.
4. Create two copies of the site backup. Keep one copy as the original, unmodified backup, the other will be modified. Rename the backups so that it is clear which is being modified.
5. Open the copy for modification and edit `entities.xml`.
6. Do four 'Search & Replaces' on each of the references to the old space key as shown below. Swap out `oldkey` and `newkey` for the actual keys

Search For	Replace With

[oldkey]	[newkey]
spaceKey=oldkey	spaceKey=newkey
[oldkey:	[newkey:
key=oldkey]	key=newkey]

7. Save the updated `entities.xml` back into the modified site XML backup file
8. Start the production instance
9. Import the modified site XML backup from Administration -> Backup & Restore. If you have any problems, revert to the original unmodified backup and redo the links manually instead

Done.

Related

Copy Space Template: There is a feature request being tracked at CONF-4538.

Customising Confluence Icons

Icons are intended to be added/customised from time to time by users to adopt a new look. The file locations are hard to locate.

These steps are under the presumption that users have access to the source code:

- Search for the `atlassian-renderer` library.
- Extract the library and search for the file `/com/atlassian/renderer/DefaultIconManager.java` where the icons are being mapped.
- Check out how the referencing are being done there and add your own icons within

The Confluence 2.10 version of CONFKB:DefaultIconManager is attached to this article.

Disable public account signups

How do I prevent just anyone from being able to sign up for new accounts?

Confluence gives you the option to enable or diable "public signup". If you want to restrict your site to a particular set of users, you can easily disable public signup from the '**General Configuration**' section of the **Administration Console**.

See Enabling or Disabling Public Signup.

RELATED TOPICS

- Hiding the People Directory
- Revoking Space Permissions
- Security Overview
- Confluence Security Advisory 2006-01-23
- Confluence Cookies
- Assigning Space Permissions
- Confluence Security
- Edit in Word Link Macro
- Space Permissions Overview
- View File Macro
- HTML Macro
- How to Hide the Referrer

FAQ Home

Displaying System Properties

Version 3.0.2 and Later

Expanded system properties are visible from `Administration >> System Information`.

Prior to Version 3.0.2

After adding memory, setting a proxy or changing other java options, it's hard to diagnose whether the system has picked them up.

To find out more about what properties are being picked up, download **systemproperties.jsp**. Place it in the <confluence-install>/confluence/admin folder. Access the URL: http://<yourbaseurl>/admin/systemproperties.jsp.

No restart of Confluence is required.

Editing your database password

To reset the password for the database user that confluence's uses to connect to the database, follow the guide below.

If your confluence instance connects directly via JDBC, then your password will be in your <CONFLUENCE_HOME>/confluence.cfg.xml file. E.g.

```
1.<property name="hibernate.connection.driver_class">com.mysql.jdbc.Driver</property>
2.<property name="hibernate.connection.password">confluencepass</property>
3.<property name="hibernate.connection.url">jdbc:mysql://localhost/confluence?autoReconnect=true&lt
;/property>
4.<property name="hibernate.connection.username">confluencedbuser</property>
5.<property name="hibernate.database.lower_non_ascii_supported">true</property>
6.<property name="hibernate.dialect">com.atlassian.hibernate.dialect.MySQLDialect</property>
```

Change the "hibernate.connection.password" property to the correct value (in the above the example replace "confluencepass" with the new password).

If you're connecting via datasource then you will see in the confluence.cfg.xml file something like:

```
1.<property name="hibernate.setup">true</property>
2.<property name="hibernate.dialect">com.atlassian.hibernate.dialect.MySQLDialect</property>
3.<property name="hibernate.connection.datasource">java:comp/env/jdbc/confluence]</property>
```

I.e. the property "hibernate.connection.datasource" is defined. If so your password is defined within your datasource. Each application server stores its information differently, but if you are using tomcat, then check your server.xml file.

Enable public anonymous access

How do I configure Confluence for public-anonymous access?

There are two different permissions that need to be set to allow anonymous access to a Confluence site. First, the '**Anonymous**' user needs the global "Use Confluence" permission, secondly you need to give '**Anonymous**' permissions in each space you want to make public. Full instructions can be found here.

RELATED TOPICS

Security Overview
Users and Groups
Confluence FAQ

Fix Out of Memory errors by Increasing Available Memory



JDK 1.4 does not provide information why the `OutOfMemory` error occurred. Since Confluence version 2.8, this JDK is not supported any longer.
JDK 1.5 or **1.6** are the recommended JDK to be used and they do provide a description of the OOM error. JDK 1.6 also has around a 20% performance improvement over 1.5

Since the default memory setting usually is around 64 or 128MB (256MB in Confluence 2.2 and later), you might have to adjust the settings to run a bigger Confluence instance with sufficient memory.

On this page:

- Diagnosis and Common Causes
- Determining the various causes of memory errors
    - java.lang.OutOfMemoryError: PermGen space

- java.lang.OutOfMemoryError: Java Heap Space
- OutOfMemoryError: unable to create new native thread
- OutOfMemoryError: GC overhead limit exceeded
- OutOfMemoryError: Requested array size exceeds VM limit

- Setting the Memory Settings
  - RELATED TOPICS

Diagnosis and Common Causes

There are several reasons that out of memory exceptions can be thrown. Either the virtual machine Confluence is using has hit its allocated memory limit, the system on which Confluence is running has run out of physical and virtual memory, or Confluence is consuming too much memory. In the first case, you should modify the maximum heap size of the virtual machine, per the instructions in this document; in the second or third cases, the solution is to identify the culprit of the memory leak.

For help determining which memory settings to choose, consult Managing Application Server Memory Settings.

If you have not yet set your memory settings and your usage has increased, it's likely that you must set your memory, described below. If your usage patterns have not changed but you've added a plugin or done an upgrade, it's likely that there is a memory leak in a plugin.

1. If you're using the in-memory database (HSQLDB), migrate to an external database. The in-memory database can use a lot of memory.
2. If you are using XML backups, disable them and move to the Production Backup Strategy. The XML backup process can be a memory hog.

To troubleshoot potential memory leaks, enter Plugin Support Mode. Take thread dumps during normal operations and during an outage, and submit this information in a support ticket.

Determining the various causes of memory errors

There are different kinds of memory limits inside the Java Virtual Machine. Each limit can be configured independently. But you must first find out which limit you have reached.

java.lang.OutOfMemoryError: PermGen space

If you get the error message: `java.lang.OutOfMemoryError: PermGen space` this means that you have exceeded Java's default 64Mb block for loading class files. This can happen if many plugins are installed. You may want to increase the PermGen memory size to suit your needs.

In the following sample, the blue parameter shows how the PermGen Memory has been set to 192 megabytes. This value should be set depending on your memory requirements. 192m should be sufficient for Confluence even when many plugins are installed.

```
JAVA_OPTS="-Xms128m -Xmx1024m -XX:MaxPermSize=192m $JAVA_OPTS -Djava.awt.headless=true "
```

Note: The other parameters in this sample are just shown to give you some context, and are not part of this example.

java.lang.OutOfMemoryError: Java Heap Space

Heap space memory errors occur when the application has to deal with large amounts of data or users. These errors will contain only a `java.lang.OutOfMemoryError`, e.g. without the reference to PermGen space as above. You should try to increase the heap size to solve this problem. This requires configuring the Xmx and Xms parameters. In the following example, the maximum heap size is set to 1024 megabytes. This should be enough for small to medium deployments.

```
JAVA_OPTS="-Xms128m -Xmx1024m -XX:MaxPermSize=256m $JAVA_OPTS -Djava.awt.headless=true "
```

Note: The other parameters are only shown to give you context and are not part of this example.

Deployments with high usage patterns may require additional memory. For high-usage deployments, it is recommended to set both Xms and Xmx as the same value (eg. -Xms1024m -Xmx1024m), _provided the memory is available_. On the other hand, adding too much memory can also cause problems (see below), so you should increment memory carefully, for example in increments of 128 megabytes.

OutOfMemoryError: unable to create new native thread

This error occurs when the operating system is unable to create new threads. This is due to the JVM Heap taking up the available RAM.

> *Big heaps take away from the space that can be allocated for the stack of a new thread*

For 32bit Linux generally the maximum heap size of the JVM cannot be greater than 2GB. Windows systems will typically split the available physical memory 50:50 as Application and Kernel/System space, so please do not allocate an amount exceeding or approaching that split. The size of the stack per thread can also contribute to this problem. The stack size can reduce the number of threads that can be created.

To fix this problem, you should reduce the size of your JVM Heap and also the size of the stack per thread.
The stack size can be changed with the following (example) parameter:

"-Xss512k"

Please refer to this guide as a reference for JVM tuning.

OutOfMemoryError: GC overhead limit exceeded

This error indicates that the JVM took too long to free up memory during its GC process. This error can be thrown from the Serial, Parallel or Concurrent collectors. It often means that the Xmx value is *too high* - you might consider lowering it. See Garbage Collector Performance Issues for more details. For more severe and persistent performance issues relating to GC, it is recommended to change to a parallel collector, and to ensure that Confluence has access to the memory demanded by its users.

The parallel collector will throw an OutOfMemoryError if too much time is being spent in garbage collection: if more than 98% of the total time is spent in garbage collection and less than 2% of the heap is recovered, an OutOfMemoryError will be thrown. This feature is designed to prevent applications from running for an extended period of time while making little or no progress because the heap is too small. If necessary, this feature can be disabled by adding the option -XX:-UseGCOverheadLimit to the command line.

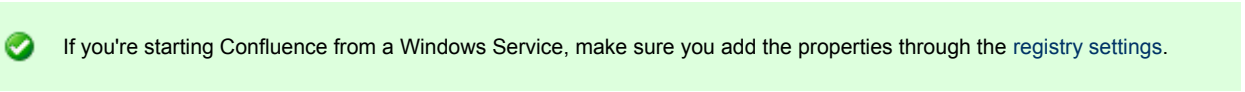
This kind of OutOfMemoryError can be caused if user requests drown the available resources in the JVM. When this occurs, performance will degrade aggressively. This will eventually require a restart or the application may recover.

OutOfMemoryError: Requested array size exceeds VM limit

This is a rare error and indicates that Confluence attempted to allocate an array that is larger than the Java heap size. More details regarding this error can be found here.
This is due to a known limitation of the JVM as documented here. We have a bug lodged against this, to better handle this exception in Confluence.

Setting the Memory Settings

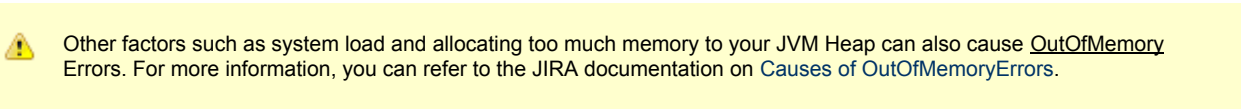
How to set the heap or permanent generation memory depends on your distribution, platform, and how you start Confluence. Refer to Configuring System Properties.



If you're starting Confluence from a Windows Service, make sure you add the properties through the registry settings.



To verify if your settings have been picked up, check Displaying System Properties.



Other factors such as system load and allocating too much memory to your JVM Heap can also cause OutOfMemory Errors. For more information, you can refer to the JIRA documentation on Causes of OutOfMemoryErrors.

RELATED TOPICS

Managing Application Server Memory Settings
Installing the Confluence EAR-WAR Edition
FAQ Home
Tomcat JVM options and Modify the Default JVM Settings
Websphere Tuning JVM
Logging A Thread Dump

Editing the Windows Registry

If you need to adjust memory settings and are using a Windows Service, you can adjust your memory settings using the instructions in this page.

You have two choices on how to set the memory settings.

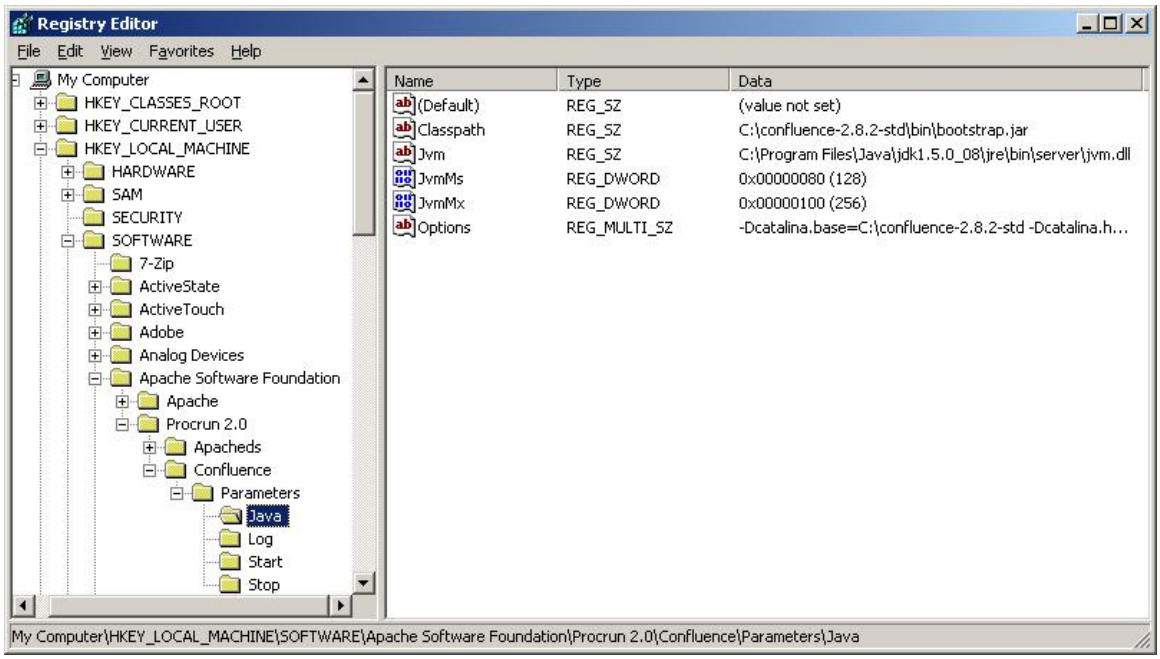
From the Command Line

Run this command from the command line:

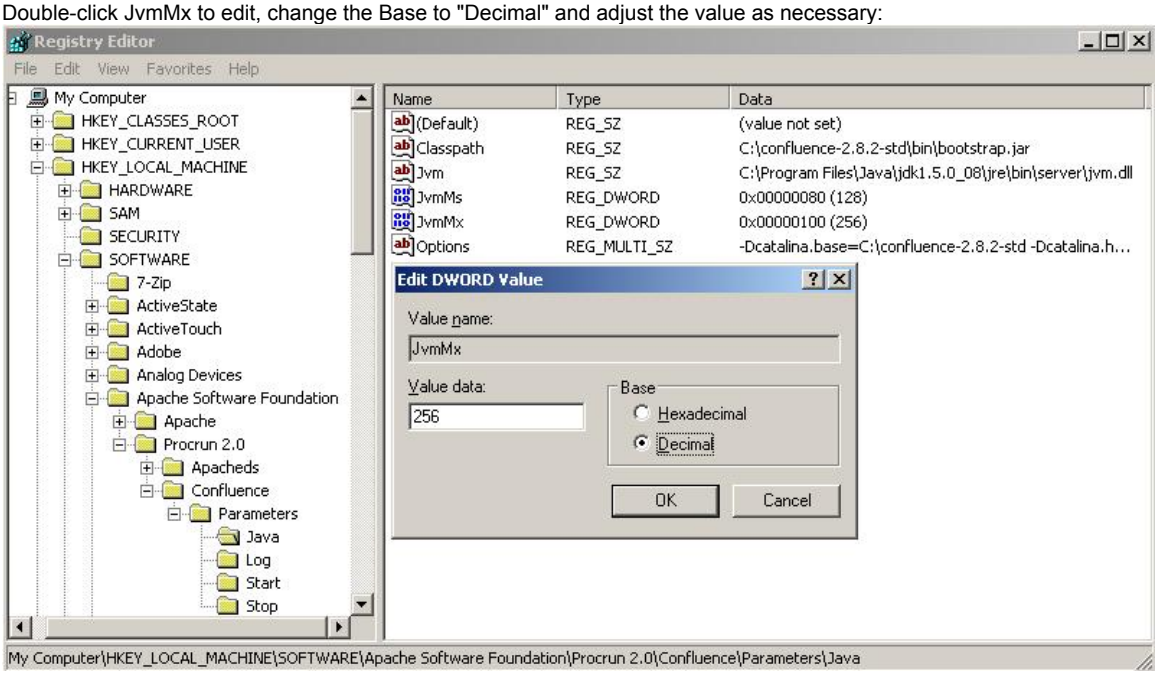
```
tomcat6 //US//Confluence --JvmMs 1024 --JvmMx 1024 ++JvmOptions="-XX:MaxPermSize=256m"
```

Edit the Registry

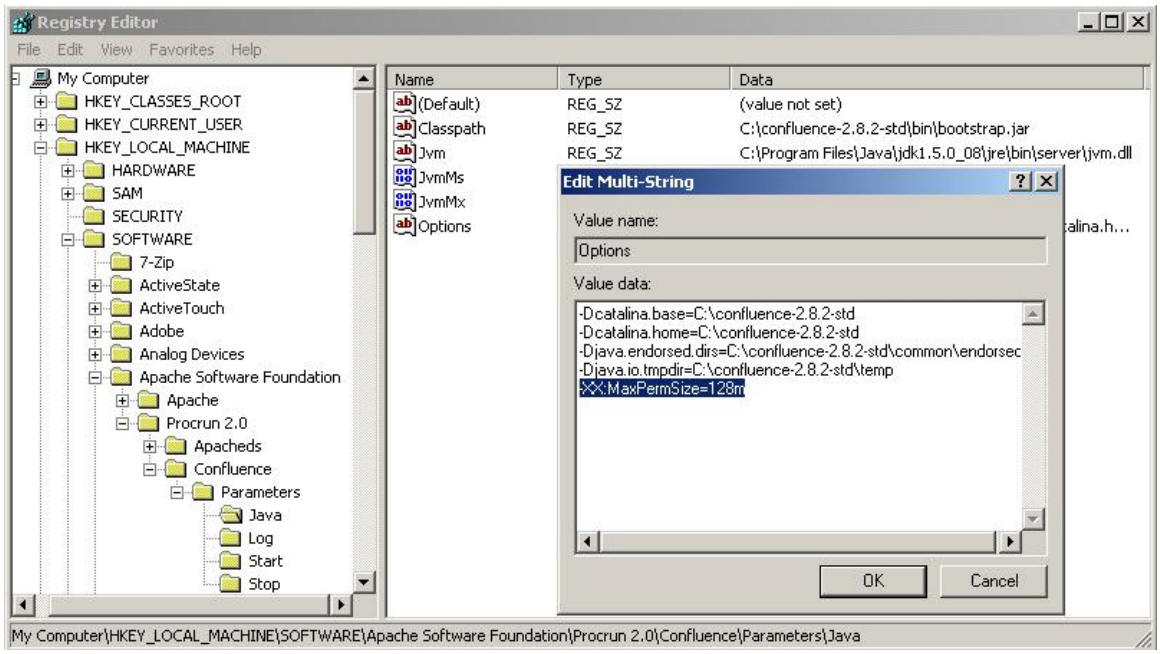
1. Shut Confluence Service.
2. Go to `Start >> Run >> Type in regedit`. (On versions of Windows prior to Windows XP, please type `regedt32` instead of `regedit`.)
3. In the registry editor, click to HKEY_LOCAL_MACHINE -> SOFTWARE -> Apache Software Foundation -> Procrun 2.0 -> Confluence -> Parameters -> Java. Here you will see an entry for JvmMx, which is the "maximum memory" setting. Choose an appropriate value based on Managing Application Server Memory Settings.



4. Double-click JvmMx to edit, change the Base to "Decimal" and adjust the value as necessary:



5. If setting the Perm Gen size, click Options and add -XX:MaxPermSize=256m:



If you want to add other JVM parameters, you can use the same method as when adding the permgen setting.

6. Restart Confluence Service.

Getting a License for a Staging Environment

If you already have a developer license, you can add it under the Administration > License Details page. You can also create a new key as detailed below:

> **Only a technical contact for your commercial/academic license is able to create a Developer license**

Atlassian supplies 'developer' licenses which can be used by existing commercial license holders who wish to deploy non-production installations of our software to use in QA/staging environments. Developer licenses are free of charge to commercial license holders and, like our commercial offerings, they include 12 months of updates starting from the date of purchase of the commercial license.

If you hold a commercial license, you can obtain a free developer license by performing the following:

1. Log in to your Atlassian account.
2. Under the "Licenses" heading, all of your licenses will be displayed. Click the plus sign next to a license to view its details.
3. Click the **'View Developer License'** link in the bottom right corner of the license detail panel, below your commercial license key.

Your new developer license will be generated and displayed in a pop up window. Repeat this process as many times as necessary for multiple developer licenses. If you're unable to create the license, contact our sales department for help.

Developer licenses are not compatible with all versions of our products. The table below indicates which versions of each product support developer licenses.

JIRA	Confluence	Bamboo	Clover for Ant	Clover for Eclipse 3	Clover for IDEA 6	Crowd	Crucible	FishEye	JIRA Perforce Plugin	JIRA VSS Plugin
3.7+	2.3+	All	1.3.14+	1.2.12+	1.0.2+	All	1.0.3+	1.3.3+	JIRA 3.11+ (Disabled)	JIRA 3.11+

If you are working with an older version that doesn't recognize developer licenses, you can use your existing commercial license in your test environment.

How can I retrieve a recently deleted space or page?

To restore a page, you may find the wanted information in the Restoring a Deleted Page documentation.

Unfortunately there is no easy way to restore a space - tell your Confluence Administrator to restore the site's daily backup and retrieve the deleted space from there.

To setup Confluence in a staging environment you may need to use a Developer License.

How do I adjust the session timeout

To change the default session timeout (which is 60 minutes) you must edit the file `web.xml`. For EAR/WAR installation, this file can be found in `<YOUR DEPLOYMENT>/confluence/WEB-INF/web.xml` and for a standalone instance, the file is located in `<YOUR CONFLUENCE INSTALLATION DIRECTORY>/conf/web.xml`

The element you want to edit in the `web.xml` file is:

```
1.<session-config>
2.          <session-timeout>60</session-timeout>
3.</session-config>
```

The value within the **session-timeout** tag defines the amount of time the session will exist, in minutes.

Note that after editing the `web.xml` file you will need to restart confluence for your change to take effect.

On a related note, to configure Confluence's internal connection pool timeout period, please tweak the **c3p0.timeout** property in your `<Confluence-Home>/confluence.cfg.xml` file:

```
<property name="hibernate.c3p0.timeout">30</property>
```

This value is an indication of the number of seconds a connection can remain pooled but unused before being discarded. Zero means idle connections never expire.
More details can be found here.

How do I change the space key?

See Copy Or Rename A Space.

How do I check which spaces have email accounts

If you have numerous spaces in Confluence and would like to check for spaces which have a mail setting, you can run a database query like the one below:

```
1.select spaces.spacename, spaces.spacekey, bandana.bandanakey
2.from spaces, bandana
3.where spaces.spacekey = bandana.bandanacontext AND bandana.bandanakey=
"atlassian.confluence.space.mailaccounts";
```

RELATED TOPICS

How do I Disable Automatic Mail Polling?

How do I configure the Plugin Repository to update its plugins information offline?

With the launch of plugins.atlassian.com, the Confluence Plugin Repository SVN metadata has been deactivated. For existing clients, requests to confluence.atlassian.com for repository metadata are now transparently redirected to plugins.atlassian.com, which provides up-to-date plugin information in the same XML format as that served from the old repository proxy.

While Confluence installations with access to confluence.atlassian.com see no impact from this change, those installations which cannot access confluence.atlassian.com need special instructions.

Download the latest plugins.atlassian.com XML for your version of Confluence

PAC makes the plugin repository client data available at the URL

`https://plugins.atlassian.com/server/legacy/confluence/xml/<bn>`

where `<bn>` is the build number of the Confluence installation. The build number is visible at **Admin -> System Information** (near the bottom of the page). Build numbers for all released Confluence builds are also available here.

Save the resulting file and place it on a Web server accessible from Confluence. This can be the same server as Confluence itself; for example, if you're running the standalone version of Confluence at `http://localhost:8090`, place the saved XML in the `confluence` directory inside your standalone installation.



Keeping the XML up to date
When a change is made on plugins.atlassian.com to a plugin's data, the generated XML will update at most one day after the change is made. Consider writing a `cron` job or similar to fetch the XML on a weekly basis and store it in your chosen location to ensure you are always up to date on plugin information.

Configure the Plugin Repository Client

Now go to **Admin -> Plugins -> Confluence Atlassian Plugin Repository** and click on the **Configure** link.

There are four plugin repository data source options:

1. **HTTP Data Source Generator** - Plain text XML file served over HTTP. Never used, except for debugging.
2. **Subversion Generator *(Deprecated)*** - Check out the metadata from SVN and generate the XML internally. As this uses the deprecated SVN metadata, it can be ignored.
3. **Local Generator *(Deprecated)*** - Read the metadata off your local filesystem and generate the XML internally. As this uses the deprecated SVN metadata, it can be ignored.
4. **Proxy Client** - pull XML data from another Confluence Plugin Repository Client that has been configured to act as a server.

In this case you would need to choose the **Proxy Client** option, as it is able to fetch XML from any server, not just plugins.atlassian.com.

Paste the URL to the XML data in the **Data Source** field and uncheck the **Data Source Proxying** option if it is checked. Then click **Save** and select **Admin -> Plugin Repository**, and the client should load the repository info as before.

RELATED TOPICS

Confluence Repository Client

How do I disable indexing of attachments

Sometimes a user can experience problems indexing large MSExcel or MSPowerPoint documents and the reindexing may cause potential `Unknown Ptg` warning messages that are harmless. There is already a request to Suppress these warnings from the re-indexing of unreadable documents by the POI library.

The error is usually not serious yet can sometimes cause problems when large attachments are used. So you may like to disable indexing of a particular type of document.

To do this, you can use one of the methods described below.

Method 1: Using the Administration Console

You can disable the relevant modules from the Attachment Extractors or Office Connector plugins, by going to **Administration -> Configuration -> Plugins** and disabling the relevant plugin modules:

- To disable the indexing of PDF attachments, go to the **Attachment Extractors** plugin and disable the following module:
  - PDF Content Extractor — For PDF attachments
- To disable the indexing of Office attachments, go to the **Office Connector plugin** and disable the following modules as required:
  - Word Content Extractor — For Word 97/2007 (`.doc` and `.docx`) attachments
  - PowerPoint 97 Content Extractor — For PowerPoint 97 (`.ppt`) attachments
  - PowerPoint 2007 Content Extractor — For PowerPoint 2007 (`.pptx`) attachments
  - Excel 97 Content Extractor — For Excel 97 (`.xls`) attachments
  - Excel 2007 Content Extractor — For Excel 2007 (`.xlsx`) attachments

The search query will ignore all attachments of the type corresponding to the disabled module.

Method 2: Editing the `atlassian-plugin.xml` files of plugins

You need to modify the content of the `atlassian-plugin.xml` file in the following JAR files and comment out the relevant **file type extractor**:

- `confluence-attachment-extractors-x.x.jar` (for PDF) or
- `OfficeConnector-x.x.jar` (for Office files)

Both of these JAR files are located in the `confluence\WEB-INF\classes\classes\com\atlassian\confluence\setup\atlassian-bundled-plugins.zip` file.

If you are unfamiliar with modifying JAR files, please refer to the Editing Files within JAR Archives document for further information.

You can identify file type extractors in `atlassian-plugin.xml` files by the occurrence of `ContentExtractor` in their `key` attribute.

Once the `ContentExtractor` for a file type is disabled, all files of that type become unsearchable.

The example below shows a **pdfContentExtractor** disabled which would prevent PDF attachments from being indexed.

```
<atlassian-plugin key="com.atlassian.confluence.plugins.attachmentExtractors" name="Attachment Extractors">
        <plugin-info>
                <description>This plugin extracts searchable text from various attachment types.</description>
                <version>1.1</version>
                <vendor name="Atlassian Pty Ltd" url="http://www.atlassian.com/"/>
        </plugin-info>

        <!--
        <extractor name="PDF Content Extractor" key="pdfContentExtractor" class="com.atlassian.bonnie.search.extractor.PdfContentExtractor" priority="1100">
                <description>Indexes contents of PDF files</description>
        </extractor>
        -->

</atlassian-plugin>
```

The following table shows the file type extractors in the `atlassian-plugin.xml` of the `OfficeConnector-x.x.jar` file, which require commenting out to prevent indexing:

Type of attachment	File Type Extractor

Word 97/2007 (`.doc` and `.docx`)	`<extractor name="Word Content Extractor" key="wordContentExtractor" class="com.atlassian.confluence.extra.officeconnector.index.word.WordTextExtractor" priority="1099">`
	`<description>Indexes contents of Word 97/2007 files</description>`
	`</extractor>`

PowerPoint 97 (`.ppt`)	`<extractor name="PowerPoint 97 Content Extractor" key="ppt97ContentExtractor" class="com.atlassian.confluence.extra.officeconnector.index.powerpoint.PowerPointTextExtractor" priority="1099">`
	`<description>Indexes contents of PowerPoint 97 files</description>`
	`</extractor>`

PowerPoint 2007 (`.pptx`)	`<extractor name="PowerPoint 2007 Content Extractor" key="ppt2k7ContentExtractor" class="com.atlassian.confluence.extra.officeconnector.index.powerpoint.PowerPointXMLTextExtractor" priority="1099">`
	`<description>Indexes contents of PowerPoint 2007 files</description>`
	`</extractor>`

Excel 97 (`.xls`)	`<extractor name="Excel 97 Content Extractor" key="excel97ContentExtractor" class="com.atlassian.confluence.extra.officeconnector.index.excel.ExcelTextExtractor" priority="1099">`
	`<description>Indexes contents of Excel 97 files</description>`
	`</extractor>`

Excel 2007 (`.xlsx`)	`<extractor name="Excel 2007 Content Extractor" key="excel2k7ContentExtractor" class="com.atlassian.confluence.extra.officeconnector.index.excel.ExcelXMLTextExtractor" priority="1099">`
	`<description>Indexes contents of Excel 2007 files</description>`
	`</extractor>`

How Do I Find My License from the File System?

If you're not sure where your license is in my.atlassian.com, you can look in `<confluence-home>/confluence.cfg.xml`.

How Do I Get More Statistics From Confluence?

Confluence has several plugins that you can use for generating statistics, such as:

- Use the SQL and Chart plugins together.
    - Our user community have contributed some great SQL queries.
- The Reporting Plugin contains macros which allow powerful and flexible reporting on Confluence content and content from other locations.
- Customware's Tracking Plugin contains macros for anonymously tracking content access. Otherwise known as *hit counting*, this macro provides the ability to count the number of times a given piece of content has been viewed. It does not count views by the most recent editor of the page.
    - Try using the `tracking-info` macro to track the number of times an attachment is downloaded
- Make user macros like countpages, which counts the number of pages in a space.
- Use the Contributors and Contributors Summary macros to get more details regarding the contributors for required pages.
- Statistical Analysis Plugin from Adaptavist, is another cluster-ready, enterprise scalable third-party plugin.

In addition, Confluence has a built-in access logging mechanism, which shows who has logged in and the URL invoked. To enable it, you need to modify a couple of configuration files and restart Confluence. The *traditionally* generated access log can then be analysed by one of the available access log analyser tools such as Webalizer, Google Analytics or AwStats which can generate useful statistics.

For more information on using Google Analytics and Confluence you may wish to refer to this blog post by David Simpson.



See our documentation on Obtaining Confluence Instance Metrics

If none of the above tools satisfy your requirements, you can create a feature request in jira. Please note that there are already several feature requests and improvements created by our customers all being collated under one umbrella issue

Please cast your vote, add your comments to the discussion and don't forget to add yourself as a watcher to be notified on progress. All our improvements and new features are implemented according to this guide.

RELATED TOPICS

Live Monitoring Using the JMX Interface
Live log viewer plugin.
Tomcat's access logs.

How Do I Identify Inactive Users in Confluence

If you want to disable inactive users and prevent them from being counted by Confluence license count it is possible to find out by running queries against your database.

This is particularly useful if you have numerous users.

This query will show you the users in reverse order according to the date they last logged in.

```
1.select u.name, p.date_val from users u
2.join OS_PROPERTYENTRY p on u.id = p.entity_ID
3.where entity_key='confluence.user.last.login.date'
4.order by date_val desc;
```

This query will show you the same information but based on their previous login:

```
1.select u.name, p.date_val from users u
2.join OS_PROPERTYENTRY p on u.id = p.entity_ID
3.where entity_key='confluence.user.previous.login.date'
4.order by date_val desc;
```

Here's another query that will show you users who have not made any edits or comments since 2007:

```
1.select u.name from users u where u.name not in (select creator from content where contenttype in
('BLOGPOST', 'COMMENT', 'PAGE') and year(creationdate) > 2007);
```

Using JIRA User Management

While Confluence does not store any of the JIRA users information in its database, Confluence still stores the login details in the `OS_PROPERTYENTRY` table, which we can refer to from the jira user id.

This means we can modify the queries above easily.

Since we are going to use two databases (assuming that they are located on the same server), you need to run this first:

```
1.use yourConfluenceDatabaseName;
```

Then you can execute the queries below, just make sure that you replace `YourJIRADATABASE` in the queries below with your own JIRA database name.

This query will show you the users in reverse order according to the date they last logged in:

```
1.select u.username, p.date_val from YourJIRADATABASE.userbase u
2.join OS_PROPERTYENTRY p on u.id = p.entity_ID
3.where p.entity_key='confluence.user.last.login.date'
4.order by p.date_val desc;
```

If you need to know the users' email addresses so you can contact them, run the following:

```
01.select u.username, psfn.propertyvalue as full_name, psem.propertyvalue as email_address,
p.date_val as last_login_date
02.from YourJIRADATABASE.userbase u
03.join OS_PROPERTYENTRY p on u.id = p.entity_ID
04.join YourJIRADATABASE.propertyentry pefn on pefn.entity_id = u.id
05.join YourJIRADATABASE.propertystring psfn on psfn.id = pefn.id
06.and pefn.entity_name = 'OSUser'
07.and pefn.property_key = 'fullName'
08.join YourJIRADATABASE.propertyentry peem on peem.entity_id = u.id
09.join YourJIRADATABASE.propertystring psem on psem.id = peem.id
10.and peem.entity_name = 'OSUser'
11.and peem.property_key = 'email'
12.where p.entity_key='confluence.user.last.login.date'
13.order by p.date_val desc;
```

This query will show you the same information but based on their previous login:

```
1.select u.username, p.date_val from YourJIRADATABASE.userbase u
2.join OS_PROPERTYENTRY p on u.id = p.entity_ID
3.where p.entity_key='confluence.user.previous.login.date'
4.order by p.date_val desc;
```

Here's another query that will show you users who have not made any edits or comments since 2007:

```
1.select u.username from YourJIRADATABASE.userbase u where u.username not in (select creator from
content where contenttype in ('BLOGPOST', 'COMMENT', 'PAGE') and year(creationdate) > 2007);
```

How do I know what Confluence version I am running?

At the bottom of a Confluence page you will see a line like this:

```
Powered by Atlassian Confluence 2.10.1, the Enterprise Wiki. Bug/feature request - Atlassian news
- Contact administrators
```

In the above it means that you are running Confluence version 2.10.1.

If you do not see that line, you can visit **http://<Your Confluence URL>/admin/systeminfo.action** and find out the version from there.

How do I make commenters added as page watchers automatically?

You will need to modify a velocity file to achieve this:

1. Open up `<confluence install>/confluence/template/includes/macros.vm` and look for `comment.notification.enable`.
2. Replace `value='false'` to `value='true'`.

It should look like this:

```
1.#if ($remoteUser && !$action.userWatchingPage && !$action.userWatchingSpace)
2.          #bodytag ("Checkbox" "name='watchPageAfterComment'" "theme='notable'"
"label='comment.notification.enable'" "value='false'")
3.                    #param ("tabindex" "203")
4.          #end
5.#end
```

 Your users will automatically be a watcher of a page after they posted a comment on the page. Even if their comment was removed they still will receive notifications if there is an update made on the page. To stop this, they should remove themselves from the watcher list.

How do I prevent personal spaces from being shown on the dashboard

To disable personal pages from being shown in the dashboard one will need to customise the recently updated content macro by giving it a different parameter.

Here's how:

1. Login as an admin and go to your Confluence administration console

2. Under the "Look and Feel" menu, click on "Layouts" (or type this in your browser: http://yourConfluenceURL.com/admin/listdecorators.action)
3. Under "Site Layouts" look for "Global Layouts" and choose "Create Custom"
4. In the layout editor, look for this line:

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

5. Enter the list of space that is allowed to shows up by inserting the SpaceKey in the line:

```
1.$helper.renderConfluenceMacro(
"{recently-updated:spaces=SpaceKey1,SpaceKey2|showProfilePic=true}")
```



If you have numerous global spaces, it would be more convenient to use **@global** instead of listing each space separately. See RecentlyUpdatedMacro-Parameters for the list of parameters avaliable for the function.

6. Save



Personal spaces will still show up in daily email update. If you don't want this to show in the email update, you need to make the space accessible only to selected groups or people.

How do I Remove a User who has Content Created

Confluence doesn't allow the removal of user who has created any content. The standard method for removal is described in Removing a User. If you want to delete the content as well, the SQL is described below. Locate the content created by the user using the queries:

```
1.SELECT * FROM CONTENT WHERE contenttype = 'COMMENT' and creator = '<Username that you wish to
remove>'
```

and

```
1.SELECT * FROM CONTENT WHERE contentid IN (
2.SELECT DISTINCT pageid FROM CONTENT WHERE contenttype = 'COMMENT' and creator = '<Username that you
wish to remove>')
```

Run the same DELETE commands after verifying the content is indeed content you wish to delete.

How do I Remove the Last Updated and Created By Text?

Open the file `<confluence-install>/confluence/decorators/includes/page-metadata.vm` in a text editor and remove the following text:

```
01.#if ($page.isLatestVersion() == true)
02.          $action.getText('added.by.user.last.edited.on.date', ["#usernameLink
($page.creatorName)","#usernameLink ($page.lastModifierName)","$action.dateFormatter.format(
$page.lastModificationDate )"])
03.          #set ($previousPage = $action.getPreviousVersion($page.previousVersion))
04.          #if ($!previousPage)
05.                    <span class="noprint"> (<a href=
"$req.contextPath/pages/diffpages.action?pageId=${page.id}&originalId=$previousPage.id"
>$action.getText('view.change')</a>)</span>
06.          #end
07.#else
08.          $action.getText('added.by.user.edited.on.date', ["#usernameLink ($page.creatorName)",
"#usernameLink ($page.lastModifierName)","$action.dateFormatter.format(
$page.lastModificationDate )"])
09.#end
```

How do I supress cluster warning message in confluence?

You might find that under Logging and Profiling in Confluence Admin, either (or both)

```
1.com.atlassian.confluence.cluster.safety
2.com.atlassian.confluence.cluster
```

are set to **DEBUG**.

Please change them to INFO and the warning messages in logs should disappear.

How to Disable Emoticons

The text markup (i) is rendered as  , a principle which applies to many more Confluence Emoticons. There is a feature request at

CONF-4884 for enabling or disabling emoticon rendering in this manner via Confluence's Administrative features. However, to disable emoticon rendering in Confluence:

For Confluence 2.8.x and earlier:

1. Open up the `wikiFiltersSubsystemContext.xml`:
   - For Confluence 2.5.x and earlier, this file is located in *WEB-INF/classes* directory
   - For Confluence 2.6.x to Confluence 2.8.x, this file is located in *WEB-INF/lib/confluence-2.x.y.jar/plugins*. Please follow the steps as suggested here to edit files in a .jar file.
2. Determine the following line and remove/comment out the content:

```
1.<ref local="emoticonRendererComponent"/>
```

For Confluence 2.9.x and later:

1. Open up the `wiki-renderer-components.xml`, located in *WEB-INF/lib/confluence-2.x.y.jar/plugins*. Please follow the steps as suggested here to edit files in a .jar file.
2. Determine the following line and remove/comment out the content:

```
1.<renderer-component key="emoticon" name="Emoticon Renderer"
2.class="com.atlassian.confluence.renderer.plugin.SpringRendererComponentFactory" weight=
"10000">
3.          <param name="componentName">emoticonRendererComponent</param>
4.</renderer-component>
```

Alternatively, manually escape the character with a "slash","\". For example:

```
\(i)
```

Code Macro and Noformat Macro is also an option as any emoticons wrapped within the macro will be disabled automatically.

How to display a banner like the Confluence Documentation space

The documentation for spaces prior to the current documentation contain a banner:

Edit the main layout for the space. After the Content div header:

```
01.## CONTENT DIV BEGINS
02.<div id="header">
03.          #quickSearch()
04.          <ul id="header-menu-bar" class="ajs-menu-bar">
05.                    #if($sitemeshPage.getProperty("page.browse-menu"))
06.                              $sitemeshPage.getProperty("page.browse-menu")
07.                    #else
08.                              #menuMacros_renderGlobalBrowseMenu()
09.                    #end
10.                    #menuMacros_renderUserMenu()
11.          </ul>
12.          $!sitemeshPage.getProperty("page.breadcrumbs")
13.</div>
```

Add the following:

```
1.<div class="noteMessage">
2.          <strong>This documentation relates to an early version of Confluence.</strong><br>
3.View <a href="http://confluence.atlassian.com/display/DOC/$title">this page in the current
documentation</a> or visit the <a href="http://confluence.atlassian.com/display/DOC/Confluence
+Documentation+Home">current documentation home</a>.
4.</div>
```

How to Force External Links to Open in a New Window

By default, external links like Atlassian and Google are opened in the same window. To force them to be opened in a new window, follow these steps:

1. Visit `Administration > Custom HTML`
2. Click `Edit`
3. In the `At end of the HEAD` field, insert this code:

```
1.<script>
2.jQuery(document).ready(function() {
3.          jQuery(".external-link").attr("target", "_blank");
4.});
5.</script>
```

4. Hit `Save`

How to get a Java Heap Dump

Getting the heap dump

If you hit the `java.lang.OutOfMemoryError: Java Heap Space` and you have elminated the usual causes you may need to get a Java heap dump, to determine the cause.

To get a heap dump add the `-XX:+HeapDumpOnOutOfMemoryError` parameter to your JAVA_OPTS

For example:

```
JAVA_OPTS="-Xms128m -Xmx1024m -XX:MaxPermSize=192m -XX:+HeapDumpOnOutOfMemoryError $JAVA_OPTS
-Djava.awt.headless=true "
```

Next reproduce the out of memory error.

The next time you have an out of memory error, a *.hprof file will be created that is approximately the size of your java heap, i.e. 1024m (according to the above example).

Please wait till its completely written out (before restarting confluence) and attach a zip of this dump to your support case.



Please note that your -Xmx should not be bigger than 1536m otherwise it is near impossible to open the heap dump.

Setting the Memory Settings

How to set the heap or permanent generation memory depends on your distribution, platform, and how you start Confluence. Refer to Configuring System Properties.



If you're starting Confluence from a Windows Service, make sure you add the properties through the registry settings.



To verify if your settings have been picked up, check Displaying System Properties.

How to Hide the Referrer

Run Confluence over SSL. Major browsers do not send the HTTP_REFERER when you are linking from an https:// site to an http:// site. To run confluence on SSL please refer to Adding SSL for Secure Logins and Page Security.

Additional Information

Severity	Medium
Regular Expression	[Hh]ide the referrer
Article ID	CONFKB162037794

Searching Confluence Knowledge Base

How to Restore Deactivated Users

Symptoms

A user that has been deactivated in the past cannot be brought back to life. Even if assigned to the right groups with confluence access.



Note that this applies only to the older mechanism for user deactivation. In newer versions of Confluence, users are removed by removing them from groups with can-use permission, as described in Removing a User.

Cause

In older versions of Confluence users have been deactivated by inserting a row into table os_propertyentry with entity_id = theUserId AND entity_key=confluence.user.deactivated
This entry is still evaluated but there is no mechanism in newer version to remove this from the DB.

Resolution

1. Backup Confluence
2. Stop Confluence
3. Find the user id (=theUserId) of the particular user in table users
4. Run the following SQL query to identify the row that contains the 'deactivated' information. The entity_name column should have an entry like 'LOC_username':

```
1.select *  from OS_PROPERTYENTRY WHERE entity_id = theUserId AND entity_key=
"confluence.user.deactivated";
```

5. Delete this row from the table
6. Restart Confluence.

Additional Information

Severity	Medium
Regular Expression	(remove|restore).*deactivated user
Article ID	CONFKB152044814

Searching Confluence Knowledge Base

How to run a SQL script on your database

This document contains some basic instructions on how to run a SQL script on your database. This document is not intended to be exhaustive of the databases we support. We still recommend that you ask your DBA to perform this task if possible.

The following examples assume a database name of `yourdb` and a script file called `myscript.sql`. Of course, the extension of the SQL script file does not have to be `.sql`. Any file can be used so long as it contains SQL statements.

MySQL

```
mysql yourdb < myscript.sql
```

PostgreSQL

```
psql yourdb < myscript.sql
```

How to Search Confluence for Uses of a Macro

Illegal Search Terms

Several terms are illegal to use when searching Confluence, like : or {. If you search for *just* the macro term, you're likely not to get the right search results. For example, searching for 'usage', which is the name of a macro but also a common English word, will not yield an accurate count.

Query the Database

If, for example you want to search for the uses of the `usage` macro, you can search the database:



For large instances of Confluence, searches across the entire bodycontent table of the database might be expensive. You can try this on a test server.



You might want to use the SQL Plugin to automate this, or better yet the Macro usage stats plugin. Note that these are third party (unsupported) plugins.

```
1.SELECT * FROM BODYCONTENT WHERE BODY LIKE '%{usage}%'
```

If you want to get the last users or the creators of the pages where the `noformat` macro is used, you might run:

```
1.SELECT CREATOR FROM CONTENT WHERE CONTENTID IN (SELECT CONTENTID FROM BODYCONTENT WHERE BODY
LIKE '%{noformat}%') GROUP BY CREATOR
2.SELECT LASTMODIFIER FROM CONTENT WHERE CONTENTID IN (SELECT CONTENTID FROM BODYCONTENT WHERE
BODY LIKE '%{noformat}%') GROUP BY LASTMODIFIER
```

Logging Uses of a Macro

This information won't tell you how often a macro is invoked - rather how often it appears on pages. For counting how often it's invoked (and measuring the time it takes to invoke it), check Identifying Slow Performing Macros.

Counting the Incidents of Invoked Logs

You can then grep the logs like:

```
grep -c "{usage" atlassian-confluence.log
```

How to turn on Debugging for indexing

There may be circumstances where you need to turn on debug logging for indexing (e.g. when automatic indexing is not occurring or you are getting errors related to indexing).

To turn on debugging for indexing:

1) Edit your `<CONFLUENCE_INSTALL>/confluence/WEB-INF/classes/log4j.properties` file and add

```
log4j.logger.com.atlassian.confluence.search.lucene=DEBUG
```

2) Restart your instance.

You should see logging in your `<CONFLUENCE_HOME>/logs/atlassian-confluence.log` file, like the following:

```
2009-10-01 08:50:07,633 DEBUG [http-8080-6] [search.lucene.queue.DatabaseIndexTaskQueue] enqueue
Enqueuing task: IndexQueueEntry{id=0, handle='com.atlassian.confluence.pages.Page-91258884',
type=Unknown, creationDate=Thu Oct 01 08:50:07 EDT 2009}
2009-10-01 08:50:07,635 DEBUG [http-8080-6] [search.lucene.queue.DatabaseIndexTaskQueue] enqueue
Enqueuing task: IndexQueueEntry{id=0, handle='com.atlassian.confluence.pages.Page-91258884',
type=Unknown, creationDate=Thu Oct 01 08:50:07 EDT 2009}
2009-10-01 08:50:08,023 DEBUG [DefaultQuartzScheduler_Worker-7]
[search.lucene.queue.DatabaseIndexTaskQueue] getUnflushedEntries Fetching index entries added
since: Thu Oct 01 08:50:06 EDT 2009
2009-10-01 08:50:08,027 DEBUG [DefaultQuartzScheduler_Worker-7]
[search.lucene.queue.DatabaseIndexTaskQueue] getUnflushedEntries Fetched 2 entries from datbase.
2009-10-01 08:50:08,027 DEBUG [DefaultQuartzScheduler_Worker-7]
[search.lucene.queue.DatabaseIndexTaskQueue] getUnflushedEntries Having excluded entries that have
previously been flushed, 2 entries remain.
```

List page- and space-related details for an attachment using the attachment's name

Occasionally, the indexing task will report some attachments that failed to be indexed correctly. If you wish to list the page- and space-related information for these problematic attachments (of which you only know their title from the logs) please query your database with the following SQL statement:

```
1.select s.spacekey, s.spacename, a.title, a.pageid, a.attachmentid,
'(confluence-home)/attachments/' + a.pageid + '/' + a.attachmentid + '/' + a.attversion as filepath
2.from attachments a
3.join content c
4.          on a.pageid = c.contentid
5.join spaces s
6.          on c.spaceid = s.spaceid
7.where a.title like '<Name of Attachment>'
```

Please substitute the attachment name in the above query depending upon your requirement.
The **filepath** column will list all attachments in the directory structure format similar to the way that they are stored in Confluence, such as `<Confluence-Home>/attachments/pageid/attachmentid/attachmentversion`

Related Unix Commands

file <Filename>

If you wish to check the filetype for attachments in your `<Confluence-Home>/atachments/` folder, run the above Unix command.
Eg:

```
file ../data/attachments/32775/98305/1
```

The result is:

```
../data/attachments/32775/98305/1: PNG image data, 1200 x 1000, 8-bit/color RGB, non-interlaced
```

open <Filename>

This command will open files from a shell. By default, opens each file using the default application for that file.
Eg:

```
open ../data/attachments/32775/98305/1 -a /Applications/Preview.app/
```

The result is:
The file is opened in Preview.

Migrate Confluence from one database to another

How do I migrate Confluence from one database to another?

First perform an XML backup of your Confluence site. You can do this from the Administration > Backup and Restore page.

Once you've made the backup file, you can set up a new Confluence instance from scratch against your new database and restore the backup that you just created. Detailed instructions can be found here.

RELATED TOPICS

Migrate to Another Database (Confluence Docs 3.2)

Confluence Configuration Guide (Confluence Docs 3.2)

FAQ Home

Need to Change Server ID for Test Install

Please note this is not essential. If your test server has the same server idea as production it will not affect your production install and is mainly used by support to help us distinguish between your servers.

If you would like to change it, you'll need to acquire a new server ID for your test environment:

1. install a new instance of Confluence (do not use the automatic installer, go with the zip file) and make sure it's pointing to a new empty home directory .
2. Start the confluence that have just been installed, this will create a confluence.cfg.xml in the home directory you specified. Look in this file for the line:

   ```
   1.<property name="confluence.setup.server.id">BIBE-YA0Z-8EK6-SS9C</property>
   ```

3. For the test server, Go to the `<<Confluence-Home>>/confluence.cfg.xml file` and replace the existing server ID with the one you've just obtained in the new installation of Confluence.
4. Start up the test server, go to `Administration > License Details` and update the license to your developer license that you generate from https://my.atlassian.com.

Page Restrictions Performance Considerations

Page Permissions and the Confluence Search Index

Modifying the page permissions requires reindexing all child pages, as well as comments and attachments on all those pages.

The reasons for this:

- Page permissions are stored on every item in the index
- When you search, a filter is applied to all items in the index which prevents you seeing content you don't have permission to see
- If the permission information on attachments wasn't updated when a page permission was changed, users who didn't have permission to view the attachment in Confluence would still see that attachment in the search results as well as some of its content
- To update any record in the index, you need to delete it from the index and re-add it (this is a limitation of Lucene)

Therefore, to update an attachment record in the index, even just to change the permissions, the attachment's content must be reindexed.

Performance Considerations

In the large majority of situations, this design should not be a problem. In large spaces with deep page hierarchies, it might be. Use performance logging for the index flush operation to assess the impact of changing a page restriction - try it on the space's home page, or a page with a lot of children, to see the performance impact of changing a page restriction.

Space Design Considerations

For some space designs, deep page hierarchies may be desirable. If possible, it's recommended to split spaces where it makes sense to do so, according to how your information is organized. One workaround – CONF-7089 – involves opening up access to just a few pages in a mostly-restricted space so as to "open" the space where space permissions close it. This may be a performance concern if the space, and attachments in the space, are large.

Workarounds

There are a few workarounds to consider:

1. Avoid page restrictions on large page hierarchies. If you have a large hierarchy you have to protect differently to other content in the space, consider moving the hierarchy to a new space. (Space permissions are applied in a manner that doesn't have this problem, but having a large number of spaces also causes scaling issues.)
2. Disable page permissions completely. You can do this on a space-by-space basis by not granting the "Restrict" permission in the space permissions screen.
3. Turn off attachment content indexing. See How do I disable indexing of attachments or Configuring Attachment Size for information on how to do attachment types or size limits.

Which Pages Have Restrictions?

A database query to show which pages in your instance have page restrictions:

```
1.SELECT DISTINCT content.contentid, content.TITLE as page_title, CONTENT_PERM.USERNAME as
page_restriction_username, CONTENT_PERM.CP_TYPE as page_restriction_type FROM CONTENT_PERM,
content WHERE CONTENTID IN (SELECT CONTENT_ID FROM CONTENT_PERM_SET WHERE ID IN (SELECT CPS_ID
from CONTENT_PERM))
```

For page permissions:

```
1.select * from CONTENT_PERM CP inner join CONTENT_PERM_SET CPS on CPS.ID = CP.CPS_ID inner join
CONTENT C on C.CONTENTID = CPS.CONTENT_ID where CP.GROUPNAME is not null;
```

For space permissions:

```
1.select distinct S.SPACEID, S.SPACENAME,PERMGROUPNAME from SPACEPERMISSIONS SP inner join SPACES
S on S.SPACEID = SP.SPACEID where PERMGROUPNAME is not null and PERMGROUPNAME <> 'confluence-users'
;
```

Passing username and password via URL

Confluence uses Seraph Authenticator which supports four methods of authentication. Among them is passing request parameters: `os_username` and `os_password`.

This allows one to login with user name and password in the URL by adding **?os_username=yourUserName&os_password=yourPassword**.

For example: **http://<Confluence url>/dashboard.action?os_username=yourUserName&os_password=yourPassword**

FAQ Home

Preventing and Cleaning up SPAM

If you have a public facing confluence site, your site may be affected by spammers.

Preventing Spammers

To prevent spammers you will have to:

1) Enable captcha
2) Run confluence behind an Apache Webserver and create rules to block the spammers IP address.

Blocking Spam at Apache or System Level

If a spam bot is attacking your Confluence site, chances are they are coming from one IP or a small range of IPs. To find the attacker's IP, it helps to 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
```

You should correlate actual spam users being created with the log entries to make sure you do not block legitimate users. By default, Apache logs the clients IP in the first field of the log line.

Once you have the offender's IP or IP range, you can add it to your firewall's blacklist. For example the popular Shorewall firewall for Linux you can simply do:

```
# echo "1.2.3.4" >> /etc/shorewall/blacklist
# /etc/init.d/shorewall reload
```

To block at the Apache level, you can update your Apache vhost config with the line:

```
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

1) Turn off public sign up
2) See CONF-1469. Your comments and vote on that issue are very much appreciated.

Deleting Spam

Profile Spam

This refers to spammers creating accounts on Confluence wikis and posting links to their profile page. This is a particularly common form of spam at the moment.

If you have had many such spam profiles created, it is easier to delete them via SQL.



Shutdown Confluence and backup your DB before doing this!

Find the last real profile

```
SELECT bodycontentid,body FROM bodycontent WHERE contentid IN
  (SELECT contentid FROM content WHERE contenttype='USERINFO')
  ORDER BY bodycontentid DESC;
```

Look through the bodies of the profile pages until you find where the spammer starts. You may have to identify an number of ranges.

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 bodycontendID <= 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);
```

Once the spam has been deleted, restart Confluence and run a rebuild of the index - which will remove any references to the spam from the search index.

Rebuild the Content Indices from scratch

Why are the Content Indices failing?

See Troubleshooting Searching and Indexing.

Flushing the Index Queue

You may be able to address the issue without a complete rebuild of the index. When you experience search problems within your Confluence instance, or cannot find things after performing an upgrade, one simple solution that can help you locate your content is to manually rebuild the index, as described in Content Index Administration. You can try either a complete rebuild, or flush the queue contents.

Rebuilding the Failed Content Indices

If you are still experiencing problems after performing the above rebuild, the next step may be to remove the index and rebuild it from scratch, as described below. For example, the rebuild procedure described above may not restore a corrupted index file. You may want to ensure that all index files are re-created.

The Space Activity feature uses the index to store data. To preserve the activity data, keep the Plugins directory and delete the rest of the index.

To remove the index and rebuild it from scratch:

1. Shut down your Confluence server.
2. Make a backup of your `<confluence-home>/index/plugins` directory if it exists. This is where the confluence usage tracking plugin stores its index for the usage statistics and it cannot be rebuilt.
3. Remove the `<confluence-home>/index` directory.
4. If in step 2, you have the `<confluence-home>/index/plugins` directory, create the `<confluence-home>/index` directory and copy in the backup of `<confluence-home>/index/plugins` directory made in step 2.
5. Restart server.
    A new **index** directory will be created upon restart of your application.
6. Now perform the manual re-indexing of your site to build a brand new index from scratch.

If you still cannot find your content, please contact confluence support.

Redirect users to a page on login

Can I set Confluence to redirect users to a space homepage immediately after login?

Yes. You can direct users to any of the space homepages when they login to the site This is configured by a site administrator from the **Administration Console**. Instructions on how to do this can be found here.

You can also modify the login page's content or establish single sign-on integration.

RELATED TOPICS

Administration Guide
Customising Space Homepage

Redirect users to a site-wide home page after a successful login

As an administrator, you can set a site-wide home page within any space, to which users are sent after logging in.

To set the site-wide home page as an administrator:

1. Go to Administration > General Configuration
2. Click Edit
3. Set **Site Homepage** to your desired home page

Setting home pages within spaces depends on the access permissions to both the space and your site.

- If your site allows anonymous access, the site homepage must also be accessible anonymously.
- If your site does not allow anonymous access, the site homepage must be accessible to the "confluence-users" group.

In Confluence 1.2 and earlier, the site homepage must be accessible anonymously, regardless of site permissions.

> **Note:** please ensure that the *View Space Goes to Space Summary* setting is set to OFF if you want users to be sent to space **homepage** instead of the **summary page**

For instructions on configuring this feature at the user level, see Redirect to a specific page (home page) within the site after logging in.

RELATED TOPICS

Redirect to a specific page (home page) within the site after logging in

Restrict Attachments Based On File Type

If a user has permission to upload attachments, you can restrict them to certain filetypes only by installing the Attachment Filter. Note that this is an unsupported plugin and has only been tested to work on Confluence 2.2 - 2.3.x. There is an open new feature request to have this updated and supported here.

You can also restrict the maximum size of any one attachment that a user can upload, or disable the indexing of attachments, which may help performance.

Search for User Properties in the Database

So long as you have not changed the user authentication from the default (i.e. you are NOT using external user management such as LDAP or JIRA) the users of confluence are stored in `OS_USER` table.

The user's properties such as fullname, email and when they previously logged in, are stored in the `OS_PROPERTYENTRY` table.

If you know the user's username (which can be found in `OS_USER` table), you can find their details using the following query:

```
select * from os_propertyentry p, os_user u where p.entity_id = u.id and
u.username='user_name_goes_here';
```

Using Firebug Lite in Internet Explorer when browsing a Confluence page

1. Open a Confluence page in IE.
2. Copy and paste the following into your IE browser URL bar and press enter:

```
1.javascript:var%20firebug=document.createElement('script');firebug.setAttribute('src','
http://getfirebug.com/releases/lite/1.2/firebug-lite-compressed.js'
);document.body.appendChild(firebug);(function(){if(window.firebug.version){firebug.init();}
else{setTimeout(arguments.callee);}})();void(firebug);
```

3. You should see a Firebug console at the bottom of the browser now.

For more information, please refer to Firebug Lite's documentation

What are the IP Address Ranges for Atlassian's Servers?

Various functions like the automatic support request and plugin repository require access to Atlassian Servers. The IP address range to configure for your firewall are as follows:

63.246.22.32/27

63.246.22.192/27
67.221.237.0/27

Where are the files that used to be in my Confluence installation directory?

With the release of Confluence 2.6.0, many files have been moved inside **`WEB-INF/lib/confluence-x.x.x.jar`** file, which means they can no longer be edited by simply opening and saving a file.

For example, you may find that some files previously were in **`<confluence install directory>/confluence/WEB-INF/classes`** are no longer there.
If you want to edit them for customisation, you will need to extract the relevant files from the **`confluence-x.x.x.jar`**, and place them in the directory where they are used to reside.

You can either use a zip application or use the `jar` tool from your JDK installation to extract the file.



Example
You want to modify **`xwork.xml`** in Confluence 2.10.2. The steps to follow are:

1. Shutdown Confluence
2. Locate **`WEB-INF/lib/confluence-2.10.2.jar`**.
3. Open the jar file using your favourite zip application and search for **`WEB-INF/lib/confluence-2.10.2.jar/xwork.xml`**. You can also use JAVA's [{{jar}}] tool if you like.
4. Extract **`xwork.xml`** and place it in **`WEB-INF/classes/`**, and modify as necessary. Any files placed in this folder will take precedence over their jarred version and you do not need to re-jar them in **`confluence-21.0.2.jar`**.
5. Restart Confluence


The problem with customization is that they might break in the new version of Confluence. Some codes in your customized file might have changed in the new version. Hence, it is best that you do not copy your customized file directly to the new installation directory. Instead, you need to apply the same customization in the new version of the file.

RELEVANT TOPICS

Editing Files within JAR Archives
Installing Patched Class Files

Where are user macros stored?

User macros are stored in the bandana table:

```
1.select * from bandana as b where b.bandanakey = 'atlassian.confluence.user.macros';
```

Backup FAQ

This section contains solutions for common issues or queries associated with backing up the content within your Confluence site or installation.


The XML backup is known to be inefficient and prone to errors with larger instances. You can [switch to an external backup process] for a reliable and efficient solution.

View one of the following issues or queries for more information:

- Are there any scripts for backup creation and restore?
- Backup will not import
- Can Confluence be restored from a backup minus attachments?
- Can XML backups be deleted automatically?
- Does running a daily XML backup slow performance?
- How can I reduce the space taken up by automatic backups?
- How to Change the Version of a Space Backup
- How to Find Attachments in Attachments Folder
- Is it Possible to Store the Confluence Home Directory on a Network Share?
- Providing MySQL database with Content Anonymised

RELATED TOPICS

Site Backup and Restore

Are there any scripts for backup creation and restore?

Check out User Submitted Backup & Restore Scripts.

Backup will not import

See Troubleshooting failed XML site backups.

Can Confluence be restored from a backup minus attachments?

Yes. First, ensure you have created a site backup that includes the attachments. Then, restore this site backup. Refer to Site Backup and Restore for more information.

Can XML backups be deleted automatically?

Windows users must manually delete any backup files. Linux users can insert a nightly or weekly automation script or cron similar to the following:

```
ls -t <path to your backup dir>/* | tail +6 | xargs -i rm {}
```

Does running a daily XML backup slow performance?

The XML backup is known to be inefficient and prone to errors with larger instances. You can switch to an external backup process for large instances.

How can I reduce the space taken up by automatic backups?

Switch to a manual backup process according to the 'Backups For Large Instances' section of Site Backup and Restore, which will give you more control over disk usage.

How to Change the Version of a Space Backup

Confluence prevents the import of space backups which aren't from the same major version. The reason for this is that any schema change between the export and imported version of Confluence will cause the import to fail, leaving you with an incomplete import. Even worse, the failure can be database-dependent, so it may work fine on one particular database but your backup will fail to import later.



Do not import a modified space backup on a production server. Import the modified space backup on a test server, then export from the test server to create a pristine space backup for the new version.

To change the version of a space backup, do the following:

- extract the space backup ZIP file
- edit `exportDescriptor.properties` in a text editor
- change the buildNumber to the buildNumber of the Confluence version you wish to import into
- zip up the modified contents of the backup into a ZIP file again.

This will allow you to import a backup into a test instance of Confluence. After checking the imported space for errors, export it cleanly from the test server and import the fresh backup into your production server.

If your import fails on the test server due to Hibernate errors, this indicates a schema incompatibility and cannot be worked around. You will need to restore your entire site on an old version of Confluence, and export the space from there. See the last section of Restoring a space for details.

How to Find Attachments in Attachments Folder

Symptoms

Attachments are stored on filesystem but there isn't any attachment in the attachments folder.

Diagnosis

You can determine the missing attachments by using the Missing Attachments Report.

Cause

The attachment naming scheme is numerical so as to avoid encoding problems with operating systems.

Resolution

You can look for the attachment detail (e.g file type, attachment name) from ATTACHMENTS table.

```
1.Select * from ATTACHMENTS where pageID='<PageID>'
```

As attachment is stored as this structure:
Confluence 2.x:

```
1.<attachments>/pageID/attachment/attachmentVersion, you may want to run the following query to
retrieve attachments of a page:
```

Confluence 3.x:
See Hierarchical File System Attachment Storage.

Is it Possible to Store the Confluence Home Directory on a Network Share?

Is it possible to house Confluence Home/Confluence on a NAS device instead of local drives?

It is possible to set up this configuration. To do so, specify the network location from <confluence-install>/confluence/WEB-INF/classes/confluence-init.properties. Atlassian does not suggest installing Confluence or hosting Confluence Home directory on NAS device because when a NAS or connection to NAS is down, Confluence cannot function correctly and you risk potential data corruption.

Providing MySQL database with Content Anonymised

Particularly for indexing issues, it is useful to provide a database backup to Support so that they can reproduce the issue.

However, some of your content may be private, so to anonymise it:

1) Take the mysql dump of your confluence database:

```
mysqldump -u username -ppassword database_name > FILE.sql
```

2) Load it into a test database

```
mysql -u username -ppassword test_database < FILE.sql
```

3) Run the following query against your test database:

```
1.update BODYCONTENT set BODY='a';
2.update USERS set PASSWORD=
'x61Ey612Kl2gpFL56FT9weDnpSo4AV8j8+qx2AuTHdRyY036xxzTTrw10Wq3+4qQyB+XURPWx1ONxp3Y3pB37A=='
```

This will update all content on pages to "a"

This only anonymizes the data on pages, comments and blog posts and user passwords. It does not anonymise the titles of pages, usernames or labels.

Configuration FAQ

This section contains solutions for common issues or queries associated with configuring Confluence.

This section focuses on providing instructions to customise Confluence's functionality and appearance by modifying its installation.

View one of the following issues or queries for more information:

- Running Tomcat on a Different Port
- Change default font, color, or spacing in Confluence
- How do I Change the Time of Daily Report Updates
- How to audit Confluence - enabling user access logging
- Share users and groups between Confluence and JIRA
- How do I Modify the Frequency of Content Indexing
- How do I change the default polling time for email in Confluence?
- How do I Configure an Automatic Refresh of the Recently Updated List
- How to Revert from Clustering to Single Node
- Adding a Site-Wide Banner
- Customise Confluence Page Exports
    - Available Velocity Context Objects in Exporters
    - Customise MS Word Exports
    - Customise PDF Exports

- Editing the Footer
- How do i completely remove the "Space Details" page from Confluence exports?
- Where does Confluence store all its data?
- Running Confluence Behind a Caching Proxy Server
- I am trying to install Confluence but the demo-site.zip file is missing
- How do I Disable Automatic Mail Polling?
- Disabling Profile Pictures on the Recently Updated Dashboard
- How to Disable Profile Pictures from the Recently Updated Section of the Dashboard
- Remove Version from Footer
- Configuring a MySQL Datasource in Apache Tomcat

Running Tomcat on a Different Port

By default Confluence's Tomcat instance listens on port **8080** for connections, and port **8005** for the command to shut down.

If you see errors like:

```
1.java.net.BindException: Address already in use:8080
```

in your logs, then you need to change the port number Confluence listens on (or you may be trying to start Confluence twice).

Confluence WAR/EAR distribution

Modify Tomcat conf/server.xml, change the ports used on the Server tag and HTTP Connector tag. If you don't know what this means, see the example in the Standalone configuration below.

Confluence Standalone 2.2 and newer

Edit conf/server.xml in your Confluence application directory, and change the ports on the Server tag and the Connector tag to unused ports on your system. The following example shows the relevant section of server.xml where they are changed to 8100 (Server) and 8180 (HTTP Connector). In this configuration you would access Confluence on the URL: http://localhost:8180/confluence.

```
1.<Server port="8100" shutdown="SHUTDOWN" debug="0">
2.          <Service name="Tomcat-Standalone">
3.                    <Connector className="org.apache.coyote.tomcat4.CoyoteConnector" port="8180"
minProcessors="5"
4.                                                  maxProcessors="75"
5.                                                  enableLookups="true" redirectPort="8444"
acceptCount="10" debug="0" connectionTimeout="20000"
6.                                                  useURIValidationHack="false" URIEncoding="UTF-8"/>
```

Confluence prior to 2.2

You can change the port number by editing the file `<confluence install directory>/conf/server.xml`.

Find the line:

```
1.<Connector className="org.apache.coyote.tomcat4.CoyoteConnector" port="8080" minProcessors="5"
maxProcessors="75"
```

and change **8080** to the port you want to use.

If you are running two Tomcat instances, you'll also need to change the shutdown port in the line:

```
1.<Server port="8005" shutdown="SHUTDOWN" debug="0">
```

Specifying the IP address to listen on

For versions 2.2 and later, edit `conf/server.xml` as described above, but add an `address` parameter. The tomcat documentation is a useful reference here. To extend the example above, you can configure tomcat to listen only on the localhost interface with this configuration:

```
1.<Server port="8100" shutdown="SHUTDOWN" debug="0">
2.          <Service name="Tomcat-Standalone">
3.                    <Connector className="org.apache.coyote.tomcat4.CoyoteConnector" port="8180"
minProcessors="5"
4.                                                  maxProcessors="75"
5.                                                  address="127.0.0.1"
6.                                                  enableLookups="true" redirectPort="8444"
acceptCount="10" debug="0" connectionTimeout="20000"
7.                                                  useURIValidationHack="false" URIEncoding="UTF-8"/>
```

Why would you want to restrict it to access via localhost? If you're running Confluence behind a proxy server on the same host, this will make sure that users can't bypass the proxy and hit the application server directly.

If your machine also runs IIS, you might encounter a problem where IIS binds to the same port (for example, port 80) for all IP addresses on the machine. This means that you will not be able to run Confluence on another web or application server through that port until you address this IIS issue first. For more information, refer to this Microsoft Knowledge Base article.

Change default font, color, or spacing in Confluence

Beginning in Confluence 2.10, you can customise your space or instance using CSS from the User Interface. See Styling Confluence with CSS for details. For earlier versions, check the instructions here.

How do I Change the Time of Daily Report Updates

You can configure Confluence to perform the daily updates at a time that is best suited to you or your organisational needs.

Time is derived from the Confluence server
The time zone is taken from the server on which Confluence is running. To check the time according to the server, do the following:

1. Go to the Confluence '**Administration Console**'. To do this:
   - Open the '**Browse**' menu and select '**Confluence Admin**'. The 'Administration Console' view will open.
2. Click '**System Information**' in the left-hand panel and look at the '**System Time**'.

Confluence uses Quartz for scheduling periodic jobs. To change the time of your daily report, you will need to edit the Quartz configuration.

To change the time of your daily reports

1. Open the Quartz configuration file `schedulingSubsystemContext.xml` located under `confluence/WEB-INF/lib/confluence-x.x.x.jar`. Where x.x.x is your Confluence version number.
2. Find the following section of the file:

```
1.<bean id="dailyReportTrigger" class=
"org.springframework.scheduling.quartz.CronTriggerBean">
2.                      <property name="jobDetail">
3.                                <ref bean="dailyReportJob"/>
4.                      </property>
5.                      <property name="cronExpression">
6.                                <value>0 0 0 * * ?</value>
7.                      </property>
8.          </bean>
```

3. The string `'0 0 0 * * ?'` sets up a Cron Trigger for the job to run at the zeroth second of the zeroth minute of the 0th hour (midnight), every day of every month, every day of the week.
4. Re-jar the file, either with a zip utility (change the title of .zip back to .jar) or a java command.
5. You can set a new time by editing this string. Note that the date and time format in this configuration file is in this order:

   Second minute hour day

6. Restart Confluence.

For example, to set the new time to twenty past ten PM, change the string to `'0 20 22 * * ?'`.
For complete details on the formatting of the cron string, please see http://www.opensymphony.com/quartz/api/org/quartz/CronTrigger.html.

RELATED TOPICS

- Changing time of Daily Backup
- How do I Modify the Frequency of Content Indexing

How to audit Confluence - enabling user access logging

Often, for auditing purposes, administrators need to know who did what. Notifications are not ideally suited for this purpose. Instead, you can generate a basic log indicating which users are accessing which pages in Confluence. Application servers are able to log the requested URL, but they cannot determine the currently logged in user. This log is not currently formatted to be accessible to web log analysis tools such as AwStats as it lacks a host and get method, so must be viewed manually.

Similar to JIRA, Confluence has a built-in access logging mechanism, which shows the user and URL invoked. To enable it, you need to modify a couple of configuration files and restart Confluence.

Configuring the AccessLogFilter

There is a simple AccessLogFilter in Confluence than can be enabled via `confluence/WEB-INF/classes/log4j.properties` and `confluence/WEB-INF/web.xml`.



Please do not modify the application-wide web descriptor, `$server/conf/web.xml`. This will be ineffective and potentially may break Confluence.

1. Uncomment these line in log4j.properties:

```
1.log4j.category.com.atlassian.confluence.util.AccessLogFilter=INFO
```

2. Enable the filter in web.xml by removing the comments around these lines:

```
1.<filter-mapping>
2.          <filter-name>AccessLogFilter</filter-name>
3.          <url-pattern>/display/*</url-pattern>
4.          <url-pattern>*.action</url-pattern>
5.</filter-mapping>
```

   Notice that the *.action pattern is added optionally to log the actions of Confluence in addition to the page views. This combination of URL patterns will work for all URLs. You can further modify the pattern by adjusting the url-pattern field.



For troubleshooting purposes, often it is useful to capture all accesses to Confluence. To do this use this filter mapping in `web.xml` instead of the above:

```
1.<filter-mapping>
2.        <filter-name>AccessLogFilter</filter-name>
3.        <url-pattern>/*</url-pattern>
4.</filter-mapping>
```

3. Restart Confluence

This will result in logging information being stored in the atlassian-confluence.log file in the confluence-home directory.

Advanced configuration

After this is working, you could redirect the access log to a **different** file by adding a new RollingFileAppender at the top of log4j.properties:

```
log4j.appender.accesslog=org.apache.log4j.RollingFileAppender
log4j.appender.accesslog.Threshold=DEBUG
log4j.appender.accesslog.File=${catalina.home}/logs/atlassian-confluence-access.log
log4j.appender.accesslog.MaxFileSize=20480KB
log4j.appender.accesslog.MaxBackupIndex=5
log4j.appender.accesslog.layout=com.atlassian.confluence.util.PatternLayoutWithStackTrace
log4j.appender.accesslog.layout.ConversionPattern=%d %p [%c{4}] %M %m%n
```

Find this line:

```
#log4j.category.com.atlassian.confluence.util.AccessLogFilter=INFO
```

Change it to this:

```
log4j.category.com.atlassian.confluence.util.AccessLogFilter=INFO, accesslog
log4j.additivity.com.atlassian.confluence.util.AccessLogFilter=false
```

The web.xml url-pattern given above only matches page views (/display/*). You could change the url-pattern, or duplicate the entire filter-mapping to log access for different kinds of access (/admin/* for admin functions, /pages/* for edits and creates, etc. Note that /pages/editpage.action* doesn't work).

What is logged

The format produced is the following values separated by spaces:

1. Username or '-' if no user
2. URL
3. VM free memory at start of request (in KB)
4. Change in free memory after request is finished (in KB)
5. Time taken for request (in ms).
6. Remote address

Example:

```
2008-08-08 10:33:05,359 INFO [atlassian.confluence.util.AccessLogFilter] init AccessLogFilter
initialized. Format is: <user> <url> <starting memory free (kb)> +- <difference in free mem (kb)>
<query time (ms)> <remote address>
2008-08-08 10:47:27,015 INFO [atlassian.confluence.util.AccessLogFilter] doFilter admin
http://localhost:8080/display/ds 42025-154 15 127.0.0.1
2008-08-08 10:47:27,187 INFO [atlassian.confluence.util.AccessLogFilter] doFilter admin
http://localhost:8080/display/ds/Confluence+Overview 41805+982 172 127.0.0.1
2008-08-08 10:47:36,296 INFO [atlassian.confluence.util.AccessLogFilter] doFilter admin
http://localhost:8080/display/ds/Breadcrumb+demonstration 42102-6660 156 127.0.0.1
2008-08-08 11:08:16,875 INFO [atlassian.confluence.util.AccessLogFilter] doFilter admin
http://localhost:8080/display/ds/test+firelite 34362-1616 188 127.0.0.1
2008-08-08 11:47:01,890 INFO [atlassian.confluence.util.AccessLogFilter] doFilter admin
http://localhost:8080/display/sand 59711-148 0 127.0.0.1
2008-08-08 11:47:02,171 INFO [atlassian.confluence.util.AccessLogFilter] doFilter admin
http://localhost:8080/display/sand/Home 59497-2302 234 127.0.0.1
2008-08-08 11:47:04,500 INFO [atlassian.confluence.util.AccessLogFilter] doFilter admin
http://localhost:8080/display/ds/Tasklist 57124+155 1266 127.0.0.1
```

The above may be preceded by additional log4j-generated text, depending on the log4j pattern which is configured.

Another option: Google Analytics

Google Analytics can be easily integrated with Confluence for access tracking.

After signing up, copy the Javascript and paste it into the 'Before end of <body>' section of Administration, Custom HTML. This will put the Javascript on every page generated by Confluence.

This might not work correctly if your users are behind a firewall or authenticated proxy.

For more information on using Google Analytics with Confluence you may wish to refer to this blog post by David Simpson.

RELATED TOPICS

Working with Confluence Logs

Share users and groups between Confluence and JIRA

How do I share users and groups between Confluence and JIRA?

To save your having to enter users into both JIRA and Confluence, you can configure Confluence to use JIRA's user database.

Instructions on how to do this can be found here.

RELATED TOPICS

Migrating users from Confluence to JIRA (Confluence Docs 3.2)

Revert from JIRA to internal user management (Confluence Docs 3.2)

Showing first 2 of 3 results

How do I Modify the Frequency of Content Indexing

Confluence Content Indexing frequency is handled using a cron job set in `schedulingSubsystemContext.xml`.

Time is derived from the Confluence server
The time zone is taken from the server on which Confluence is running. To check the time according to the server, do the following:

1. Go to the Confluence '**Administration Console**'. To do this:
    - Open the '**Browse**' menu and select '**Confluence Admin**'. The 'Administration Console' view will open.
2. Click '**System Information**' in the left-hand panel and look at the '**System Time**'.

Confluence uses Quartz for scheduling periodic jobs. To change the time of your content indexing, you will need to edit the Quartz configuration.

To change the time of your content indexing

1. Open the Quartz configuration file `schedulingSubsystemContext.xml` located under `confluence/WEB-INF/lib/confluence-x.x.x.jar`. Where x.x.x is your Confluence version number.
   For Confluence earlier than 2.6, the index cron job is located in `<install dir>\confluence\WEB-INF\classes\schedulingSubsystemContext.xml`.
2. Find the following section of the file:

```
1.<bean id="indexQueueFlushTrigger" class=
"org.springframework.scheduling.quartz.CronTriggerBean">
2.          <property name="jobDetail">
3.                    <ref bean="indexQueueFlushJob"/>
4.          </property>
5.          <property name="cronExpression">
6.                    <value>0 0/5 * * * ?</value>
7.          </property>
8.</bean>
```

3. The string `'0 0/5 * * * ?'` sets up a Cron Trigger for the job to run every 5 minutes.
4. Re-jar the file, either with a zip utility (change the title of .zip back to .jar) or a java command.
5. You can set a new time by editing this string. Note that the date and time format in this configuration file is in this order:

   Second minute hour day

6. Restart Confluence.

For example, to set the new time to twenty past ten PM, change the string to `'0 20 22 * * ?'`.

For complete details on the formatting of the cron string, please see http://www.opensymphony.com/quartz/api/org/quartz/CronTrigger.html.

How do I change the default polling time for email in Confluence?

Modify the **cronExpression** property of the mailPollTrigger bean in `schedulingSubsystemContext.xml` file.

In Confluence installation directory, this file is jarred in **<confluence installation directory>/confluence/WEB-INF/lib/confluence-x.x.jar**.
Unzipping the file and placing it **WEB-INF/classes** folder will over ride the configuration from the jarred version.

For example, the expression below will fire at 12pm (noon) every day:

```
1.<property name="cronExpression">
2.     <value>0 0 12 * * ?</value>
3.</property>
```

Cron expression
You may also like to view an information on how to create Cron expressions to provide the ability to specify complex time combination.
http://quartz.sourceforge.net/javadoc/org/quartz/CronTrigger.html

How do I Configure an Automatic Refresh of the Recently Updated List

To have the dashboard refresh automatically:

1. Modify the Main Layout at Administration->Layouts->Main Layout.
2. Add in the <META HTTP-EQUIV="REFRESH" CONTENT="5"> tag in the html header tag
3. This example will refresh the browser in every 5 seconds

```
1.<html>
2.<head>
3.<META HTTP-EQUIV="REFRESH" CONTENT="5">
4.....................................................
```

How to Revert from Clustering to Single Node

Disabling a Cluster

If reverting from a Clustered configuration, one solution is to back up the data, install a new stand-alone, and restore the data. This is the recommended approach. An alternative is:

1. Download the standard (non-clustered) distribution from the Confluence Downloads
2. Point confluence-init.properties to the existing confluence-home directory.
3. Set in confluence.cfg.xml:

```
<property name="confluence.cluster">false</property>
```

To check to see if clustering has been disabled, look in the logs after the xml during startup. In your catalina.out, you have:

```
INFO [KB:main] [KB:confluence.cluster.tangosol.TangosolClusterManager] startCluster Bringing up
cluster service
```

This line won't exist if you start it up with the config we gave above. That's how to test it.

Adding a Site-Wide Banner

Confluence administrators can add a site-wide banner, i.e. a message or alert that will appear at the top of every page on your Confluence site.

Screenshot: Example of a Site-Wide Banner



To add a site-wide banner,

1. Go to the Confluence '**Administration Console**'. To do this:
    - Open the '**Browse**' menu and select '**Confluence Admin**'. The 'Administration Console' view will open.
2. Click '**Custom HTML**' in the '**Look and Feel**' section of the left-hand panel.
3. Click '**Edit**'.
4. Add the following code to the '**At end of the HEAD**' textbox.

```
1.<!-- Message Banner -->
2.<div style="background-color: yellow; border: 2px solid red; margin: 4px;
padding: 2px; font-weight: bold; text-align: center;">
3.Your important message...
4.</div>
```

5. Click '**Save**'.

If you want the banner across the bottom of the page, you should add the code to the '**At end of the BODY**' textbox instead.

RELATED TOPICS

Customising Look and Feel Overview

Customise Confluence Page Exports

Modify the style or content of the following page exports:

- Available Velocity Context Objects in Exporters
- Customise MS Word Exports
- Customise PDF Exports

Available Velocity Context Objects in Exporters

Since the export functionality is not implemented as a WebWork action, it does not inherit the Velocity context of an action. It creates its own context and populates it with a separate list of components. All exporters will inherit this context.

Velocity Reference	Description	Javadoc Reference
$generalUtil	A `GeneralUtil` object with several useful methods, including URL encoding.	GeneralUtil
$textUtil	Common utilities for string manipulation.	TextUtils
$rendererBean	Mostly for internal use, but can also be used for manipulating page or space exports.	WikiExporter
$exportDate	A `java.util.Date` created when the export is actually performed.	Date
$exportContent	An `ExportContext` object which holds the context in which the export is performed, such as the user performing the export.	ExportContext
$dateFormatter	Provides a date and time formatter suitable for the exporting user's locale and environment.	DateFormatter
$baseUrl	The base URL of the Confluence installation (`http://<server>:<port>/contextPath`).	String

Specific exporters may have additional context members. See their pages for details.

Customise MS Word Exports

Overview

Confluence exports a Confluence page as a Microsoft Word document by:

1. Generating the HTML rendering of a page from Confluence wiki markup
2. Overriding some Confluence page styles with MS Word styles using a CSS wrapper

The wrapper is generated using a Velocity macro to provide CSS style information specific to MS Word exports.

Modifying Content

Please do not attempt to modify the output without some experience in CSS and HTML markup. You may also need to review Velocity template language.

- W3Schools CSS Tutorial
- W3Schools HTML Tutorial
- Velocity Template Overview

Before modifying any styles, you should always take a backup of both any files you are modifying and your entire Confluence install directory. If you require assistance with your customisation attempt, you should post your modified files along with a technical description to our Confluence Developer Forum.

Important Files

The formatting is defined using the CSS styles in main Confluence stylesheet and overridden by any styles in the Word export wrapper.

Main Stylesheet

This file formats the default appearance of all Confluence content. Changes to this document will be shown when viewing a page from within Confluence, as well as in HTML or MS Word exports. You should avoid changing this document unless you wish a style to be changed throughout Confluence. The file is located under the Confluence install directory:

```
1....\confluence\WEB-INF\classes\styles\site-css.vm
```

Export Wrapper

This file overrides the main Confluence stylesheet, so styles outlined in this file will be specific to MS Word exports. It set the font and style for body text, some macro and grids, and a default background colour. The file is located under the Confluence install directory:

```
1....\confluence\pages\exportword.vm
```

How To Modify Styles

If the CSS tag already exists in the Word export wrapper, you can modify it directly. If the CSS tag is defined by the main stylesheet, copy it into the export wrapper as a new entry in the `<style>` element and update it there:

1. Identify its CSS tag in the main stylesheet. You may need to use trial and error to identify which cascading property is relevant
2. Transfer the tag into the Word export wrapper so that it overrides the main stylesheet
3. Modify the tag style in the wrapper

Common Modifications

You can easily set the body font or background colour in the export wrapper.

Modify Body Font

The wrapper sets the body font:

```
1.body, p, td, table, tr, .bodytext, .stepfield {
2.          font-family: Verdana, arial, sans-serif;
3.}
```

Modify Background Colour

The wrapper sets the background colour:

```
1.<body style="background-color: white; padding: 10px;">
2.          <h1>$page.title</h1>
3.          $renderedPageContent
4.</body>
```

Customise PDF Exports

To customise Confluence's PDF output, you can edit the CSS stylesheets used by the PDF exporter. See Editing the PDF Stylesheet.

Editing the Footer



'Powered by Atlassian Confluence'
Atlassian requires that the text 'Powered by Atlassian' be displayed in the footer on every page, as specified in the license agreement. Other than that, you can customise the footer text.

To change the footer text, follow the instructions in Modify Confluence Interface Text. You can specify additional configurations in `/confluence/decorators/includes/footer.vmd`.

If you need to revert to a former version, for example to restore the 'Powered By Atlassian Confluence' text, you can refer to the attached footer file.

RELATED TOPICS

Modify Confluence Interface Text
Customising Look and Feel Overview

How do i completely remove the "Space Details" page from Confluence exports?

1. Export to html file extension, and customize the layout from Administration->Layouts->Export Layouts->Space Export Layout
2. Export to PDF file extension, take a look at WEB-INF\classes\com\atlassian\confluence\spaces\Space.pdfexport.vm

Where does Confluence store all its data?

Attachments, extensions and configuration files are stored in the **Confluence Home Directory** that is configured when Confluence is first installed. All remaining data resides in the configured database.

See Important Directories and Files for more information.

RELATED TOPICS

 Important Directories and Files (Confluence Docs 3.2)

FAQ Home

Running Confluence Behind a Caching Proxy Server

> *One major concern is Confluence's ability to withstand a Slashdot, and someone told us that Atlassian had basically said that Confluence could not handle the load of such an event/attack.*
>
> *Ideally I would want to put a Squid cache directly infront of Conflunce, set the default policy to cache content of **normal pages** for ~5 minutes (at least) and then pass-through more of the dynamic pages (like the editor & such).*

This is, in fact, the case. We don't have any deployed Confluence sites that have the requirement of being Slashdot-proof, but this is probably one of those chicken-and-egg things.

The problem is not one of simple scaleability. We're currently working on "Confluence Massive", a clusterable Confluence that will scale to handle whatever load you feel like throwing at it. But if your aim is to protect the server against sudden, transient loads, throwing a cluster at the problem that will then spend 99% of its time not being utilised is probably a waste. Thus, the best solution is to have some kind of caching reverse-proxy that will divert load away from Confluence itself.

The main problem with the reverse-proxy solution is that every Confluence page is built dynamically for whichever user is currently accessing it. This affects obvious stuff like the "You are logged in as *username*" notice, less obvious stuff like the "edit" and "attachments" links that appear or disappear based on whether the user has permission to perform the action on the other end of the link, and even less obvious stuff like wiki-links to spaces the user can't see, or in-page macros that output their content based on the user's identity.

To run Confluence behind a caching reverse-proxy, you'd need one of:

1. A proxy that understood the user's identity, or
2. A Confluence site that removed all the personalised content for cacheable pages.

If you had (1), you could tell the proxy to cache content only for anonymous users (since all anon content is the same, and to survive a slashdotting you only really have to worry about the sudden influx of non-logged-in users). That said, (1) is quite tricky, as it relies on the existence of some SSO mechanism that both Confluence and Squid can be hooked into. If such a mechanism existed, though, it'd be a really neat solution.

In the absence of SSO, you've got (2), which involves.

- Theme Confluence so that the 'view page' 'view blog post' and 'view mail' pages contain no personalised content: no profile link or user identity, and all links to other functions available whether the user has permission to access them or not.
- Ensure that all wiki pages on the server are meant to be visible to anonymous users
- Disable (or avoid the use of) macros that deliver different content based on user identity
- Introduce an interceptor into Confluence that would provide If-Modified-Since/Last-Modified conditional get support for wiki pages
- Configure Confluence so the site root URL points to a page, rather than the dashboard.
- Configure Squid to cache the 'view page' URLs (/display/* /pages/viewpage.action /pages/viewblogpost.action)

This is assuming that only the site root or a regular wiki page would ever be the victim of a direct slashdotting, but I figure this is a reasonable enough assumption to make.

With conditional get supported, you could have Squid configured to query the server to see if a page has changed, and just put in some kind of sensible defaults for the maximum time to cache any page (5 minutes or so would be fine, since pages could contain dynamic content), and the minimum gap between if-modified queries (15 seconds would easily prevent the server from being overloaded, while making sure that in regular use you wouldn't get many situations where you edited a page, but couldn't see your own changes).

I am trying to install Confluence but the demo-site.zip file is missing

The `demo-site.zip` is normally located in the `WEB-INF/classes/com/atlassian/confluence/setup` directory.

There are some cases where the extraction utility used to extract the Confluence installation file will recursively extract the contents of all zip files contained within the installation file. If the `demo-site.zip` file has been extracted, you will see an `entities.xml` file in this directory instead.

RELATED TOPICS

Confluence Setup Guide
Confluence FAQ

How do I Disable Automatic Mail Polling?

Disabling mail polling in Confluence will prevent Confluence from checking for new mail automatically, but mail retrieval will still work if a user triggers "Check for new mail" manually.

To disable automatic mail polling you would need to restart Confluence with the following parameter in your JVM's `JAVA_OPTS`:

```
1.-Dconfluence.disable.mailpolling=true
```

RELATED TOPICS

How do I check which spaces have email accounts
Recognised System Properties

Disabling Profile Pictures on the Recently Updated Dashboard

To prevent Profile Pictures from being displayed in the Recently Updated Dashboard:

1. Open your `confluence/decorators/global.vmd` file.
2. Edit the following line (change true to false)

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

There is no need to stop or restart confluence. The change should work on the next refresh of the page.

How to Disable Profile Pictures from the Recently Updated Section of the Dashboard

Description

From confluence 2.6, recently updated section of the dashboard displays the profile pictures of the authors next to recently created pages and comments.

Resolution

To disable this, please follow the instructions outline in disabling profile pictures.

Additional Information

Severity	Medium
Regular Expression	[Dd]isabl(e|ing) [Pp]rofile picture
Article ID	CONFKB132382971

Searching Confluence Knowledge Base

Remove Version from Footer

1) Disable the footer content to be displayed
Remove/comment the following line in the main.vmd file which could also be accessed via Administration > Layouts > Main Layout (if you are using the "Default theme"):

```
#parse ("/decorators/includes/footer.vmd")
```

2) Modify the message rendered in the footer
This can be done in the ConfluenceActionSupport.properties file. See Modifying Confluence Interface Text for a step-by-step guide. The following line, is a mapping key to render the footer's message:

```
powered.by.atlassian.confluence
```

Searching Confluence Knowledge Base

Configuring a MySQL Datasource in Apache Tomcat

Instructions on setting up a MySQL DataSource connection for Confluence Standalone or EAR/WAR.



autoReconnect=true is required
The Confluence database connection URL must have `autoReconnect=true` added to the end to prevent disconnection issues.

Shut down Tomcat

- Run `bin/shutdown.sh` or `bin/shutdown.bat` to bring Tomcat down while you are making these changes

Install the Drivers

- download the MySQL JDBC drivers from http://www.mysql.com/downloads/api-jdbc-stable.html.
- After unpacking the file you have downloaded, you'll find a file called something like `mysql-connector-java-3.0.10-stable-bin.jar`.
- Copy this file into the `common/lib` directory of your Tomcat installation. Be aware that this directory may be just `lib` for Tomcat version 6 and beyond.

Configure Tomcat

The configuration is different for Conflence 2.2 onwards due to an upgrade to Tomcat 5.5

For users of Confluence 2.2 and later

- If you are using the Standalone distribution. edit the `conf/server.xml` file in your Tomcat installation. Users running their own Tomcat instance must edit the xml file where they declared the Confluence Context descriptor.
- If editing `conf/server.xml`, find the following lines:

```
1.<Context path="" docBase="../confluence" debug="0" reloadable="true">
2.                                                  <!-- Logger is deprecated in Tomcat 5.5.
Logging configuration for Confluence is specified in
confluence/WEB-INF/classes/log4j.properties -->
```

- Within the `Context` tags, directly after the opening <Context.../> line, insert the DataSource `Resource` tag:

```
1.<Resource name="jdbc/confluence" auth="Container" type="javax.sql.DataSource"
2.                      username="yourusername"
3.                      password="yourpassword"
4.                      driverClassName="com.mysql.jdbc.Driver"
5.                      url="jdbc:mysql://localhost:3306/confluence?autoReconnect=true"
6.                      maxActive="15"
7.                      maxIdle="7"
8.                      validationQuery="Select 1" />
```

** Replace the username and password parameters with the correct values for your database
  - In the url parameter, replace the word 'confluence' with the name of the database your confluence data will be stored in.
  - If you plan to use non-latin characters, you will also need to add "`&useUnicode=true&characterEncoding=utf8`" on the end of the above URL. These options are not required for any database other than MySQL.

The configuration properties for Tomcat's standard data source resource factory (org.apache.tomcat.dbcp.dbcp.BasicDataSourceFactory) are as follows:

- driverClassName - Fully qualified Java class name of the JDBC driver to be used.
- maxActive - The maximum number of active instances that can be allocated from this pool at the same time.
- maxIdle - The maximum number of connections that can sit idle in this pool at the same time.
- maxWait - The maximum number of milliseconds that the pool will wait (when there are no available connections) for a connection to be returned before throwing an exception.
- password - Database password to be passed to our JDBC driver.
- url - Connection URL to be passed to our JDBC driver. (For backwards compatibility, the property driverName is also recognized.)
- user - Database username to be passed to our JDBC driver.
- validationQuery - SQL query that can be used by the pool to validate connections before they are returned to the application. If specified, this query MUST be an SQL SELECT statement that returns at least one row.

For users of Confluence 2.1.x and earlier

- Edit the `conf/server.xml` file in your Tomcat installation
- Find the following lines:

```
1.<Context path="" docBase="../confluence" debug="0" reloadable="true">
2.                  <Logger className="org.apache.catalina.logger.FileLogger" prefix=
"atlassian-confluence."
3.                  suffix=".log" timestamp="true"/>
```

- Directly after the <Logger.../> line, (*before* the next </Context> line), insert the following:

```
01.<Resource name="jdbc/confluence" auth="Container" type="javax.sql.DataSource"/>
02.<ResourceParams name="jdbc/confluence">
03.          <parameter>
04.                      <name>factory</name>
05.                      <value>org.apache.commons.dbcp.BasicDataSourceFactory</value>
06.          </parameter>
07.          <parameter>
08.                      <name>maxActive</name>
09.                      <value>20</value>
10.          </parameter>
11.          <parameter>
12.                      <name>maxIdle</name>
13.                      <value>10</value>
14.          </parameter>
15.          <parameter>
16.                    <name>maxWait</name>
17.                    <value>10000</value>
18.          </parameter>
19.          <parameter>
20.                    <name>username</name>
21.                    <value>yourusername</value>
22.          </parameter>
23.          <parameter>
24.                    <name>password</name>
25.                    <value>yourpassword</value>
26.          </parameter>
27.          <parameter>
28.                    <name>driverClassName</name>
29.                    <value>com.mysql.jdbc.Driver</value>
30.          </parameter>
31.          <parameter>
32.                    <name>url</name>
33.                    <value>jdbc:mysql:
//localhost:3306/confluence?autoReconnect=true</value>
34.          </parameter>
35.</ResourceParams>
```

** You may want to choose different maxActive and maxIdle values: these are how many total database connections will be allowed at one time, and how many will be kept open even when there is no database activity

- Replace the username and password parameters with the correct values for your database
- In the url parameter, replace the word 'confluence' with the name of the database your confluence data will be stored in.
- If you plan to use non-latin characters, you will also need to add &useUnicode=true&characterEncoding=utf8 on the end of the above URL. These options are not required for any database other than MySQL.

Configure the Confluence web application

- Edit `confluence/WEB-INF/web.xml` in your confluence installation
- Go to the end of the file and just before `</web-app>`, insert the following:

```
1.<resource-ref>
2.                    <description>Connection Pool</description>
3.                    <res-ref-name>jdbc/confluence</res-ref-name>
4.                    <res-type>javax.sql.Datasource</res-type>
5.                    <res-auth>Container</res-auth>
6.          </resource-ref>
```

Configure Confluence

- **If you have not yet set up Confluence**
  - Follow the steps in the Confluence Setup Guide
  - In the Database Setup section, choose the "Datasource Connection" option.
  - Set the JNDI name to `java:comp/env/jdbc/confluence`
  - Set the Database dialect to MySQL.
- **If you are changing an existing Confluence installation over to using a Tomcat datasource**
  - Find your ConfluenceHome directory (see: Confluence Home Directory if you don't know where it is).
  - Edit the confluence.cfg.xml file
  - Delete any line that contains a property that begins with `hibernate.`
  - Insert the following at the start of the `<properties>` section.

```
1.<property name="hibernate.setup"><![CDATA[true]]></property>
2.          <property name="hibernate.dialect"
><![CDATA[net.sf.hibernate.dialect.MySQLDialect]]></property>
3.          <property name="hibernate.connection.datasource"
><![CDATA[java:comp/env/jdbc/confluence]]></property>
```

Restart Confluence

- Run `bin/startup.sh` or `bin/startup.bat` to start Tomcat with the new settings.

F.A.Q.

I see "Can't call commit when autocommit=true" during an upgrade of Confluence or otherwise.

Check out http://jira.atlassian.com/browse/CONF-1710 This error can be fixed by adding:
&relaxAutoCommit=true
to the end of your JDBC url.

Example:
jdbc:mysql://localhost/confdb?autoReconnect=true&relaxAutoCommit=true

This JDBC url can be found in your confluence.cfg.xml file located in your confluence home folder.

Installation FAQ

This section contains solutions for common issues encountered when installing and starting Confluence, including solutions to common queries about this process.

If necessary, review your logs by opening the Confluence install directory and checking the `/logs/catalina.out` and `/logs/catalina.out` files for errors you may encounter.

View one of the following issues or queries for more information:

- The Confluence window closes immediately when started
- How do I re-trigger the setup wizard
- Confluence starts but a problem prevents me from accessing the dashboard
- I receive a BUILD FAILED message when trying to create an EAR file in Confluence 2.6 or 2.7
- How much disk space does Confluence need?
- How Do I Make Confluence Accessible from the Root Context with a Tomcat EAR WAR configuration
- How To Run Confluence Standalone and Apache on Port 80 (Different IP Addresses)

The Confluence window closes immediately when started

An error is preventing Confluence from starting.

1. Open a command prompt. On Windows, do this by clicking on your Start menu, then click Run. In the Run box, type `cmd` and press OK
2. From the command prompt, go to your Confluence install directory
3. Go into the `bin` subdirectory
4. Run `startup.bat` and read the error message
5. Find a solution to the error below:
    - java.lang.NoClassDefFoundError IntraHibernateAttachmentCopier
    - Exception in thread "main" java.lang.NoClassDefFoundError: ...
    - Error creating Confluence Home directory
    - JAVA_HOME environment variable is not defined correctly
    - Port 8080 is in use
    - Error creating bean with name 'scheduler'
    - Error registering bean with name 'fileSystemAttachmentDataDao'

java.lang.NoClassDefFoundError IntraHibernateAttachmentCopier

If you are seeing "java.lang.NoClassDefFoundError: com/atlassian/confluence/pages/persistence/dao/hibernate/AbstractHibernateAttachmentDao$IntraHibernateAttachmentCopier", you have unzipped Confluence using a program that cannot handle long filenames. You must delete your install directory and go back to the point in the instructions that covers unzipping Confluence using a third-party unzip program.

Exception in thread "main" java.lang.NoClassDefFoundError: ...

If http://localhost:8080 goes to an error page when you run Confluence, go to the install directory and find the /logs/catalina.out. If this file contains a single line error starting with "Exception in thread "main" java.lang.NoClassDefFoundError", your install path may contain spaces. The solution is to rename your install directory path so that there are no spaces, then restart Confluence.

Error creating Confluence Home directory

The confluence.home variable specified in confluence-init.properties cannot be created. To fix this, edit /confluence/WEB-INF/classes/confluence-init.properties and check the path specified exists. Make sure all the directory slashes use / and that the # has been removed. If you're stuck, try using confluence.home=c:/confluence/data

JAVA_HOME environment variable is not defined correctly

You have not installed the Java Development Kit, or not set the %JAVA_HOME% to the directory of the JDK. You should re-check your steps in Stage 2.

Port 8080 is in use

If you have another Tomcat running on the same machine, you must edit <INSTALL>/conf/server.xml and change both 8080 and 8005 to ports that do not conflict with your other Tomcat instance according to these instructions.

Error creating bean with name 'scheduler'

You will need to adjust your system time.

Error registering bean with name 'fileSystemAttachmentDataDao'

If you did not use one of the recommended unzipping tools, and happened to use for instance the default Windows XP extractor, one of the classes required for starting up Confluence may not be located due to the lengthy `file-path`. This is due to the reason that the default Windows tool silently fails to extract files with long names. We suggest you to use other tools such as WinZIP and 7zip and redo the deployment process.

Cause:
org.springframework.beans.factory.BeanDefinitionStoreException: Error registering bean with name 'fileSystemAttachmentDataDao' defined in class path resource [applicationContext.xml]: Class that bean class [com.atlassian.confluence.pages.persistence.dao.FileSystemAttachmentDataDao] depends on not found; nested exception is java.lang.NoClassDefFoundError: com/atlassian/confluence/pages/persistence/dao/FileSystemAttachmentDataDao$FileSystemAttachmentNamingStrategy
at org.springframework.beans.factory.xml.DefaultXmlBeanDefinitionParser.parseBeanDefinition(DefaultXmlBeanDefinitionParser.java:366)
caused by: java.lang.NoClassDefFoundError: com/atlassian/confluence/pages/persistence/dao/FileSystemAttachmentDataDao$FileSystemAttachmentNamingStrategy
at java.lang.Class.forName0(Native Method)

How do I re-trigger the setup wizard

To re-trigger the setup wizard:

1. Ensure the application server (for example, Apache Tomcat) running Confluence has been stopped.
2. Delete `<confluence-home>/confluence.cfg.xml`.
3. Re-start Confluence's application server and then Confluence.

Confluence starts but a problem prevents me from accessing the dashboard

Find a solution to one of these problems below:

- localhost:8080 times out
- localhost:8080 goes to the Tomcat start homepage
- Logins fail at the login screen

localhost:8080 times out

Check the server logs for errors. If you are running Confluence Standalone on Windows, error messages will be printed to the console window that opened when you ran `startup.bat`. On Unix systems, Confluence will log messages to `logs/catalina.out`.

localhost:8080 goes to the Tomcat start homepage

The CATALINA_HOME environment variable is set to another instance of Tomcat. You should run shutdown.sh, remove the CATALINA_HOME reference to the other Tomcat version, and run startup.sh again.

Logins fail at the login screen

If you try to login with the correct username and password but are always returned to the login screen without any error messages, and you are running Zone Alarm, please check that it is not blocking the Confluence server.

I receive a BUILD FAILED message when trying to create an EAR file in Confluence 2.6 or 2.7

When trying to create a EAR file, you may encounter with the following error:

```
BUILD FAILED
file:/users/a237185/confluence/confluence-2.7/build.xml:67: Warning: Could not find file
/users/a237185/confluence/confluence-2.7/etc/ear-application.xml to copy.
```

This is due to the `build.xml` file being incorrect for **Confluence 2.6.x and 2.7.0**. This problem will be fixed for future releases. However for these releases, please open the `build.xml` file and change the following two lines:
From:

```
1.<copy preservelastmodified="true" file="${ant.confluence.etc}/ear-application.xml" tofile=
"${ant.confluence.build.ear}/META-INF/application.xml" overwrite="yes"/>
2.<copy preservelastmodified="true" file=
"${ant.confluence.dist}/${ant.confluence.name}-${ant.confluence.version}.war" tofile=
"ant.confluence.nameild.ear}/Confluence War Distribution.war" overwrite="yes"/>
```

To:

```
1.<copy preservelastmodified="true" file="ear-application.xml" tofile=
"${ant.confluence.build.ear}/META-INF/application.xml" overwrite="yes"/>
2.<copy preservelastmodified="true" file=
"${ant.confluence.dist}/${ant.confluence.name}-${ant.confluence.version}.war" tofile=
"${ant.confluence.build.ear}/Confluence War Distribution.war" overwrite="yes"/>
```

How much disk space does Confluence need?

For the recommended minimum amount of disk space, please see the 'Requirements' section in the Installation Guide.

Note that the actual disk space needed will depend significantly on the number of attachments (i.e. files which users attach to Confluence pages), and on the sizes of the attachments. You can calculate the hard drive requirements as you would with any standard file server.

RELATED TOPICS

Important Directories and Files (Confluence Docs 3.2)

FAQ Home

How Do I Make Confluence Accessible from the Root Context with a Tomcat EAR WAR configuration

Tomcat uses the ROOT.xml file to describe the root context. To make Confluence run at the Root, name the file that contains the context descriptor for Confluence ROOT.xml in the conf/Catalina/localhost/ directory. This is described in Installing the Confluence EAR-WAR Edition .

How To Run Confluence Standalone and Apache on Port 80 (Different IP Addresses)

To Configure Confluence to run on port 80 alongside another web service already on Port 80, you can add an "address" attribute to the connector element where address="<IP address on which Tomcat-Standalone is running >"

Change the `server.xml` file by adding the element:

> *<Connector className="org.apache.coyote.tomcat4.CoyoteConnector"*
> *address="192.168.1.1"*
> *port="8080" minProcessors="5" maxProcessors="75"*
> *enableLookups="true" redirectPort="8443"*
> *acceptCount="100" debug="0" connectionTimeout="20000"*
> *useURIValidationHack="false" disableUploadTimeout="true" />*

Also change the hostname from localhost to the relevant host name, in the file above, on the same IP.

JIRA Integration FAQ

This section contains solutions for common issues or queries associated with Confluence's JIRA integration features, such as its JIRA Issues Macro or JIRA User Management integration.

View one of the following issues or queries for more information:

- The JIRA Issues Macro behaves problematically
- The JIRA Issues Macro generates an error
- When setting up JIRA user management in Confluence, the Confluence login page loads but login fails
- When setting up JIRA user management in Confluence, the Confluence login page loads with an error

The JIRA Issues Macro behaves problematically

When the JIRA Issues Macro is used, one of the following problems may occur:

- Cannot access issues when JIRA uses HTTPS
- Issue icons are not displayed
    - RELATED TOPICS

Cannot access issues when JIRA uses HTTPS

Take a look at the Confluence Knowledge Base article on problems connecting to SSL services.

Issue icons are not displayed

Check that you are using the correct macro version for JIRA.

RELATED TOPICS

JIRA Issues Macro
Setting Up Trusted Communication between JIRA and Confluence
The JIRA Issues Macro generates an error

The JIRA Issues Macro generates an error

When the JIRA Issues Macro is used, one of the following errors is generated:

- The message 'Error Rendering Macro' is displayed, or either no or not all issues are displayed
- The message 'Error rendering macro: java.io.IOException: Could not download' is displayed
- The message 'Error rendering macro: java.io.IOException: Error on line -1: Premature end of file' is displayed
    - RELATED TOPICS

The message 'Error Rendering Macro' is displayed, or either no or not all issues are displayed

You may need to set up trusted communication between JIRA and Confluence. Please refer to the section on restricted viewing.

The message 'Error rendering macro: java.io.IOException: Could not download' is displayed

You may need to configure Confluence to acknowledge your proxy server before it is able to download the feed.

The message 'Error rendering macro: java.io.IOException: Error on line -1: Premature end of file' is displayed

Did you select an existing JIRA filter when you built the macro? If you did, your issues filter URL will end with `.../12345/SearchRequest-12345.xml` where 12345 is any number. You should retry using a new and unsaved filter.

RELATED TOPICS

JIRA Issues Macro
Setting Up Trusted Communication between JIRA and Confluence
The JIRA Issues Macro behaves problematically

When setting up JIRA user management in Confluence, the Confluence login page loads but login fails

When setting up JIRA user management in Confluence, the Confluence login page loads but login fails with one of the following errors:

- The Confluence login page loads but login fails with a 'Username and password are incorrect' error and the output log shows 'Access denied for user'
- The Confluence login page loads but login fails with a 'Username and password are incorrect' error and the output log shows 'Cannot create JDBC driver'
- I cannot get my JIRA integration to work, where can I get technical support?

The Confluence login page loads but login fails with a 'Username and password are incorrect' error and the output log shows 'Access denied for user'

This is caused by an incorrect database URL or login in the datasource.

The Confluence login page loads but login fails with a 'Username and password are incorrect' error and the output log shows 'Cannot create JDBC driver'

If your log outputs an error with "FATAL [user.provider.jdbc.JDBCCredentialsProvider] Could not list users. org.apache.commons.dbcp.SQLNestedException: Cannot create JDBC driver of class '' for connect URL 'null'", you are using the incorrect Tomcat format for specifying the Datasource resource. You should check your Tomcat version and use the alternative format.

I cannot get my JIRA integration to work, where can I get technical support?

See Requesting External User Management Support for information on logging a support request.

When setting up JIRA user management in Confluence, the Confluence login page loads with an error

When setting up JIRA user management in Confluence, the Confluence login page loads with one of the following errors:

- An error is encountered when JIRA is using Microsoft SQL Server
- The Confluence login page loads with a 'NullPointerException' system error
- The Confluence login page loads with 'HTTP Status 404' and the output log shows a 'java.lang.ClassNotFoundException' error for the driver, eg 'com.mysql.jdbc.Driver'

An error is encountered when JIRA is using Microsoft SQL Server

There is a known issue when JIRA is using Microsoft SQL Server and the schema name is 'jira'. Unfortunately in Confluence, you cannot define this schema name. Please refer to Knowledge Base article CSP-19533.

The Confluence login page loads with a 'NullPointerException' system error

If the page loads with 'System error' with cause 'java.lang.NullPointerException at com.atlassian.user.impl.osuser.OSUUserManager.getOpensymphonyUser(OSUUserManager.java:85)' and output logs shows 'user.provider.jdbc.BaseJDBCProvider] init Could not look up DataSource using JNDI location' error, either the Resource for the DataSource is not being loaded by the application server, or it is being loaded but the resource names do not match. Check the names first, so if '/confluence/WEB-INF/classes/osuser.xml' specifies a datasource named 'java:comp/env/jdbc/JiraDS', the datasource specified in server.xml or confluence.xml must be 'jdbc/JiraDS'.

The Confluence login page loads with 'HTTP Status 404' and the output log shows a 'java.lang.ClassNotFoundException' error for the driver, eg 'com.mysql.jdbc.Driver'

The database driver library is missing from your Confluence/application server installation. In your Confluence Standalone or Apache Tomcat installation, download the database jar to the common/lib directory.

LDAP FAQ

This section contains solutions for common issues or queries associated with LDAP integration in Confluence.

View one of the following issues or queries for more information:

- Are all users in LDAP visible in Confluence administration and can they can be assigned permissions or to groups?
- Can we use LDAP and Confluence groups simultaneously, as a 'mixed mode', where some groups are kept in Confluence and others in LDAP?
- Confluence Domino LDAP FAQs
- Confluence integration with LDAP and Active Directory FAQs
- How are LDAP or Active Directory users counted toward my license limit?
- How can I assign an LDAP user a Confluence account?
- How can I enable LDAP?
- How does Confluence handle user deletions from LDAP? Is the user's assignment to one or more groups still visible?
- I am having a problem with Confluence LDAP integration
- I enabled LDAP and some users are now returned twice under the user browser
- If a Confluence user had a lowercase username, but an LDAP user has the same username in UPPERCASE, does it matter which one I use?
- If a user already exists in Confluence and an LDAP user with the same username is added, which account's password gets used?

Other Questions

For troubleshooting, please create a problem report. General enquiries should be posted to a support ticket

RELATED TOPICS

LDAP User Management

Are all users in LDAP visible in Confluence administration and can they can be assigned permissions or to groups?

All LDAP users with 'Can Use' permission can be viewed from the user browser, even if they have never logged in. When an LDAP user logs in for the first time, a Confluence user account is created automatically to store their information. You have read-only access to LDAP groups, and can add/remove Confluence internal groups to any user.

Can we use LDAP and Confluence groups simultaneously, as a 'mixed mode', where some groups are kept in Confluence and others in LDAP?

Yes.

Confluence Domino LDAP FAQs

Find an answer to one of these questions below:

- Why are individual users always shown as not belonging to any groups?

Why are individual users always shown as not belonging to any groups?

Domino servers allow user groups to be set as 'mail-only', 'access control' and 'multi-purpose'. If the groups are set to 'mail-only', when Confluence queries the Domino LDAP server about a given user, Domino will return null. Groups that are created as 'multi-purpose' seem to work fine.

Confluence integration with LDAP and Active Directory FAQs

Find an answer to one of these questions below:

- Can Confluence make distinctions between security and distribution groups, or group types?
- Can group memberships be retrieved from multi-domain forests?
- Can Confluence support multiple Active Directory repositories?
- Can Confluence handle nested groups?

Can Confluence make distinctions between security and distribution groups, or group types?

No, Confluence has no group types. However, you can configure Confluence to only recognise some of these groups over others. For example, you can configure Confluence to only recognise distribution groups. this is done by adjusting the `groupSearchFilter` in your atlassian-user.xml file.

Can group memberships be retrieved from multi-domain forests?

Yes, you can do this by configuring multiple repositories: one for each domain. More instructions on how to do this can be found here.

Can Confluence support multiple Active Directory repositories?

Yes.

Can Confluence handle nested groups?

No, each child group must be individually specified instead. You may wish to vote towards support for nested groups at USER-101.

How are LDAP or Active Directory users counted toward my license limit?

Your user count is determined by the number of internal users plus the number of LDAP users who can potentially login. LDAP users that are a member of an LDAP group with 'Can Use' permission granted in Confluence can all potentially login, which means that all members of groups with this permission granted will be counted towards your license. To manage your license usage, only grant login permission to Active Directory groups where all members need accounts. You may like to setup a special confluence LDAP group if no combination of your existing groups is suitable.

How can I assign an LDAP user a Confluence account?

LDAP groups or users granted 'Can Use' permission under 'Global Permissions' can login to Confluence.

How can I enable LDAP?

See Add LDAP Integration.

How does Confluence handle user deletions from LDAP? Is the user's assignment to one or more groups still visible?

Users are not deleted from Confluence, but their logins are disabled within one hour as they expire in the cache. Only non-LDAP groups are retained. Refer to the overview for more detail.

I am having a problem with Confluence LDAP integration

Find a solution to one of these problems below:

- I just added LDAP integration, why can't I login using my original account?
- Why do my LDAP users sees 'Not Permitted' screens when they login?
- Confluence fails to start with error 'Error creating bean with name 'userManager' defined in class path resource [atlassianUserContext.xml]'?
- Editing a user under Administration -> Manage Users throws an error 'org.apache.velocity.exception.MethodInvocationException'
- After setting up LDAP, I cannot see LDAP users or groups from the Confluence user or group browser
- I cannot see an LDAP/AD group in Confluence

- I cannot get my LDAP to work, where can I get technical support?

I just added LDAP integration, why can't I login using my original account?

If there is an LDAP user with the same username as your administrator account, you must now use their password to login. LDAP logins override internal logins.

Why do my LDAP users sees 'Not Permitted' screens when they login?

To login, the user must be a member of one or more groups that have been granted 'Can Use' permission from the Administration -> Global Permissions -> Group Permissions.

Confluence fails to start with error 'Error creating bean with name 'userManager' defined in class path resource [atlassianUserContext.xml]'?

Your atlassian.xml file may contain filters with characters that must be escaped from XML. Check here for details.

Editing a user under Administration -> Manage Users throws an error 'org.apache.velocity.exception.MethodInvocationException'

If you see an error:
"org.apache.velocity.exception.MethodInvocationException: Invocation of method 'isUserDeactivated' in class com.atlassian.confluence.user.actions.ViewUserAction threw exception class java.lang.NullPointerException : null"
You should open \confluence\WEB-INF\classes\atlassian-user.xml and check that your Hibernate Repository is not wrapped in a comment tag (<!-- and -->). The line to uncomment is:
<hibernate name="Hibernate Repository" key="hibernateRepository" description="Hibernate Repository" />

After setting up LDAP, I cannot see LDAP users or groups from the Confluence user or group browser

Are your users or groups located in subtrees beneath the directory returned by the search filter? If so, you may need to add `<usersearchalldepths>TRUE</usersearchalldepths>` or `<groupsearchalldepths>TRUE</groupsearchalldepths>` to your `atlassian-user.xml` See Map LDAP Users and Groups for details.

I cannot see an LDAP/AD group in Confluence

Is the group in a subtree? If so, you will need to edit atlassian-user.xml and add a `groupSearchAllDepths=true` parameter to the LDAP repository to set Confluence to search subtrees of the base group namespace. See Map LDAP Users and Groups for details.

I cannot get my LDAP to work, where can I get technical support?

See Requesting External User Management Support for information on logging a support request.

I enabled LDAP and some users are now returned twice under the user browser

When LDAP is enabled, it is normal for the user browser to return two copies of users who have both LDAP and internal users accounts. If you are interested in a fix for this, please vote towards 'User browser shows duplicate accounts when a user exists both locally and in LDAP'.

If a Confluence user had a lowercase username, but an LDAP user has the same username in UPPERCASE, does it matter which one I use?

As mentioned in another FAQ, LDAP login has priority over the confluence login. However, only the password is taken into account here. You can log in with either the lowercase or UPPERCASE username.

If a user already exists in Confluence and an LDAP user with the same username is added, which account's password gets used?

The LDAP login has priority over the Confluence login. If LDAP 'Can Use' permission is removed or the user is deleted, the Confluence login will still work.

Mail Archiving FAQ

This section contains solutions for common issues or queries associated with Confluence's Mail Archiving feature that was introduced in Confluence 1.3.

View one of the following issues or queries for more information:

- Can Confluence replace my regular mail client?
- How do I get mail into Confluence?
- How do I use the mail archive?

- Okay, I've imported the mail, but where is it?

The emphasis on the DR3 release was to:

- import email into Confluence
- have Confluence monitor POP mailboxes
- view email
- search email

However, any further suggestions for this feature are welcome. Please let us know about them by filing an issue in JIRA, commenting on the forum, or just dropping us an email.

Can Confluence replace my regular mail client?

No.

Confluence's mail archive is designed to supplement the way you currently handle email, not to replace it. This is why Confluence deliberately does not come with features common in email clients. For example, you can not mark emails as read or unread, you can not reply to emails from within Confluence, and so on.

How do I get mail into Confluence?

All mail messages belong to a particular space.

From the space administration screen, space administrators can:

- Configure Confluence to poll a POP mailbox for incoming mail
    1. Go to Space Admin
    2. Choose "Mail Accounts"
    3. Add Pop Account
- Import mail from an mbox-format mail file

Confluence will delete mail from a POP box as it reads it. Do not point Confluence to an account unless you are happy with it removing all the mail you have stored there.

How do I use the mail archive?

Some suggested scenarios include:

Project-related conversations

Say you are using a Confluence space to organise a project. The project lead and the customer have a long conversation (via email) clarifying the project's goals and requirements. Rather than have that conversation lost in their individual mailboxes, if they *CC'ed* their mail to a POP box being monitored by Confluence, all that information will be archived alongside the rest of the project's documentation.

Customer Support Tracking

All incoming and outgoing sales and support email is diverted (at the mail-server) to be read by Confluence. Staff can then use Confluence's features to find all previous communications with particular customers.

And?

Of course, the real fun will come from seeing how this feature can be put to other uses.

Okay, I've imported the mail, but where is it?

Because of the typically overwhelming volume of email, especially compared to the more sedate pace of wiki updates, we do not notify you of recently arrived mail in the same places we notify you of changed Confluence content.

Mail will not appear in the recent changes list on the dashboard or space summary pages. Similarly, mail will not appear by default in search results.

You can view mail...

1. In order of arrival from the *Mail Archive* section, found under the *Content* tab of the Space Summary screen.
    - 'Browse' Dropdown menu > Mail
2. You can also explicitly select Mail (or All Content) in the search page to include mail in your search results.

New User FAQ

This section contains solutions for common issues or queries encountered by new Confluence users.

View one of the following issues or queries for more information:

- Can I use CamelCaseLinks like they do on WardsWiki?
- Can Users Edit Individual Sections Within a Page?
- How does Confluence differ from a wiki?

Can I use CamelCaseLinks like they do on WardsWiki?

Yes you can. Camelcasing is not enabled by default but a site administrator can easily enable it from the administration screens.

See Enabling CamelCase Linking.

RELATED TOPICS

CamelCase linking

Showing first 1 of 2 results

FAQ Home

Can Users Edit Individual Sections Within a Page?

Some wiki software allows the editing of sections within a page (sectional editing). This functionality is currently not available in Confluence, but we are looking to include it in a future release. This issue is being tracked on the Confluence JIRA project: CONF-5913.

In the meantime, for pages that are getting long enough to be hard to edit in a single block, you can get an approximation of sectional editing by using the {include} macro. For example:

```
h3. [Section One] ([edit|///pages/editpage.action?spaceKey=SPACE&pageTitle=Section One])

{include:Section One}

h3. [Section Two] ([edit|///pages/editpage.action?spaceKey=SPACE&pageTitle=Section Two])

{include:Section Two}
```

The links to the edit pages can be simplified by using the {link-to:page edit} macro available in David Peterson's Linking Plugin.

Technical Stuff

The problem lies in the complexity of Confluence's wiki markup. We made a couple of proof-of-concept implementations of sectional editing as part of our ShipIt Day program, and while it's quite easy to come up with a solution that works with 90% of pages, there are a lot of edge-cases where it's actually quite hard to determine precisely where a given section starts and finishes.

Next time we perform a significant overhaul of our wiki markup processing engine, we'll be looking specifically to add functionality that will make sectional editing work properly.

How does Confluence differ from a wiki?

Essentially, Confluence is a wiki. Our aim was to build an application that was built to the requirements of an enterprise knowledge management system, without losing the essential, powerful simplicity of the wiki in the process.

From the wiki, we took the following lessons:

- It should be easy for anyone to create and edit pages
- It should be easy for anyone to link pages together
- It should be easy to see what has changed recently
- The site should be searchable
- Users should have the tools to organise and group pages without having any particular structure imposed upon them

On top of that, we added professional features, such as the partitioning of content into separately managed spaces, user- and group-based access control, automated refactoring, PDF exporting, searchable attachments, a comprehensive remote API, easy installation and a professional and easy-to-use presentation; all wrapped up in Atlassian's "Legendary Service".

RELATED TOPICS

More about Confluence on Atlassian's website
More about wikis

FAQ Home

RSS Feeds FAQ

This section contains solutions for common issues or queries associated with RSS Feeds and the RSS Feed Macro.

View one of the following issues or queries for more information:

- Create an RSS feed for mail from only specified mail accounts
- How do I fix a "Could not download (Feed URL) - Connection timed out (errno238)" error?
- How do I fix a "Could not retrieve (Feed URL) - Not Permitted" error?
- How do I fix an "Error formatting 'macro rss java.lang.NullPointerException'" error?
- How do I fix an "Unable to retrieve (Feed URL) - Connection refused - connect" error?
- How do I force authentication for public feeds?
- Is it possible to delete a feed?
- I want to remove RSS Feeds completely

Create an RSS feed for mail from only specified mail accounts

This is not possible, but you can vote towards tagging incoming mail with labels on arrival. Once tagged, feeds could monitor all new mail with that label.

How do I fix a "Could not download (Feed URL) - Connection timed out (errno238)" error?

The feed source may be offline, or the firewall may be blocking access either between the Confluence server any your computer. Confirm that you can access the feed URL from your browser. If it cannot, your firewall settings may be blocking access to Confluence. For example, your server may be configured to block outgoing requests.

How do I fix a "Could not retrieve (Feed URL) - Not Permitted" error?

You must append a valid login to Private Feeds as described in the Usage section of the RSS Feed Macro.

How do I fix an "Error formatting 'macro rss java.lang.NullPointerException'" error?

The link is not a valid feed, so check your URL. If stuck, you can recreate internal Confluence feeds Using the RSS Feed Builder.

How do I fix an "Unable to retrieve (Feed URL) - Connection refused - connect" error?

The URL is invalid. If the link appears correct, confirm that you can access Confluence. Paste the feed into a third-party RSS feed reader and confirm that it can access it. If it cannot, your firewall settings may be blocking access to Confluence. For example, your server may be configured to block outgoing requests.

How do I force authentication for public feeds?

With anonymous access enabled, you can force user authentication when creating the feed by checking 'Authorised'. If anonymous access is disabled, all feeds will require user authentication.

Is it possible to delete a feed?

No, because RSS feeds are based on the view permissions for pages and spaces. RSS is an extension of normal page viewing functionality, so if you can view a page, you can receive an RSS feed for it. The only way to remove an RSS feed is to prevent all access to a page for that user, so that no content will be delivered.

I want to remove RSS Feeds completely

While Confluence does not have this functionality, there is a work around to remove RSS feeds completely. Refer to How do I Disable RSS Feeds?.

Upgrade FAQ

This section contains solutions for common issues and queries encountered when upgrading Confluence.

View one of the following issues or queries for more information:

- I cannot find the "Rich Text" editor. Is the editor part of Confluence 1.4.3?
- Server ID FAQ
- Upgrade My Trial To A Commercial Version

I cannot find the "Rich Text" editor. Is the editor part of Confluence 1.4.3?

The Rich Text Editor (aka WYSIWYG editor) is available in Confluence 2.0 and upwards. Rich Text editing is enabled by default.

If you wish to upgrade your Confluence installation, instructions can be found here.

RELATED TOPICS

Enabling Rich Text Editing Option (Confluence Docs 3.2)

Making Rich Text Editing default (Confluence Docs 3.2)

Showing first 2 of 3 results

FAQ Home

Server ID FAQ

What causes this Server ID to be generated? Is it tied to the hardware, OS, or Confluence instance?

The Server ID:

- is generated when you install Confluence for the first time
- exists for the life of the Confluence instance
- survives an upgrade
- is held in the database
- is not bound to a specific licence
- is the same for all servers in a cluster.

What's the policy on re-associating licenses with server IDs?

There's no need to do this. Once you have a Server ID associated to a license, you can leave it as is.

What happens when I need to reinstall (quickly) on a different system?

Because the server ID is held in the database, it travels with the instance when the database or XML backup is restored on the new system. You need not generate a new Server ID for your new system.

What do I do when the license screen from my.atlassian.com is asking for my Server ID?

The Server ID is located on the license screen. If you have only a License ID, you may bypass the requirement to enter a Server ID - just look for the link from my.atlassian.com on my.atlassian.com after choosing "associate server ID." Check Unable to Find Server ID for Confluence 2.5.4 or Before for further info.

Additional Information

Severity	Low
Article ID	CONFKB151519275

Searching Confluence Knowledge Base

Upgrade My Trial To A Commercial Version

First, you will need to purchase Confluence to receive your commercial license key. If you already have a free 30 day evaluation then you can easily convert this to the commercial version, or setup your commercial instance on another system and transfer your trial data across.

Upgrade A Trial To The Commercial Version

If you wish to change your trial into a commercial version while keeping the same hardware, login in as a Confluence administrator and paste in your commercial license key under the Administrator > License Details screen. The conversion to a commercial version is instantaneous and can be done regardless of whether your evaluation period has expired or is still ongoing.

Migrate Your Trial Data To A Server

If you installed your Confluence trial on a PC but wish to host your commercial version on another system, such as a server managed by your IT department, you can transfer the trial data across. Begin by creating an XML backup on the trial PC, then install Confluence on the

commercial server. During installation, you will be presented with the Confluence Setup Wizard. The wizard gives you the opportunity to insert your commercial license key and also to import an XML backup. If you import the XML backup from your trial, your commercial instance will be setup already containing all your trial wiki content and any internal users.

Usage FAQ

This section contains solutions for common issues or queries about the everyday use of Confluence.

View one of the following issues or queries for more information:

- Add many files to a page at once
- Create a page by passing parameters to a template
- How do I obtain content that hasn't been modified in a certain period of time
- How to Add a Quick Search for Firefox
- How to disable PDF Export
- How to Find Pages with no Label
- How to Make Confluence Open a New Tab when Clicking on the Attachments Link
- How to Reset a Custom Layout
- Redirect to a specific page (home page) within the site after logging in
- Setup email notifications of page updates

Add many files to a page at once

Confluence offers several ways to add many files to a page at once:

- Using of Confluence's **Drag-and-Drop** functionality, simply drag the files onto any Confluence page. Refer to Drag-and-Drop for more information.
- Confluence pages and attachments can also be mounted as a network drive, and files can be dragged and dropped into Confluence using the WebDAV Plugin.
- Alternatively, users who do not wish to use WebDAV can write a custom script to attach all PDF files in a directory to a Confluence page using the addAttachment function in the Remote API. This script can be adapted from one of the Remote API script examples.

Create a page by passing parameters to a template

Confluence supports populating wiki content through templates. Check out Page Templates if you would like to create a new page by filling in a graphical, form-based template. If you have an existing page and would like to pass text as parameters to a macro that fills in the blanks in a template, create the template as a User Macro and call it from inside your Confluence page.

How do I obtain content that hasn't been modified in a certain period of time

Archive plugin

The Archiving Plugin is a great tool for managing outdated content.

Via SQL

This can be achieved by running the following SQL query on your Confluence database.

```
1.select * from content as c, spaces as s where c.spaceid = s.spaceid and s.spacename='INSERT SPACE
NAME HERE' and c.LASTMODDATE < 'INSERT DATE HERE';
```

How to Add a Quick Search for Firefox

Description

To add a quick search term into Firefox's address bar, add the following link to your bookmarks:

```
http://confluence.atlassian.com/dosearchsite.action?quickSearch=true&searchQuery.queryString=%s
```

Make a keyword for it, and you can search a confluence instance using the keyword, a space and then the search term. For example, if you use the keyword cac with the above link, you can search confluence using:

```
cac searchterm
```

How to disable PDF Export

Modifying PDF Links Behaviour

If you want to remove the links regarding Export to PDF from your instance, you can make changes to the underlying velocity templates to remove the options from the list.

From: <confluence-install>/confluence/pages/viewinfo.vm
Remove this code:

```
1.<td nowrap>
2.                <a href=
"$req.contextPath/pages/doexportpage.action?pageId=$confPage.id&type=TYPE_PDF" rel=
"nofollow">$action.getText('pdf')</a>
3.                <a href="$req.contextPath/exportword?pageId=$confPage.id" rel="nofollow"
>$action.getText('msword')</a>
4.                </td>
```

Disabling export to PDF

From <confluence-install>/confluence/template/includes/macros.vm:

```
1.#if ($sitemeshPage && $req && $pageId && $permissionCheckDispatcher.isPermitted(
"/pages/doexportpage.action?pageId=$pageId&type=TYPE_PDF"))
2.                                  <a href=
"$req.contextPath/pages/doexportpage.action?pageId=$pageId&type=TYPE_PDF" rel=
"nofollow">
3.                                  <img src="$req.contextPath/images/icons/attachments/pdf.gif"
height="16" width="16" border="0" align="absmiddle" title="$action.getText('export.as.pdf')"></a>
4.                          #end
```

How to Find Pages with no Label

You can use the Label Management Plugin to find out pages that contain no labels.

Severity	Low
Regular Expression	([Uu]nlabeled pages|pages with no label|pages without label)
Article ID	CONFKB154239468

Searching Confluence Knowledge Base

How to Make Confluence Open a New Tab when Clicking on the Attachments Link

Modifying the template file in confluence-attachments-plugin-x.x.jar (in this case, I try on confluence-attachments-plugin-2.10.jar). Please edit attachmentsmacro.vm file. This file is located in confluence-attachments-plugin-2.10.jar file. Extract this file by using the [Editing files within .jar archives guide] guide.

If you are able to extract it successfully, you can locate to attachmentsmacro.vm file and find the following lines:

```
1.<td><a name=
"$generalUtil.urlEncode($page.title)-attachment-$generalUtil.urlEncode($attachment.fileName)"
>#parse ("/pages/includes/attachment_icon.vm")</a> <a href=
"$req.contextPath$attachment.downloadPathWithoutVersion">$attachment.fileName</a></td>
```

and change the above code to be:

```
1.<td><a name=
"$generalUtil.urlEncode($page.title)-attachment-$generalUtil.urlEncode($attachment.fileName)"
>#parse ("/pages/includes/attachment_icon.vm")</a> <a target="_blank" href=
"$req.contextPath$attachment.downloadPathWithoutVersion">$attachment.fileName</a></td>
```

After making changes, please repack the file by following the steps [here].

Additional Information

Severity	Medium
Regular Expression	attachments in a new tab

Article ID	CONFKB160792804

Searching Confluence Knowledge Base

How to Reset a Custom Layout

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.

Redirect to a specific page (home page) within the site after logging in

As a user, you can set the home page to which you are sent after logging in.

To set your home page as a user:

1. View your profile via the profile link on the top right
2. Click the **Edit My Profile** on the right
3. Set **Site Homepage** to your desired home page

For instructions on configuring this feature at the administrator level, see Redirect users to a site-wide home page after a successful login.

RELATED TOPICS

Redirect users to a site-wide home page after a successful login

Setup email notifications of page updates

Confluence sends notifications of changes to spaces or pages to users who are watching pages. Users choose to watch pages or spaces, an administrator does not force notifications upon them.



There is one way to automatically have users watch a given page or news item. The Autowatch Plugin can be set up to have anyone who contributes to a space watch the content they have contributed to. Eg. if they comment on a page, they will be set to watch that page for further updates. This should be used with care.

To set up a watch on a specific page or space as a user, see:

- Watching a Page or
- Watching a Space

For instructions on configuring administrator-level email notifications, see Configuring Confluence to send email notifications.

RELATED TOPICS

Configuring Confluence to send email notifications

Unsupported Platform Information



The information in this section relates to platforms which are not supported for Confluence. Consequently, Atlassian **can not guarantee providing any support for it**. Please be aware that this material is provided here for your information only and using it is done so at your own risk.

For details on supported platforms for Confluence, please refer to the Supported Platforms topic.

View one of the following topics for more details:

- Setting up Confluence with IIS
  - Connecting Confluence with IIS 5.1 or 6
  - Connecting Confluence with IIS 7
- Using the IBM 64bit J9 JDK

Setting up Confluence with IIS



The content on this page relates to platforms which are not supported for Confluence. Consequently, Atlassian **can not 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.

This page describes how to install Confluence Standalone with IIS using the Apache jk connector.



If you are using JIRA as well as Confluence, please see this document instead.

On this page:

- Step 1. Install IIS
  - IIS 6
  - IIS 7
- Step 2. Install Confluence Standalone
- Step 3. Configure Tomcat
- Step 4. Configure the Tomcat Connector
- Step 5. Connect Confluence with IIS

Step 1. Install IIS

If you are running Windows Server 2003, you will only be able to use IIS 5.1 or 6. If you are using Windows Server 2008, you might like to install IIS 7.

IIS 6

If you are using Windows Server 2003 or XP professional, follow these instructions for installing IIS 6.

After installation is complete you should be able to go to http://localhost/iishelp/iis/misc/default.asp in your browser and see the IIS Getting Started page.

IIS 7

Similar to the previous Windows versions, IIS is not installed by default in Windows Server 2008 so you need to install it manually.

1. Start your `Server Manager`.
2. Click 'Roles'.
3. In the right hand panel, click "Add Roles".
4. A new window will pop up. Select the 'Web Server (IIS)' option.
5. Click 'Next' until you see another set of checkbox options to install the required 'Roles Services' for the web server (IIS).
6. Scroll down to 'Application Development' and tick the following:
   - CGI
   - ISAPI Extensions
   - ISAPI Filters
7. Click 'Next'.
8. And lastly, click 'Install'.

To check that IIS has been installed successfully, you can direct your browser to http://localhost/ and see the IIS 7 logo.

You can learn more about IIS 7 from this website.

Step 2. Install Confluence Standalone

Do a normal Confluence installation, after which you should be able to use confluence as usual through the URL http://localhost:8080.

Step 3. Configure Tomcat

Add another connector to your `server.xml` file, directly after the existing `<Connector ... />` tag:

```
1.<Connector port="8009" enableLookups="false" redirectPort="8443"
2.protocol="AJP/1.3" URIEncoding="UTF-8" />
```

Restart Confluence.
In the logs/catalina.YYYY-MM-DD.log file you should see the **Jk** is running:

```
1.INFO: Starting Coyote HTTP/1.1 on http-8080
2.7/09/2006 14:40:04 org.apache.jk.common.ChannelSocket init
3.INFO: JK: ajp13 listening on /0.0.0.0:8009
4.7/09/2006 14:40:04 org.apache.jk.server.JkMain start
5.INFO: Jk running ID=0 time=0/31&amp;nbsp; config=null
```

Step 4. Configure the Tomcat Connector

These instructions are based on the Tomcat Connector, IIS Configuration documentation.

1. Download the **isapi_redirect.dll** from the apache tomcat download page - click '**browse download area**' to search for the file.
    - **For example**, you will find the win32 binaries here: http://apache.wildit.net.au/tomcat/tomcat-connectors/jk/binaries/win32/ and if the current version is 1.2.27, you will download this file: http://apache.wildit.net.au/tomcat/tomcat-connectors/jk/binaries/win32/jk-1.2.27/isapi_redirect-1.2.27.dll.

    ⚠ Make sure you rename the file to `isapi_redirect.dll` before using it otherwise it *will not work*.
2. Place the **isapi_redirect.dll** file in a directory `c:\ajp_iis` (the name of the directory isn't important, but if you use a different one make sure to take account of this in the instructions which follow)
3. Create a **isapi_redirect.properties** file in the same directory as you put the DLL. You can use this sample isapi_redirect.properties file if you have used `ajp_iis` as the directory name.
    - Note that this sample properties file assumes that the dll is named `isapi_redirect.dll`. If you want to name your DLL something else, you must edit this file.
4. Create **workers.properties** and **uriworkermap.properties** files. You can use the sample workers.properties file and the sample uriworkermap.properties file.
5. Create an **empty** file named **rewrites.properties** in `c:\ajp_iis`.

Step 5. Connect Confluence with IIS

Connect Confluence with IIS, depending on your version of IIS:

- To connect Confluence with IIS 5.1 or 6 please refer to Connecting Confluence with IIS 5.1 or 6.
- For IIS 7 please refer to Connecting Confluence with IIS 7.

RELATED TOPICS

- Connecting Confluence with IIS 5.1 or 6
- Connecting Confluence with IIS 7
- JIRA's documentation on Configuring IIS with Tomcat, including how to integrate both Confluence and JIRA with the same IIS instance. The Troubleshooting section there is relevant to Confluence as well as JIRA.
- Supported Platforms

Connecting Confluence with IIS 5.1 or 6

⚠ The content on this page relates to platforms which are not supported for Confluence. Consequently, Atlassian **can not 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.

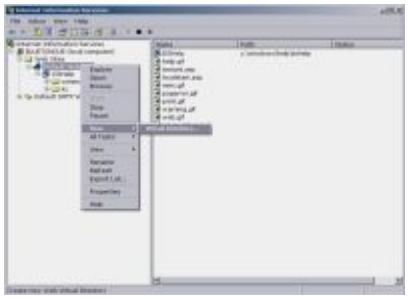
This documentation is part of the Setting up Confluence with IIS documentation.

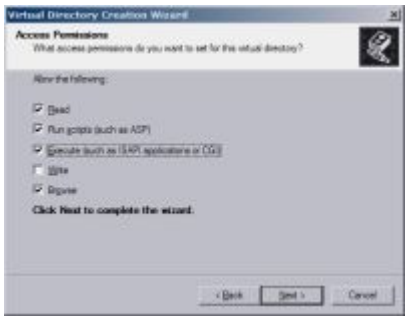
On this page:

- Connecting Confluence with IIS
- Other Configuration
- IIS 6.0
- Troubleshooting

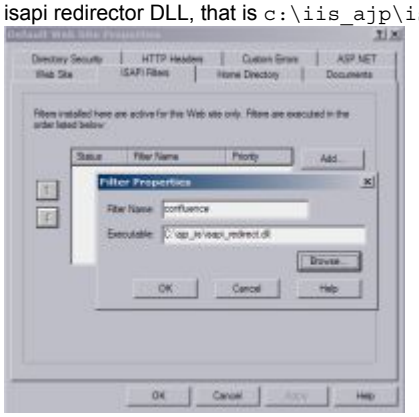
Connecting Confluence with IIS

1. Using the IIS management console (Internet Information Services in Administrative Tools), add a new **Virtual Directory** to your IIS web site. The name of the `virtual directory` must be **jakarta**, as it must correspond with the first part of the **extension_uri** setting specified in the **isapi_redirect.properties** file. Its physical path should be the directory where you placed **isapi_redirect.dll** (in the example it is `C:\ajp_iis`). When creating this new virtual directory, give it *execute* access as well.

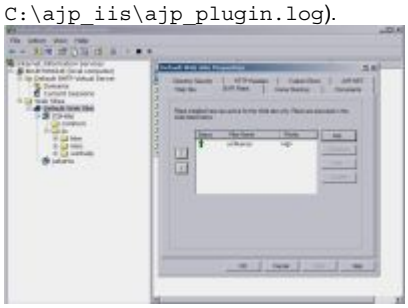




2. Using the IIS management console, add **isapi_redirect.dll** as a filter to your IIS web site. To do this, right click on the Web Sites icon from the left hand pane of the Internet Services Manager (or IIS management console), select Properties and then the `ISAPI Filters` tab. The name of the filter should reflect its task, for example confluence, and its executable must be the full path to the ajp isapi redirector DLL, that is `c:\iis_ajp\isapi_redirect.dll`



3. Restart IIS (stop and start the IIS service - not just the web server -- do this by right-clicking on 'Local Computer' in the IIS Manager and choosing All Tasks, Restart IIS...). Ensure that the confluence filter is marked with a green up arrow verifying that it is loaded and initialized correctly. If the ajp redirector did not initialize properly, check the log file for errors messages (`C:\ajp_iis\ajp_plugin.log`).



You can now go to http://localhost and see the Confluence Dashboard.

Other Configuration

If you want to run Confluence on a named context, rather than the root context (i.e. access it via http://host/confluence/ instead of just http://host) you need to:

1. Change the path attribute of the Context tag in server.xml from "" to "/confluence".
2. Change the line `/*=ajp13w` in `uriworkermap.properties` to `/confluence/*=ajp13w`.
3. Note that 'http://host/confluence' gives a 404 error, but 'http://host/confluence/' works. You need to create a virtual directory so that requests *without the trailing slash* still work. If you are using Confluence you would want to name the alias as confluence (if you are using JIRA, name it as jira). The physical directory can be anywhere and does not need to contain anything.

IIS 6.0

1. If using IIS 6.0 you will also need to add the *Jakarta Isapi Redirector* to the **Web Service Extension's** .
2. Right-click on Web Service Extensions and choose Add a new Web Service Extension...
3. Enter tomcat for the Extension Name and then add the `isapi_redirect.dll` file to the required files.
4. Check the Set extension status to Allowed and then click on OK.
5. Also add the *Jakarta Isapi Redirector* to the **ISAPI Filters** for the website

Troubleshooting

If you have problems, look in the System Event Log, the c:\iis_ajp\ajp_plugin.log and your confluence logs.

When requesting support for IIS configuration problems, please include:

1. A zip of your logs directory

2. A zip of your `c:\iis_ajp` directory
3. Your `conf/server.xml` file
4. A screen shot of the Virtual Directory tab of the Properties window for your 'jakarta' virtual directory.
5. A screen shot of the ISAPI Filters tab of the Properties window of your 'Web Sites' icon.

RELATED TOPICS

- JIRA's documentation on Configuring IIS with Tomcat, including how to integrate both Confluence and JIRA with the same IIS instance. The Troubleshooting section there is relevant to Confluence as well as JIRA.
- Connecting Confluence with IIS 7
- Supported Platforms

Take me back to Setting up Confluence with IIS.

Connecting Confluence with IIS 7



The content on this page relates to platforms which are not supported for Confluence. Consequently, Atlassian **can not 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.

This documentation is part of the Setting up Confluence with IIS documentation.

On this page:

- Setting up Confluence with IIS 7
- Known Issues
  - 64 bit IIS
  - Double Escaped Character
- Troubleshooting

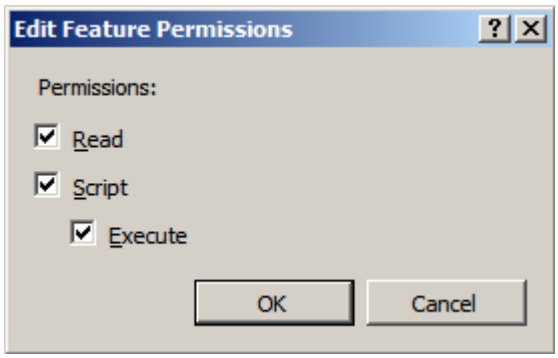
Setting up Confluence with IIS 7

After you have installed IIS 7, you need to configure the Tomcat Connector. Then follow these steps:

1. Open IIS 7 Manager
2. Navigate to your host. In the picture below, your host would be the one highlighted in blue just below "Start Page".



3. Double click on the `ISAPI and CGI Restrictions` icon
4. On the right hand panel, click on **Add...**. Click here for picture
5. Point the path to your `isapi_redirect.dll` file and give it a description eg. tomcat. Give it an execute permission by clicking on the `Allow extension path to execute` tick box.
6. Now, click on the `Default Website` and double click `ISAPI Filter`.
7. On the right hand panel, click on **Add...** and point to your `isapi_redirect.dll` file, give the filter a name (eg. tomcat)
8. Navigate to your Default Website again.
9. Right click on `Default Website` to create a virtual directory and name it `jakarta`, point this to your `c:\ajp_iis` directory.
10. Click on the newly created virtual directory jakarta and double click `Handler Mappings`.
11. Click on `Edit Feature Permissions` and tick the *execute* permission.



12. Add another virtual directory and name it `confluence`. Point the physical path to `c:\confluence`.
13. Set a context path in Confluence's xml configuration. For example, if you are using Confluence Standalone distribution you need to edit `<confluence install directory>/conf/server.xml` and edit your context path to this:

```
1.<Context path="/confluence" docBase="../confluence" debug="0" reloadable="false">
```

> The reason for creating this virtual directory is so that requests *without the trailing slash* still work. For example, if you are deploying Confluence under http://www.example.com/confluence/ without the virtual directory, then requests to http://www.example.com/confluence will fail.

14. Finally, navigate to your host context again and do a restart. Confluence should now be accessible via http://localhost/confluence

Known Issues

64 bit IIS

If you are running a 64 bit OS, please use a 64 bit version of the Tomcat IIS connector.

Customer submitted solution:
If you must use a 32 bit IIS connector, you can do so by clicking `Application Pools > Advanced Settings > Allow 32bit applications`.

Double Escaped Character

IIS 7 blocks double escaped character sequences by default. This will cause problems with Confluence pages with spaces. To fix this, please see http://support.microsoft.com/kb/942076.

Troubleshooting

The information in the Windows System Event Log can be useful for troubleshooting, followed by the `c:\ajp_iis\ajp_plugin.log` and your confluence logs.

When requesting support for IIS configuration problems, please include:

1. A zip of your logs directory
2. A zip of your `c:\ajp_iis` directory
3. Your `conf/server.xml` file

RELATED TOPICS

- JIRA's documentation on Configuring IIS with Tomcat, including how to integrate both Confluence and JIRA in the same IIS server. The Troubleshooting section there is relevant to Confluence as well as JIRA.
- Connecting Confluence with IIS 5.1 or 6
- Supported Platforms

Take me back to Setting up Confluence with IIS

Using the IBM 64bit J9 JDK

> The content on this page relates to platforms which are not supported for Confluence. Consequently, Atlassian **can not 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.

This JVM must be started with the argument: -Dsun.reflect.inflationThreshold=0

Otherwise you will see an error message like:

```
1.bucket.core.InfrastructureException: java.lang.NoClassDefFoundError:
com.atlassian.confluence.spaces.Space$$EnhancerByCGLIB$$58d74b80
2.      at
com.atlassian.confluence.util.XWorkTransactionInterceptor.intercept(XWorkTransactionInterceptor.java:
151)
3.caused by: java.lang.NoClassDefFoundError:
com.atlassian.confluence.spaces.Space$$EnhancerByCGLIB$$58d74b80
4.      at sun.reflect.GeneratedMethodAccessor311.invoke(Unknown Source)
```

Support Policies

Welcome to the support policies index page. Here, you'll find information about how Atlassian Support can help you and how to get in touch with our helpful support engineers. Please choose the relevant page below to find out more.

- Bug Fixing Policy
- How to Report a Security Issue
- New Features Policy
- Patch Policy
- Security Advisory Publishing Policy
- Security Patch Policy
- Severity Levels for Security Issues

To request support from Atlassian, please raise a support issue in our online support system. To do this, visit support.atlassian.com, log in (creating an account if need be) and create an issue under Confluence. Our friendly support engineers will get right back to you with an answer.

Bug Fixing Policy

Summary

- Atlassian Support will help with workarounds and bug reporting.
- Critical bugs will generally be fixed in the next maintenance release.
- Non critical bugs will be scheduled according to a variety of considerations.



Raising a Bug Report

Atlassian Support is eager and happy to help verify bugs — we take pride in it! Please open a support request in our support system providing as much information as possible about how to replicate the problem you are experiencing. We will replicate the bug to verify, then lodge the report for you. We'll also try to construct workarounds if they're possible.

Customers and plugin developers are also welcome to open bug reports on our issue tracking systems directly. Use http://jira.atlassian.com for the stand-alone products and http://studio.atlassian.com for JIRA Studio.

When raising a new bug, you should rate the priority of a bug according to our JIRA usage guidelines. Customers should watch a filed bug in order to receive e-mail notification when a "Fix Version" is scheduled for release.

How Atlassian Approaches Bug Fixing

Maintenance (bug fix) releases come out more frequently than major releases and attempt to target the most critical bugs affecting our customers. The notation for a maintenance release is the final number in the version (ie the 1 in 3.0.1).

If a bug is critical (production application down or major malfunction causing business revenue loss or high numbers of staff unable to perform their normal functions) then it will be fixed in the next maintenance release provided that:

- The fix is technically feasible (i.e. it doesn't require a major architectural change).
- It does not impact the quality or integrity of a product.

For non-critical bugs, the developer assigned to fixing bugs prioritises the non-critical bug according to these factors:

- How many of our supported configurations are affected by the problem.
- Whether there is an effective workaround or patch.
- How difficult the issue is to fix.
- Whether many bugs in one area can be fixed at one time.

The developers responsible for bug fixing also monitor comments on existing bugs and new bugs submitted in JIRA, so you can provide feedback in this way. We give high priority consideration to security issues.

When considering the priority of a non-critical bug we try to determine a 'value' score for a bug which takes into account the severity of the

bug from the customer's perspective, how prevalent the bug is and whether roadmap features may render the bug obsolete. We combine this with a complexity score (i.e. how difficult the bug is). These two dimensions are used when developers self serve from the bug pile.

Further reading

See How to Get Legendary Support from Atlassian for more support-related information.

How to Report a Security Issue

Finding and Reporting a Security Vulnerability

If you find a security bug in the product, please open an issue on http://jira.atlassian.com in the relevant project.

- Set the priority of the bug to 'Blocker'.
- Provide as much information on reproducing the bug as possible.
- Set the security level of the bug to 'Developer and Reporters only'.

All communication about the vulnerability should be performed through JIRA, so that Atlassian can keep track of the issue and get a patch out as soon as possible.

Further reading

See How to Get Legendary Support from Atlassian for more support-related information.

New Features Policy

Summary

- We encourage and display customer comments and votes openly in our issue tracking systems, http://jira.atlassian.com and http://studio.atlassian.com.
- We do not publish roadmaps.
- Product Managers review our most popular voted issues on a regular basis.
- We schedule features based on a variety of factors.
- Our Atlassian Bug Fixing Policy is distinct from our Feature Request process.
- Atlassian provides consistent updates on the top 20 feature/improvement requests (in our issue tracker systems).

How to Track what Features are Being Implemented

When a new feature or improvement is scheduled, the 'fix-for' version will be indicated in the JIRA issue. This happens for the upcoming release only. We maintain roadmaps for more distant releases internally, but because these roadmaps are often pre-empted by changing customer demands, we do not publish them.

How Atlassian Chooses What to Implement

In every major release we *aim* to implement highly requested features, but it is not the only determining factor. Other factors include:

- **Direct feedback** from face to face meetings with customers, and through our support and sales channels.
- **Availability of staff** to implement features.
- **Impact** of the proposed changes on the application and its underlying architecture.
- How **well defined** the requested feature is (some issues gain in popularity rapidly, allowing little time to plan their implementation).
- Our long-term **strategic vision** for the product.

How to Contribute to Feature Development

Influencing Atlassian's release cycle
We encourage our customers to vote on feature requests in JIRA. The current tally of votes is available online in our issue tracking systems, http://jira.atlassian.com and http://studio.atlassian.com. Find out if your improvement request already exists. If it does, please vote for it. If you do not find it, create a new feature or improvement request online.

Extending Atlassian Products
Atlassian products have powerful and flexible extension APIs. If you would like to see a particular feature implemented, it may be possible to develop the feature as a plugin. Documentation regarding the plugin APIs is available. Advice on extending either product may be available on the user mailing-lists, or at our community forums.

If you require significant customisations, you may wish to get in touch with our partners. They specialise in extending Atlassian products and can do this work for you. If you are interested, please contact us.

Further reading

See How to Get Legendary Support from Atlassian for more support-related information.

Patch Policy

Patch Policy

Atlassian will only provide software patches in extremely unusual circumstances. If a problem has been fixed in a newer release of the product, Atlassian will request that you upgrade your instance to fix the issue. If it is deemed necessary to provide a patch, a patch will be provided for the current release (e.g. JIRA 4.0) and the last maintenance release of the last major version (e.g. JIRA 3.13.5) only.

Patches are issued under the following conditions:

- The bug is critical (production application down or major malfunction causing business revenue loss or high numbers of staff unable to perform their normal functions).
- A patch is technically feasible (ie it doesn't require a major architectural change)
  OR
- The issue is a security issue, and falls under our Security Policy .

Atlassian does not provide patches for non-critical bugs.

Provided that a patch does not impact the quality or integrity of a product, Atlassian will ensure that patches supplied to customers are added to the next maintenance release. Customers should watch a filed bug in order to receive e-mail notification when a "Fix Version" is scheduled for release.

Patches are generally attached to the relevant http://jira.atlassian.com issue.

Further reading

See How to Get Legendary Support from Atlassian for more support-related information.

Security Advisory Publishing Policy

Publication of Security Advisories

When a security issue in an Atlassian product is discovered and resolved, Atlassian will inform customers through the following mechanisms:

- A security advisory will be posted in the documentation.
- A copy of the advisory will be sent to the product mailing-lists. These lists are mirrored on our forums.
- If the person who reported the issue wants to publish an advisory through some other agency (for example, CERT), Atlassian will assist in the production of that advisory, and link to it from our own.

Further reading

See How to Get Legendary Support from Atlassian for more support-related information.

Security Patch Policy

Unable to render {include} Couldn't find a page to include called: Support:Security Patch Policy

Further reading

See How to Get Legendary Support from Atlassian for more support-related information.

Severity Levels for Security Issues

Severity Levels

Atlassian security advisories include a severity level, rating the vulnerability as one of the following:

- Critical
- High
- Moderate
- Low

Below is a summary of the factors which we use to decide on the severity level, and the implications for your installation.

Severity Level: Critical

We classify a vulnerability as critical if most or all of the following are true:

- Exploitation of the vulnerability results in root-level compromise of servers or infrastructure devices.
- The information required in order to exploit the vulnerability, such as example code, is widely available to attackers.
- Exploitation is usually straightforward, in the sense that the attacker does not need any special authentication credentials or knowledge about individual victims, and does not need to persuade a target user, for example via social engineering, into performing any special functions.

Severity Level: High

We give a high severity level to those vulnerabilities which have the potential to become critical, but have one or more mitigating factors that make exploitation less attractive to attackers.

For example, given a vulnerability which has many characteristics of the critical severity level, we would give it a level of high if any of the following are true:

- The vulnerability is difficult to exploit.
- Exploitation does not result in elevated privileges.
- The pool of potential victims is very small.

Note: If the mitigating factor arises from a lack of technical details, the severity level would be elevated to critical if those details later became available. If your installation is mission-critical, you may want to treat this as a critical vulnerability.

Severity Level: Moderate

We give a moderate severity level to those vulnerabilities where the scales are slightly tipped in favour of the potential victim.

The following vulnerabilities are typically rated moderate:

- Denial of service vulnerabilities, since they do not result in compromise of a target.
- Exploits that require an attacker to reside on the same local network as the victim.
- Vulnerabilities that affect only nonstandard configurations or obscure applications.
- Vulnerabilities that require the attacker to manipulate individual victims via social engineering tactics.
- Vulnerabilities where exploitation provides only very limited access.

Severity Level: Low

We give a low severity level to those vulnerabilities which by themselves have typically very little impact on an organisation's infrastructure.

Exploitation of such vulnerabilities usually requires local or physical system access. Exploitation may result in client-side privacy or denial of service issues and leakage of information about organisational structure, system configuration and versions, or network topology.



Original ranking compiled by the SANS Institute
Our vulnerability ranking is based on a scale originally published by the SANS Institute.

Further reading

See How to Get Legendary Support from Atlassian for more support-related information.

Troubleshooting Problems & Requesting Technical Support

This document tells you how to troubleshoot problems and obtain technical support.

On this page:

- Troubleshooting a Problem
- Raising a Support Request
  - Raising a Support Ticket via the Internet
  - Support Utility
  - Raising a Support Ticket via the Confluence Support Request Form
- Logging a Bug Report
  - STEP 1. Check your Bug is Undiscovered
  - STEP 2. Confirm the Bug

Troubleshooting a Problem

If you have a problem with Confluence, please follow these steps:

1. If you are not a Confluence administrator, report your problem to the person in charge of your Confluence site and ask them to follow up on the issue.
2. Check our Frequently Asked Questions and our Knowledge Base for a solution to your problem.
3. Check the appropriate Confluence component in JIRA for known bugs.
4. If you are having problems configuring a feature, please take a look at the appropriate guides:
   - Confluence Installation Guide
   - Confluence Setup Guide
   - Confluence Administrator's Guide
   - Confluence Configuration Guide
   - Database Configuration
5. If your issue is related to your database server, please refer to the documentation within the Known Issues For Supported Databases section.
6. If your issue is related to your application server, please refer to the documentation within the Installing the Confluence EAR-WAR

Distribution section.
7. If your problem may be related to a plugin, you can enter Plugin Support Mode by briefly disabling the third party plugins.

If the above documentation does not solve your problem, you should create a support request and attach your support zip. If you believe you are experiencing a bug, you may wish to create a bug report instead. Instructions for both are given below.

Raising a Support Request

Please check the plugin support
If you have a plugin-related issue, please check whether the plugin is supported by Atlassian. Visit the plugin's home page in the Extension space, and check for the 'Atlassian Supported' logo. If the plugin is not supported by Atlassian, you will need to contact the author directly.

There are two ways to raise a support request with Atlassian:

For pre confluence 3.1 (confluence 2.10, 3.0.x etc) instances please skip to the support request form via the confluence administration

- (Recommended) Raise a support ticket directly via our support site on the Internet and then create a support zip using the Support Utility below. The advantage of this method is that it includes all the relevant files confluence support need. You can also be sure that the support case has been created and includes your logs.
- Complete the support request form via your Confluence **Administration Console**. The disadvantage of this method is that your mail may not be forwarded correctly due to restrictions imposed by your mail server. For example, the zip of your log files might be too large for your mail server to forward them on.

Both of these methods are described below.

Raising a Support Ticket via the Internet

If your Confluence instance is not configured with SMTP mail or your Confluence instance is not running, you can raise a support ticket via the Atlassian Support System:

1. Create a support zip to attach to the ticket. If your instance does not start up, refer to Working with Confluence Logs for information about the log files.
2. If your problem concerns user management or performance, please take a look at the additional requirements in Requesting External User Management Support or Requesting Performance Support.
3. If you do not already have a free Atlassian support account, create one here.
4. Log in to https://support.atlassian.com and select '**Create New Issue**'.
5. Lodge a detailed description of your problem in the new support ticket.
6. Fill in all applicable information about your system, such as application server, database, etc.
7. If Confluence is running, go to the '**System Information**' screen in your **Administration Console** and copy the text of your system information into the ticket.
8. Once your ticket is lodged, wait to be notified by email of updates. If your production instance of Confluence is experiencing a critical problem, jump on Live Support and ask to have your issue reviewed immediately.

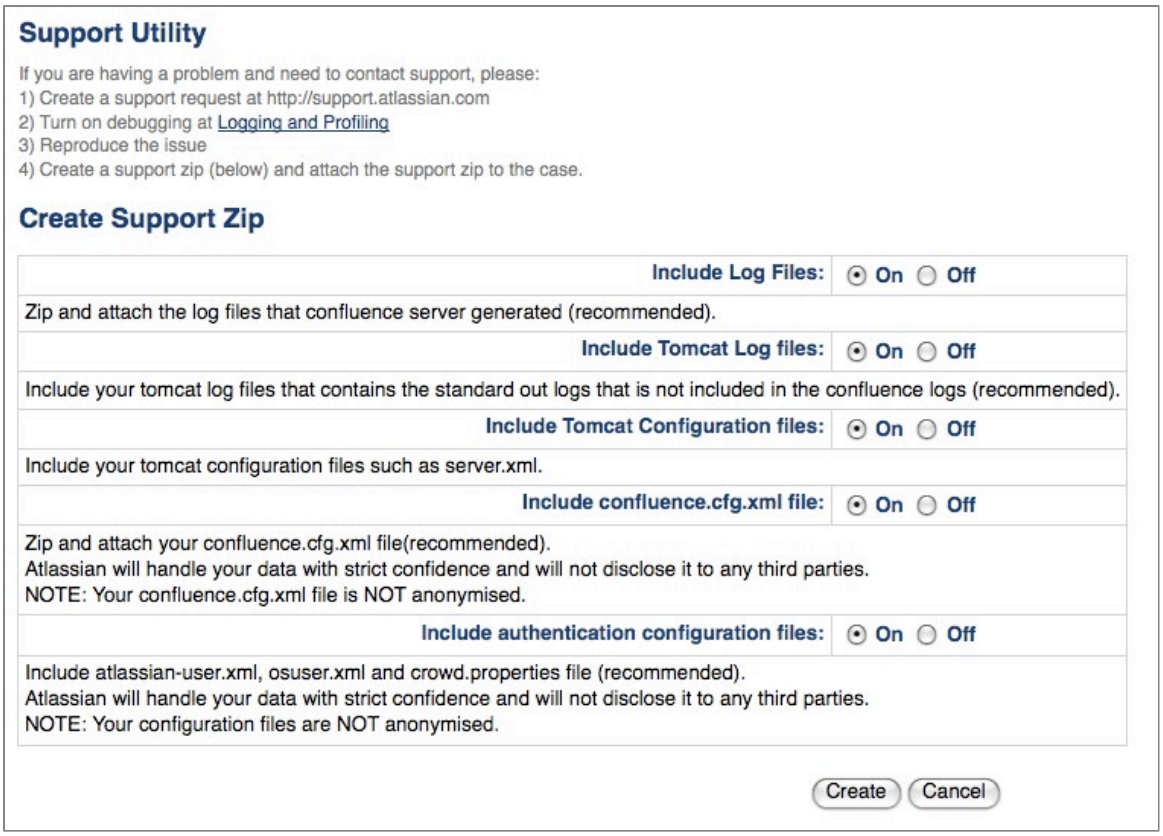
Support Utility

It is recommended that you add a support zip to every interaction with confluence support. The utility will also dump your system information to the logs before zipping them.

You can also use this method to append system information to an existing support ticket.

1. Log in as a user with System Administrator or Confluence Administrator access.
2. Go to the **Administration Console** and click '**Support Utility**' under '**Administration**' in the left-hand panel. Ensure that everything is checked, then click on "Create" button.
3. Attach the created support zip, to the support case you raised.

The Create Support Zip form is shown below:



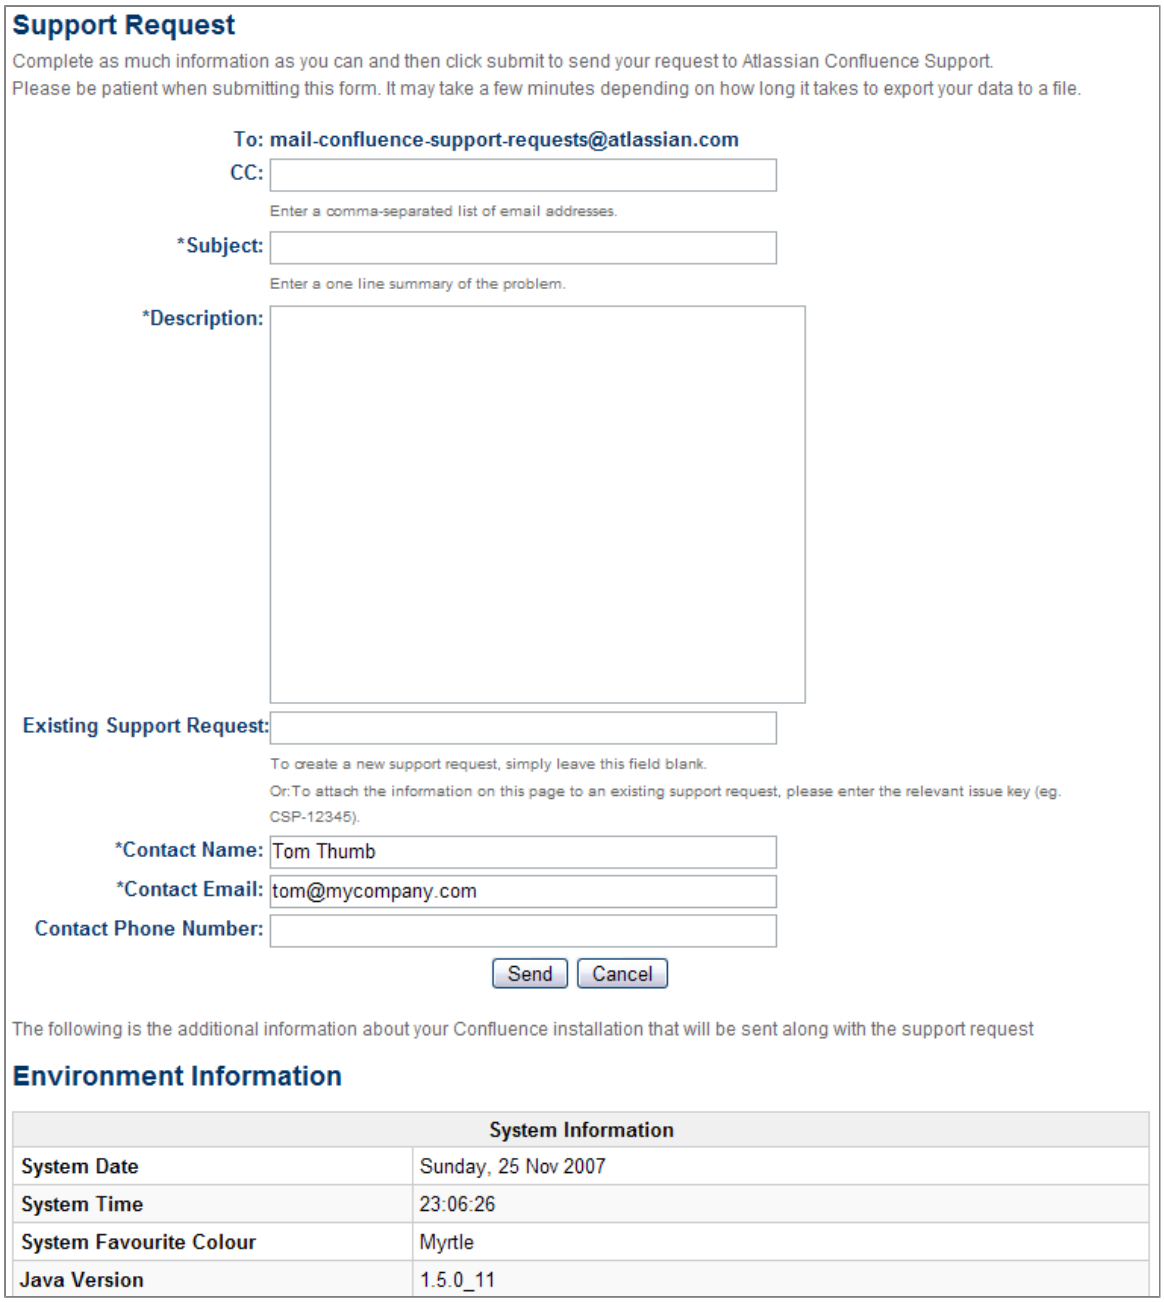
Raising a Support Ticket via the Confluence Support Request Form

Ensure that SMTP email is set up on your Confluence instance and your mail server allows zip files.

The advantage of this method is that it is convenient, however the disadvantage is that your mail may not be forwarded correctly due to an issue (e.g. zip file too large) or security restriction on your mail server.

You can also use this method to append system information to an existing support ticket.

1. Log in as a user with System Administrator or Confluence Administrator access.
2. Go to the **Administration Console** and click '**Support Utility**' under '**Administration**' in the left-hand panel. Then click on "create a support request via the confluence administration" link. The '**Raise Support Request**' form will appear. Part of the form is shown below:



Support Request

Complete as much information as you can and then click submit to send your request to Atlassian Confluence Support.
Please be patient when submitting this form. It may take a few minutes depending on how long it takes to export your data to a file.

To: mail-confluence-support-requests@atlassian.com

CC:

Enter a comma-separated list of email addresses.

***Subject:**

Enter a one line summary of the problem.

***Description:**

Existing Support Request:

To create a new support request, simply leave this field blank.
Or:To attach the information on this page to an existing support request, please enter the relevant issue key (eg. CSP-12345).

***Contact Name:** Tom Thumb

***Contact Email:** tom@mycompany.com

Contact Phone Number:

Send Cancel

The following is the additional information about your Confluence installation that will be sent along with the support request

Environment Information

System Information	
System Date	Sunday, 25 Nov 2007
System Time	23:06:26
System Favourite Colour	Myrtle
Java Version	1.5.0_11

3. Please provide as much information as possible, following these guidelines:
    - '**To**' — This is an email address, named the Site Support Address and configured on the '**General Configuration**' screen of your Confluence instance. It points to a JIRA instance (usually the Atlassian Support System) which is configured to receive and handle support requests by email.
    - '**CC**' — Any email address(es) entered here will receive a copy of the support request, including all system information. You can enter more than one email address, separated by commas (e.g. `joe@mycompany.com, sally@mycompany.com, jane@myothercompany.com`).
    - '**Subject**' — Enter a short and meaningful description of the problem.
    - '**Description**' — Please enter as much information as possible, including any error messages that are appearing and any steps the support team can take to reproduce the problem.
    - '**Existing Support Request**' — If you have previously raised a support request for the problem, please type the issue key here (e.g. CSP-12345). The information on this form will be appended to the existing support ticket.
    - '**Contact Name**' — This will default to the name of the logged-in user.
    - '**Contact Email**' — This will default to the email address of the logged-in user.
      Note: This email address will be used to find your support account on the Atlassian Support System. If no matching account is found, a new account will be created. Confluence will also send all further notifications and updates to this address.
    - '**Contact Phone Number**' — Please enter a telephone number where our support staff can reach you. Include international and city codes.
4. Click the '**Send**' button.
5. Confluence will submit your request via email to the JIRA instance referenced by the 'To' email address on the form. If you do not already have a support account, Confluence will automatically request one for you. The submitted request will include all the system and environment information which you see on the support request form. It will also include a zipped copy of your Confluence log file. Refer to Working with Confluence Logs for information about the log files. JIRA will create a support ticket including the submitted information.

   Log files can be very big — it is possible that your email server may bounce the message if too large. With the default log4j configuration, the log file could be up to 20Mb in size. If you have customised the log settings, the maximum size could be much

larger still. Please check whether the email message has been successfully sent, and consult your email administrator if you need special provisions for this email message.
6. Once you have submitted your support request, you will receive email updates about its progress. These emails will give you the support ticket number.

You can view the status of your support request and add any additional information required by visiting the Atlassian Support System at any time.

Logging a Bug Report

If you have found a bug, the easiest way to report it is to:

- Create numbered instructions on how to reproduce the bug.
- Log them as a support request.

The Atlassian support team will confirm your bug and lodge a bug report. Alternatively, you can log a bug report directly by confirming it according to these instructions:

STEP 1. Check your Bug is Undiscovered

1. Visit the Confluence bug tracker.
2. On the left under '**Text Search**', type keywords for your problem into the Query field.
3. Click View and browse the summaries of the unresolved bugs. If any summary appears to describe your problem, check that the bug is not a match. If it is the same, you may wish to set a watch to be notified of updates or apply your vote towards having it resolved. If the problem does not already appear to have been logged, the next step is to confirm that the problem is a bug.

STEP 2. Confirm the Bug

Check the headings below. If one of the headings matches your problem, follow the instructions. If the problem does not fall under any category, follow the general instructions instead.

Bug Affects Page Rendering or Content

If you are having issues with Wiki Markup or page content not being shown as expected:

1. Create a new page in the Confluence Sandbox and try to duplicate the issue there.
2. If the problem recurs, log the new bug here.
3. Paste the web address (URL) of the Sandbox page along with the process you used to duplicate the problem.
4. If the issue does not occur, this is may not be a bug and you should log the problem as a support request instead.

This is a temporary page visible to **Atlassian staff only**. At the release of Confluence 3.1, its contents will be copied to the parent page (and this temporary one deleted).

Bug Prevents Confluence from Starting

Please lodge a support request with your configuration information and numbered instructions on how to reproduce the issue.

Bug in External User Management

Please lodge a support request with your user management configuration and numbered instructions on how to reproduce the bug.

General Bug Confirmation

1. If Confluence will not run, please log the problem as a support request instead.
2. Attempt to replicate the bug:
    - Download the latest version of Confluence Standalone.
    - Install the Standalone with the appropriate database and the demonstration site.
    - Try to duplicate your problem on the default setup with the demonstration data.
3. If the issue does not occur, you should open a support ticket.
4. If the issue does occur, log the new bug here along with the information you used to duplicate it.
5. Once your issue is lodged, wait to be notified by email of updates. If your production instance of Confluence is experiencing a critical problem, jump on Live Support and ask to have your issue reviewed immediately.

500 page

Sometimes it may be useful to include the result of the `500page.jsp`

```
https://<domain><host>:<port>/500page.jsp
```

RELATED TOPICS

General Support Enquiries
Requesting External User Management Support
Requesting Performance Support
Feature Requests
Configuring the Site Support Address
Site Configuration
Administrators Guide Home Confluence Documentation Home

Content Anonymiser for Data Backups

Introduction



A Jira data anonymiser is also available.

Atlassian may request a copy of the `entities.xml` file from a customer's exported zip file, in order to diagnose database corruption, or to find a bug in Confluence.

If your data is confidential, you can run this program over your `entities.xml` file, removing all your data and leaving only the structure of the export.

Usage

To run the anonymiser on your backup:

1. Download the anonymiser JAR.
2. Extract the `entities.xml` file from your zipped backup file to the same directory as the JAR.
3. Use the command prompt to go to the directory where all three files are located.
4. To create `cleaned.xml`, run the command:

```
1.java -jar confluence-export-cleaner-1.1-jar-with-dependencies.jar entities.xml cleaned.xml
```

How it works

The Content Anonymiser application replaces all the text content in file `entities.xml` with 'x' characters. For example, the word "Atlassian" will be transformed to "xxxxxxxxx". The resulting `cleaned.xml` file is expected to have the same size of the original file.

This release of the export cleaner uses STX, a fast and efficient XML transformation technology. It should not require a lot of memory to run, even for a large backup.

Development

For Atlassian developers:

- Source code
- Maven repository.

Editing or Deleting a Page That Won't Render



You may be able to access the edit page URL by hitting ctrl+e

If you have a page that you can't access (for example, due to an incompatible plugin that won't render a macro), you can delete or edit the page by manually entering the appropriate URL. The URL looks like this:

```
http://<baseurl>/pages/removepage.action?pageId=<pageID>
http://<baseurl>/pages/editpage.action?pageId=<pageID>
http://<baseurl>/pages/editblogpost.action?pageId=<pageID>
```

Substitute your page ID for the one you wish to delete. To determine the page ID, you may be able to access it from the edit page URL by hitting ctrl+e. If not, you can obtain this information from the database using an SQL query like this:

```
SELECT CONTENTID FROM content WHERE TITLE = '<pagename>' AND VERSION = '1';
```

This may return multiple results if there are pages with the same name in different spaces, so you may have to further determine the correct one.

To delete an attachment manually, you can use a URL like:

```
http://<baseurl>/pages/removeattachment.action?pageId=32787&fileName=harbour.jpg&version=1
```

To view the attachments on a page:

```
http://<baseurl>/pages/viewpageattachments.action?pageId=<pageId>
```

Get the page ID similarly.

To get the wiki markup from the database directly, try:

```
1.SELECT BODY FROM BODYCONTENT WHERE CONTENTID IN (SELECT CONTENTID FROM content WHERE TITLE =
'<insert name of page or blog post>');
```

Enabling detailed SQL logging

Confluence uses the open source persistence framework Hibernate. This page tells you how to configure Confluence's logging to report individual SQL requests being sent to the database by Hibernate. It is useful for troubleshooting:

- XML site backups that fail to import
- Exceptions caused by an illegal database operation

Enable SQL logging via the Administration Console
Since the 2.7 release, you can also enable SQL logging at runtime via the Administration Console — read the instructions. This runtime option does not allow you to enable logging of SQL parameter values.

To enable detailed SQL logging in Confluence, you need to modify `log4j.properties`, located in `confluence/WEB-INF/classes`.

After you have enabled hibernate logging, please replicate the action that is causing the error in the first place. This is the best way to ensure that the Confluence log file contains relevant SQL logging.

If you require support assistance with a database related problem, it is advisable to enable detailed SQL logging **before** sending us the log files. This will assist us in determining what SQL queries were running during the reported problem.

If the entries mentioned below are not defined in the `log4j.properties file`, you can manually add the entries to the file in the 'Hibernate Logging' section.

To Log SQL Queries

Stop Confluence, then uncomment the following lines in `log4j.properties`:

```
## log hibernate prepared statements/SQL queries (equivalent to setting 'hibernate.show_sql' to
'true')
log4j.logger.net.sf.hibernate.SQL=DEBUG, confluencelog
log4j.additivity.net.sf.hibernate.SQL=false
```

To Log SQL Queries with Parameters

Stop Confluence, then uncomment the following lines in `log4j.properties`:

```
## log hibernate prepared statement parameter values
log4j.logger.net.sf.hibernate.type=DEBUG, confluencelog
log4j.additivity.net.sf.hibernate.type=false
```

To Disable Batched Updates for Simpler Debugging

Stop Confluence, then edit `databaseSubsystemContext.xml`:

- In Confluence 2.5.x and earlier, the `databaseSubsystemContext.xml` file is at `confluence/WEB-INF/classes/databaseSubsystemContext.xml`

- From Confluence 2.6.x, the `databaseSubsystemContext.xml` file is available in the `confluence-2.6.0.jar` file located in the <confluence-install>/WEB-INF/lib directory.

Uncomment the `<prop>` line in the following location:

```
<!-- it can be useful to disable batching during debugging, as HSQLDB doesn't report the exact
statement which fails in batch mode -->
<prop key="hibernate.jdbc.batch_size">0</prop>
```

RELATED TOPICS

Troubleshooting SQL exceptions
Working with Confluence Logs

General Support Enquiries

For information on Confluence features and configuration.

Online Documentation

- Confluence Setup Guide
- Confluence Administrator's Guide
- Confluence Configuration Guide
- Online Forum. For general discussion, plugin development, customisation, new features and issues.
- Subscribe to the Mailing List. Mailing list posts are archived on the Online Forum and posts to the forum online are sent to the Mailing List, so you can use either method.

Support Requests

Check out Troubleshooting Problems & Requesting Technical Support

Feature Requests

View Creating a Feature Request

Live Support

Live Support is available during US, Malaysian and Australian business hours, 19 hours a day, Monday to Friday. Downtime is 9am to 2pm GMT. Click to see 9am and 2pm GMT in your local time. Live Support is unavailable on weekends.

If you are experiencing a problem, you should always create a support request before using Live Support.

Generating a Thread Dump

- Stack Traces and Security
- Generating a Thread Dump Externally when Confluence stops responding
- Generating a Thread Dump via the Administration Console
- Scheduling Thread Dumps via the Administration Console

If Confluence is performing poorly, behaving unexpectedly or stops responding and you can generate a thread dump to help diagnose the cause of the problem. Furthermore, if you wish to contact Atlassian Support for assistance about it, you should include a thread dump in your support enquiry to help the Support team determine the cause of the problem.

A thread dump will show the state of each thread in the JVM, including a stack trace. Thread dumps are only useful if they are taken at the appropriate time. They normally need to be taken at or close to the time when the application is experiencing problems.
Information about what locks are being held and waited upon by a thread are not produced by Confluence's Thread Dump tool. If you require this information, then generate a thread dump externally.

Stack Traces and Security

To help debug support cases and provide legendary support, Confluence provides stack traces through the web interface when an error occurs. These stack traces include information about what Confluence was doing at the time, and some information about your deployment server.

Only non-personal information is supplied such as operating system and version and Java version. With proper network security, this is not enough information to be considered dangerous. No usernames or passwords are included.

Generating a Thread Dump Externally when Confluence stops responding

If Confluence stops responding or you require information on locks being held and waited upon by threads, then use one of methods

described in Generating a Thread Dump Externally.

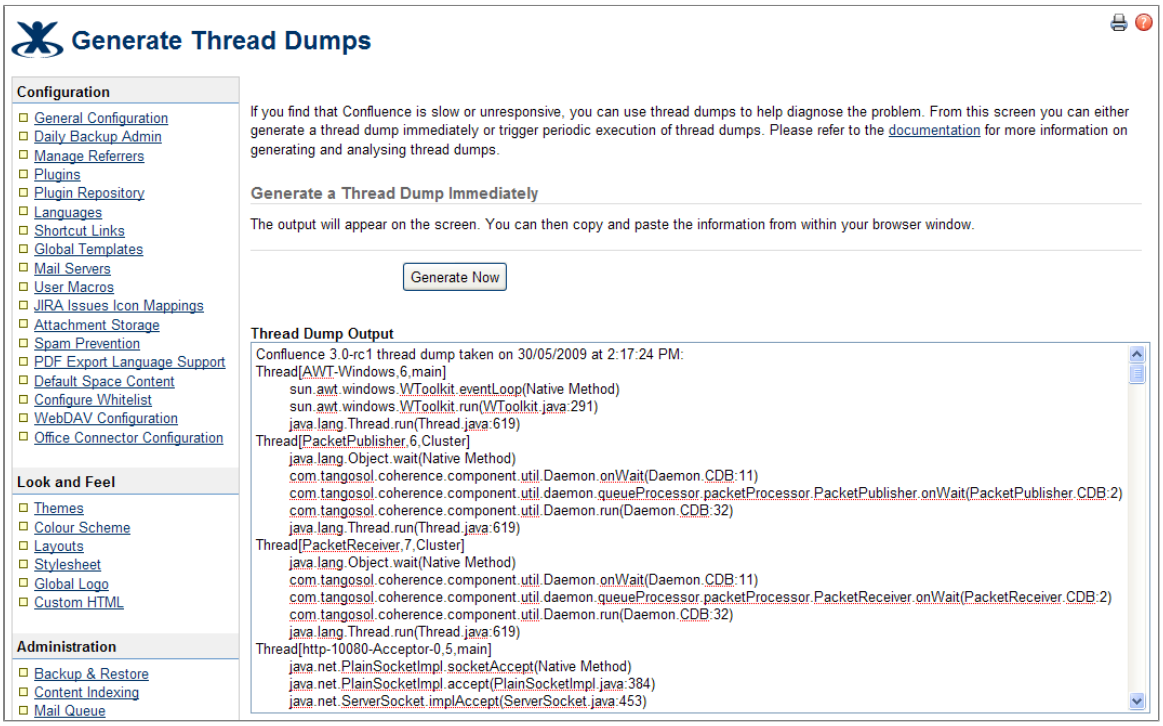
Generating a Thread Dump via the Administration Console

> **For Confluence 2.10.3 or below**
> This feature was introduced in Confluence 3.0. if you are using a prior version then please consult this documentation on Generating a Thread Dump Externally.

To generate a Thread Dump from the Administration Console,

1. Go to the Confluence '**Administration Console**'. To do this:
    - Open the '**Browse**' menu and select '**Confluence Admin**'. The 'Administration Console' view will open.
2. Select '**Thread Dump**' in the left-hand panel.
3. Click the '**Generate Now**' button in the centre of the page. The output is displayed in a new text box that appears just below the button.
4. Copy the contents of the thread dump in the text box and save it to a text file.

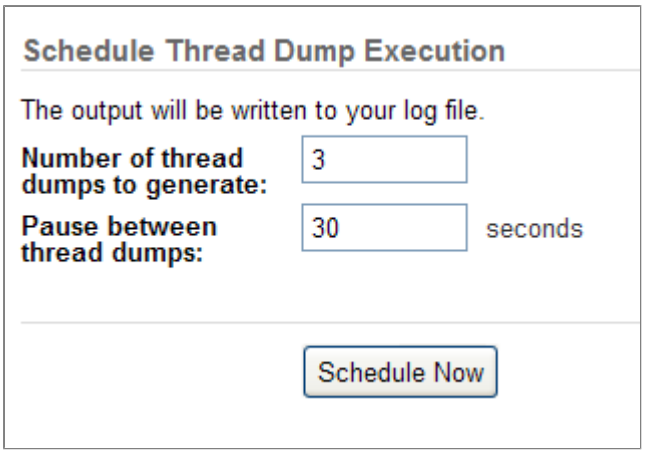
Screenshot: Example of a generated thread dump from the Confluence administration console



Scheduling Thread Dumps via the Administration Console

If you were asked by Atlassian Technical Support to generate regular thread dumps, please set the Thread Dump Scheduler to take 2 to 3 thread dumps with a 30 seconds time interval in between so the Support team can observe any important patterns that may assist with the diagnosis of the problem. Attach the log file to the support ticket.

Example: Scheduling thread dumps from the Confluence administration console



Generating a Thread Dump Externally

If Confluence stops responding and you cannot access its integrated Generate Thread Dump feature, it is possible to create thread dumps outside the application. External thread dumps are also useful if you require information on locks being held or waited upon by threads.

Generating a Thread Dump on Windows

To take a thread dump from Windows:

1. Visit http://www.adaptj.com/root/main/download and click Launch
2. Click Run for any security warnings
3. Under Process Id, select the '...' button.
4. From the drop-down list, select the Confluence process. Users running Confluence Standalone, select the 'Java (Tomcat) ...' option. Users running Confluence WAR should select their application server process.
5. Under Process -> Properties, Ensure that the "Thread dump" and "Keep Remote Thread Running" is selected.
6. Select Process -> Thread Dump
7. Save the output to a file, eg 'threaddump.log'

If you were asked by Atlassian technical support to create the thread dump, please take 2 to 3 thread dumps with a time interval in between (eg. 30 seconds) so we can see some patterns. Attach the log file to the support ticket.

Alternatively, if you are **not** running Confluence as a service, click on the console and press `<CTRL>+BREAK`

Generating a Thread Dump on Linux, including Solaris and other Unixes

Find the process ID of the JVM and use the **ps** command to get list of all processes:

```
kill -3 <pid>
```

Note: This will not kill your server (so long as you included the "-3" option, no space in between).
The thread dump will be printed to Confluence's standard output (catalina.out).

Generating Thread Dumps using jstack

If you have trouble using kill - 3 <pid> to obtain a thread dump, or are on a Windows machine without internet access then try using jstack: a java utility that will output stack traces of Java threads for a given process.

1) Identify the **java** process that JIRA is running in:

On *nix and Solaris
This can be achieved by running a command similar to:

ps -ef | grep javaThe process will appear similarly as follows:

```
adam 22668 0.3 14.9 1691788 903928 ? Sl Jan27 9:36 /usr/lib/jvm/java-6-sun-1.6.0.14/bin/java
-Djava.util.logging.config.file=/home/adam/Products/installs/atlassian-jira-enterprise-4.0.1-standalone/conf/logging.properties
-XX:MaxPermSize=256m -Xms128m -Xmx1048m -Djava.awt.headless=true -Datlassian.standalone=JIRA
-Dorg.apache.jasper.runtime.BodyContentImpl.LIMIT_BUFFER=true -Dmail.mime.decodeparameters=true
-Datlassian.mail.senddisabled=false -Datlassian.mail.fetchdisabled=false
-Djava.util.logging.manager=org.apache.juli.ClassLoaderLogManager
-Djava.endorsed.dirs=/home/adam/Products/installs/atlassian-jira-enterprise-4.0.1-standalone/common/endorsed -classpath
/home/adam/Products/installs/atlassian-jira-enterprise-4.0.1-standalone/bin/bootstrap.jar
-Dcatalina.base=/home/adam/Products/installs/atlassian-jira-enterprise-4.0.1-standalone
-Dcatalina.home=/home/adam/Products/installs/atlassian-jira-enterprise-4.0.1-standalone
-Djava.io.tmpdir=/home/adam/Products/installs/atlassian-jira-enterprise-4.0.1-standalone/temp
org.apache.catalina.startup.Bootstrap start
```

On windows
Launch the task manager by, pressing `Ctrl + Alt + Del` and find the Process ID of the JAVA (Confluence) process

2) Run jstack <pid> to Capture a Single Thread Dump:

This command will take one thread dump of the process id <pid>, in this case the pid is 22668,

```
1.adam@track:~$ jstack 22668
```


if the jstack executable is not in your $PATH, then please look for it in your <JDK_HOME>/bin directory

3) **Take Multiple Thread Dumps**

Typically you'll want to take several dumps about 10 seconds apart, in which case you can generate several dumps and output the stack traces to a single file as follows:

```
1.adam@track:~$ jstack 22668
2.adam@track:~$ jstack 22668
3.adam@track:~$ jstack 22668
```

Output

Standard logging for Confluence Stand-alone is sent to the atlassian-confluence.log, in the confluence-home directory, not in the confluence-install directory. Thread dumps are an exception since they dump the threads of the entire application server - they'll appear in the catalina.out file in the application directory's logs folder. You can search for the term "thread dump" in the log file for the beginning of the dump. Submit this along with the atlassian-confluence.log in your support ticket.

Thread Dump Tools

- Samurai
- Thread Dump Analyzer TDA TDA 1.0 Final can be obtained from the `java.net`

Plugin Support Mode


Incompatible with Confluence Clustered
At this time, Plugin Support Mode will not work correctly in a Confluence cluster.

Beginning with Confluence version 2.9 the Plugin Repository Client now offers "Support Mode" for troubleshooting plugin-related issues. It will temporarily disable all unsupported plugins and re-enable them with a single click.

1. From `Administration > Plugin Repository`, choose the `Admin` tab.
2. Click "Enter Support Mode."

To re-enable the plugins, click the same link to restore the plugins to their previous state.

If you get errors when trying to re-enabled plugins, this is because some third-party plugins cannot be disabled and then re-enabled at runtime. Restart Confluence to restore your complete list of plugins.

Profiling using the YourKit Plugin

There is a plugin for Confluence 2.2 and later which allows easy profiling using the YourKit profiler. No license is required to generate a memory or CPU snapshot, but you will need at least an evaluation license to analyse the results.

JIRA also has a plugin to profile JIRA's CPU and memory usage with YourKit.

On this page:

- Configuring YourKit in your JVM
  - Windows Configuration
  - Linux/Mac OS X Configuration
- Performance Impact
- Installing the YourKit Plugin
- Why would I do this?
- Plugin Source Code

Configuring YourKit in your JVM

Download YourKit 6.0 for your platform and follow the installation instructions to install it.

Note: YourKit version 7 is not compatible with the Confluence yourkit plugin.

The following instructions apply to Confluence Standalone and Tomcat installations with Sun JDK 1.5. They should be easily applicable to other application servers and JVMs, however. The YourKit documentation covers this in more detail.

Windows Configuration

On Windows, add to your PATH environment variable the `bin/win32/` directory underneath the YourKit installation directory. For example, you might add "C:\Program Files\YourKit Java Profiler 6.0.12\bin\win32" to your PATH, via Control Panel, System, Advanced, Environment Variables.

To configure Confluence's JVM to use the YourKit agent, you need to add a parameter to JAVA_OPTS in the `bin/setenv.bat` file in your Confluence application directory. This file controls the startup parameters for Tomcat, so you'll need to restart it after making the changes.

Add the 'agentlib' parameter to the end like this:

```
set JAVA_OPTS=%JAVA_OPTS% -Xms128m -Xmx256m -agentlib:yjpagent
```

Linux/Mac OS X Configuration

On Unix-based systems, include the installation directory in the library path environment variable, as shown below:

- For the Mac: `export DYLD_LIBRARY_PATH=$DYLD_LIBRARY_PATH:/path/to/yourKitAgent`
- For other Unix-based systems: `export LD_LIBRARY_PATH=$LD_LIBRARY_PATH:/path/to/yourKitAgent`

In general, to configure a Sun 1.5 JDK, you add the `agentlib` parameter:

```
java -agentlib:yjpagent ...
```

You can add this to Tomcat's `bin/setenv.sh` like this:

```
JAVA_OPTS="-Xms128m -Xmx256m $JAVA_OPTS -Djava.awt.headless=true -agentlib:yjpagent "
```

Performance Impact

Running YourKit can have detrimental effects on performance.

To minimize performance problems use the following modifications to the agentlib parameter:

```
-agentlib:yjpagent=disablecounts,disablealloc,disablej2ee
```

See also Profiling overhead: how to reduce or avoid in the YourKit documentation.

Installing the YourKit Plugin

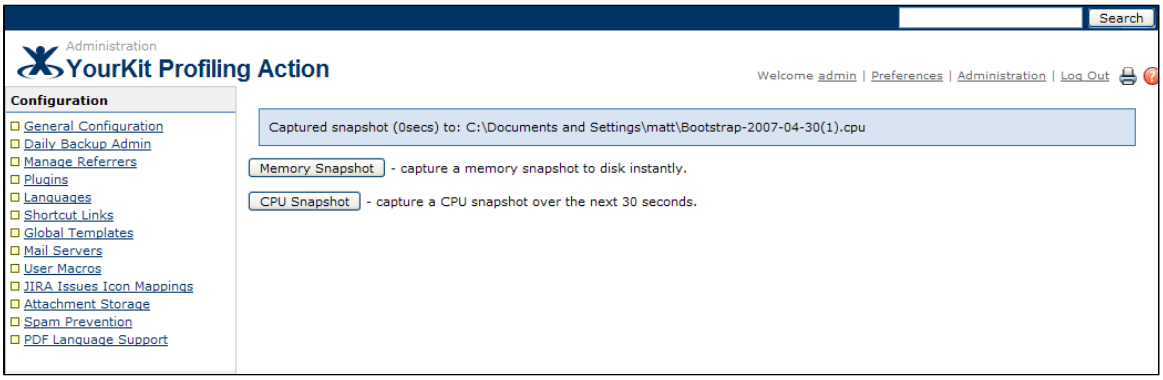
Download the plugin and upload it into Confluence through the Administration, Plugins page.

A new menu option will appear under the 'Administration' heading. Click it and you should see the options to take a memory or CPU snapshot.



YourKit Profiling menu item

This profiler dump will be saved to a local temp directory, and the path shown once it is complete. For the CPU snapshot, this will take at least 30 seconds. For the memory snapshot, 10-15 seconds.



You can take either a memory or CPU snapshot

Why would I do this?

Analysing a profiler dump enables Atlassian Support (or you, if you are interested) to see exactly what is happening in your application: what classes are using the memory, what is using CPU and so on. This can help us debug tricky performance problems which would otherwise be impossible to analyse remotely.

Take a CPU snapshot if:

- Confluence is sometimes unresponsive
- Pages take a long time to load
- Confluence's CPU usage is peaking.

Take a memory snapshot if:

- Confluence's memory usage is higher than you expect
- You are getting OutOfMemoryError's in your logs.

If you run into situations where Confluence is unresponsive and you are not able to trigger a memory snapshot, please ensure that you start Confluence with the `onexit=memory` parameter in the JVM options (as in the example below) and simply shut down Confluence. Before shutting down a memory snapshot will be created.

```
-agentlib:yjpagent=onexit=memory
```

Plugin Source Code

The source code for this Confluence plugin is available from Subversion and browseable in Fisheye. The JAR produced by 'mvn package' includes a copy of the YJP redistributable bundled in META-INF/lib/.

Confluence Resources

Resources for Evaluators

- Free Trial
- Feature Tour

Resources for Administrators

- Confluence Knowledge Base
- Confluence FAQ
- Tips of the Trade
- Guide to Installing an Atlassian Integrated Suite
- The big list of Atlassian gadgets

Downloadable Documentation

- Confluence documentation in PDF, HTML or XML formats

Plugins

- Documentation for the Confluence SharePoint Connector
- Atlassian Plugin Exchange
- Library of Confluence Plugins and Extensions

Support

- Atlassian Support
- Support Policies

Training

- Atlassian Training

Forums

- Confluence Announcements | subscribe
- Confluence General Forum | subscribe
- Confluence Developers Forum | subscribe

Feature Requests

- Issue Tracker and Feature Requests for JIRA

Tips of the Trade

Below are some links to external blog posts, videos and articles containing technical tips and instructions on setting up and using Confluence. This page presents an opportunity for customers and community authors to share information and experiences.

The references here are specific to Confluence wiki and are technical 'how to' guides written by bloggers who use Confluence. For general information on wiki comparisons, wiki adoption, best practices and business cases, please refer to the Atlassian website and to our evaluator resources.



Please be aware that these are external blogs and articles.
Most of the links point to external sites, and some of the information is relevant to a specific release of Confluence. Atlassian provides these links because the information is useful and relevant at the time it was written. Please check carefully whether the information is still relevant when you read it, and whether it is relevant to your version of Confluence. Unless explicitly stated, Atlassian does not offer support for third-party extensions or plugins. The information in the linked blog posts has not been tested or reviewed by Atlassian. We recommend that you test all solutions on a **test** server before trying it on your production site.

On this page:

- Tracking Atlassian Confluence usage with Google Analytics
- Moving Confluence from Windows to (Ubuntu) Linux
- Plugging Memory Leaks in Confluence
- Using a wiki for technical documentation
- Wiki docs --- release management
- Using a wiki for online help
- Content re-use on a wiki
- Starting out with your technical documentation on a wiki
- Universal Wiki Converter - Now with SSL Support
- Scroll converts wiki to DocBook and PDF
- Playing with DITA2Confluence part 1 and part 2
- The Confluence Reporting HOWTO
- Drawing diagrams on a wiki page
- Organisation is Key
- Creating FAQs
- Styling Tabs in Confluence 2.10

- How to determine the context your macro is being rendered in
- Video: Confluence overview
- Video: Macro browser

Administration

Tracking Atlassian Confluence usage with Google Analytics

And Using the Google Analytics Javascript API to show pageviews from Atlassian Confluence

- By: David Simpson, on blog 'david simpson'
- About: Setting up Google Analytics for Confluence
- Date: 18 March 2009 and 11 September 2009
- Related documentation:
    - How Do I Get More Statistics From Confluence?
    - How to audit Confluence - enabling user access logging

Moving Confluence from Windows to (Ubuntu) Linux

- By: Ricky Sheaves, on blog 'flimflam' (calebscreek)
- About: Moving Confluence to its own dedicated environment: Ubuntu 8.04 with a MySQL backend
- Date: 19 October 2008
- Related documentation: Migrating Confluence Between Servers

Plugging Memory Leaks in Confluence

- By: Don Willis, on blog 'Atlassian developer blog'
- About: Identifying memory leaks in Confluence and fixing them
- Date: 1 October 1007
- Related documentation: Performance Tuning

Confluence for Technical Documentation

Using a wiki for technical documentation

- By: Sarah Maddox, on blog 'ffeathers'
- About:
  - Overview — what a wiki is and does.
  - Workflow — draft, review, publish.
  - Tracking — page history, notification of updates, reverting to a previous version.
  - Permissions.
  - Adding structure to your documentation — table of contents, left-hand navigation bar, logical page ordering, content re-use.
  - Release management on a wiki.
  - Using spaces for version control.
  - How a wiki is useful in agile development.
- Date: 21 November 2009

Wiki docs --- release management

- By: Sarah Maddox, on blog 'ffeathers'
- Date: 17 November 2007
- About:
  - Using spaces for version control
  - Release management on a wiki
  - Archiving documentation on a wiki
- Related documentation: The Copy Space plugin

Using a wiki for online help

- By: Sarah Maddox, on the 'Atlassian Blog'
- About: Pointing online help links to version-controlled wiki documentation spaces
- Date: 13 December 2007

Content re-use on a wiki

- By: Sarah Maddox, on blog 'ffeathers'
- About: Content reuse and defining an inclusions library
- Related documentation:
  - Excerpt Macro
  - Excerpt Include Macro
  - Include Page Macro
- Date: 29 July 2008

Starting out with your technical documentation on a wiki

- By: Sarah Maddox, on blog 'ffeathers'
- About: Choosing your wiki and planning your documentation
- Date: 4 November 2007

Content Conversion

Universal Wiki Converter - Now with SSL Support

- By: Laura Kolker, on the 'Atlassian Blog'
- About: Configuring the UWC for two new features:
  - A Trac Converter module
  - SSL support
- Date: 6 March 2009
- Related documentation: Importing Content from another Wiki

Scroll converts wiki to DocBook and PDF

- By: Sarah Maddox, on blog 'ffeathers'
- About: Using the Scroll Wiki Exporter plugin to convert Confluence content to DocBook XML and to PDF
- Date and Confluence version: 18 May 2009; Confluence 2.10
- Related documentation: The Scroll Wiki Exporter plugin

Playing with DITA2Confluence part 1 and part 2

- By: Sarah Maddox, on blog 'ffeathers'
- About: Using the DITA2Confluence tool to convert documentation from DITA XML to Confluence pages
- Date and Confluence version: October 2008; Confluence 2.9
- Related documentation: The DITA2wiki project on SourceForge

Usage Tips

The Confluence Reporting HOWTO

- By: Jim Severino and John Rotenstein, Atlassian Internal Systems, on the 'Atlassian Blog'
- About:
  - Using Confluence as a reporting and business intelligence tool
- Date and Confluence version: August 2009; Confluence 3.0
- Related documentation: The Confluence Reporting HOWTO

Drawing diagrams on a wiki page

- By: Sarah Maddox, on blog 'ffeathers'
- About:
  - Using the Gliffy plugin to draw diagrams on a Confluence page
  - Links to other tools for displaying flowcharts, graphs etc based on editable content in the wiki page
- Date and Confluence version: 4 July 2009; Confluence 3.0
- Related documentation: The Gliffy plugin

Organisation is Key

- By: Matt Hodges, on the 'Atlassian Blog'
- About: Designing the structure of a Confluence space using an inclusions library, macros and tabbed pages
- Date and Confluence version: 17 March 2009; Confluence 2.10

Creating FAQs

- By: Matt Hodges, on the 'Atlassian Blog'
- About: Designing the FAQ (frequently asked questions) section of your Confluence space
- Date and Confluence version: 2 April 2009; Confluence 2.10

Styling and Customisation

Styling Tabs in Confluence 2.10

- By: Jens Schumacher, on the 'Atlassian Blog'
- About: Using CSS to change the look of the tabs in Confluence
- Date and Confluence version: 12 January 2009; Confluence 2.10
- Related documentation: Styling Confluence with CSS

Development

How to determine the context your macro is being rendered in

- By: Cheryl Jerozal, on the 'Atlassian Blog'
- About: Discovering find out the current render context (including PDF document, feed reader, email notification, etc) from within your macro
- Date and Confluence version: 25 June 2009; Confluence 3.0
- Related documentation: Macro Module

Videos

To speed up the loading of the page and ensure correct export to PDF, HTML and XML formats, we will just link to the videos rather than including them into the wiki page.

Video: Confluence overview

- By: Matt Hodges, on the Atlassian website
- About:
  - Confluence Overview
  - Create and edit
  - Tracking updates and blogs
  - Security and permissions
  - Attachments and the Office Connector
  - Search and discover
  - Plugins
- Date and Confluence version: July 2009; Confluence 3.0
- Related documentation: Confluence documentation

Video: Macro browser

- By: David Cook, on the 'Atlassian Blog'
- About: Using the new Confluence macro browser in Confluence 3.0
- Date and Confluence version: 18 June 2009; Confluence 3.0
- Related documentation: Working with the Macro Browser

Have you written a technical tip for Confluence?
Add a comment to this page, linking to your blog post or article. We will include it if the content fits the requirements of this page.

Feedback?
Your first port of call should be the author of the linked blog post. If you want to let us know how useful (or otherwise) a linked post is, please add a comment to this page.

Other Sources of Information

Confluence documentation
Evaluator resources
Atlassian website
Atlassian forums
Atlassian blog
Confluence plugins

Information For Evaluators

Useful reading for anyone new to our wiki:

Basics

Confluence Evaluator Resources
Example Sites From Customers
Independent Reviews
Live Webinar Demonstration
Confluence Presentation Material
Case Studies

Customisation

Plugins
Document Management
Knowledge Base
Solutions Tour

Confluence Presentation Material

Feel free to use or adapt some of our presentation materials to help WOW the boss and build a business case for Confluence.

You can find more useful content in the Confluence Evaluator Resources space.

Related Content

Information For Evaluators
Live Demonstration Webinars

TreeNavigation

Index

TreeNavigationVersions

Click for all versions
Confluence 3.2.x
Confluence 3.1.x
Confluence 3.0.x
Confluence 2.10.x
Confluence 2.9.x
Confluence 2.8.x
Confluence 2.7.x
Confluence 2.6.x
Confluence 2.5.4 to 2.5.8
Confluence 2.0 to 2.5.3

.bookmarks

Confluence Bookmarks
This page is a container for all the bookmarks in this space. Do not delete or move it or you will lose all your bookmarks.
Bookmarks in Confluence | Links for Confluence

The 15 most recent bookmarks in Confluence Docs 3.2

There are no bookmarks to display.