MICRO FOCUS

Retain

Retain

July 2017

Legal Notices

For information about legal notices, trademarks, disclaimers, warranties, export and other use restrictions, U.S. Government rights, patent policy, and FIPS compliance, see https://www.novell.com/company/legal/.

Copyright © 2017 Micro Focus Software, Inc. All Rights Reserved.

Contents

Preface 13

1 Retain 4.2.0.0 15

1.1 Overview . . . 15
- 1.1.1 Intended Audience . . . 15
- 1.1.2 What Retain does . . . 15
- 1.1.3 How Retain works . . . 15
- 1.1.4 Architecture . . . 15
- 1.1.5 Best Practices for Component Placement . . . 17
- 1.1.6 How Retain Stores the Archives . . . 17
- 1.1.7 Other components that Retain depends on . . . 17
- 1.1.8 Design Considerations . . . 18
- 1.1.9 SQL Database Server . . . 18
- 1.1.10 Retain Server . . . 18
- 1.1.11 Retain Worker . . . 18
- 1.1.12 Retain Router . . . 18
- 1.1.13 Messaging System Address Book . . . 19
- 1.1.14 Retain Functions . . . 19

1.2 Retain Services . . . 20
- 1.2.1 Email Systems . . . 20
- 1.2.2 Smart Phone Communications . . . 21
- 1.2.3 Social Platforms . . . 21
- 1.2.4 Micro Focus GWAVA Family of Message Handling Products . . . 21
- 1.2.5 Technical Support . . . 22
- 1.2.6 Sales . . . 22
- 1.2.7 Copyright Notice . . . 23

2 Upgrading to Retain 4 25

2.1 Upgrading to Retain 4 . . . 25

2.2 Upgrading from Retain 3.5.1 or 4.x to 4.y . . . 25
- 2.2.1 Upgrade changes . . . 25
- 2.2.2 Requirements . . . 25
- 2.2.3 GroupWise Module . . . 26
- 2.2.4 Upgrade Process . . . 28
- 2.2.5 Post-Install Tasks . . . 29
- 2.2.6 Tool Compatibility with Retain 3.x (Lucene) Indexes . . . 35

3 Install 37

3.1 Retain Planning and Design Best Practices . . . 37
- 3.1.1 Retain Architecture . . . 37
- 3.1.2 Archive Job Flow . . . 38
- 3.1.3 Hardware Design . . . 40

3.2 System Requirements . . . 51
- 3.2.1 Minimum System Requirements . . . 51
- 3.2.2 Single Server . . . 52
- 3.2.3 Remote Database . . . 52
- 3.2.4 *High Availability Indexer . . . 53
- 3.2.5 Software installed by Retain . . . 53
- 3.2.6 Supported Linux Versions . . . 53

3.2.7 Supported Windows Versions . . . 53
3.2.8 Microsoft Exchange Requirements . . . 53
3.2.9 Novell GroupWise Requirements . . . 54
3.2.10 Android Requirements . . . 54
3.2.11 CellTrust Secureline . . . 54
3.2.12 BlackBerry Requirements . . . 54
3.2.13 Supported Social platforms . . . 54
3.2.14 Supported Search . . . 54
3.2.15 Supported SQL Database Management Software . . . 54
3.2.16 Supported Browsers . . . 55
3.3 Ports in use . . . 55
3.3.1 Retain Server . . . 55
3.3.2 Retain Worker . . . 55
3.3.3 Retain Message Router . . . 56
3.3.4 Retain Reporting & Monitoring Server . . . 56
3.3.5 Retain Stubbing Server . . . 56
3.3.6 Retain for Social Media (RSM) Proxy Server . . . 56
3.3.7 From RSM WAN IP to Untrusted, all TCP/UDP ports . . . 56
3.3.8 Database Management Systems . . . 57
3.3.9 Blackberry . . . 57
3.3.10 Exchange . . . 57
3.3.11 Google Apps Module . . . 57
3.3.12 GroupWise . . . 58
3.3.13 Mobile . . . 58
3.3.14 Office 365 . . . 58
3.3.15 CellTrust Secureline . . . 58
3.4 Estimating Storage Requirements . . . 58
3.4.1 Considerations for Storage Requirements . . . 59
3.4.2 Storage Minimums . . . 60
3.4.3 Indexes During Migration . . . 60
3.4.4 Change the Log File Location . . . 61
3.5 Starting and Stopping Tomcat . . . 62
3.6 Tomckat Memory Tuning . . . 63
3.6.1 General Recommendations . . . 63
3.6.2 Tomcat Configuration Parameters . . . 63
3.6.3 Checking memory . . . 66
3.7 SQL Database Guides . . . 68
3.7.1 SQL Quick Start Guide . . . 68
3.8 Create a new database for Retain . . . 69
3.8.1 MySQL . . . 69
3.8.2 MySQL Connector/J (JDBC driver for MySQL) Installation . . . 70
3.8.3 MS SQL Server 2008 R2, 2012, 2014 . . . 71
3.8.4 ORACLE 10/11g/12c . . . 71
3.8.5 Installation . . . 72
3.8.6 TCPIP Listener . . . 72
3.8.7 Database Creation . . . 72
3.8.8 User Account Creation . . . 73
3.8.9 Tuning . . . 73
3.8.10 Postgres 9 . . . 73
3.8.11 Further tuning . . . 74
3.9 Linux Installation . . . 74
3.9.1 Prerequisites . . . 74
3.9.2 Retain Install . . . 74
3.9.3 Custom Storage Path Rights . . . 79
3.9.4 Installing Just a Worker to a Server . . . 80
3.9.5 Installing Additional Workers . . . 81
3.10 Windows Installation . . . 82
3.10.1 Prerequisites . . . 82
3.10.2 Webserver Install . . . 83

3.10.3 Installing ISAPI extensions and filters with IIS . . . 83
3.10.4 Retain Install . . . 84
3.10.5 Controlling Tomcat . . . 93
3.10.6 Installing Just a Worker to a Server . . . 93
3.10.7 Installing Additional Workers . . . 94
3.11 Post-Install Tasks . . . 95
3.11.1 If using MySQL . . . 96
3.11.2 Initial Setup Overview . . . 96
3.12 Initial Configuration . . . 97
3.12.1 Moving Retain . . . 97
3.12.2 Run the Configuration Wizard . . . 97
3.12.3 About Storage Paths . . . 97
3.12.4 Standard Storage Path . . . 98
3.12.5 Advanced Storage Options . . . 98
3.13 Front Page and First Login . . . 103
3.13.1 Logging In . . . 103
3.13.2 Status and Updates . . . 104
3.13.3 Saving Changes . . . 106
3.13.4 Notifications . . . 106
3.14 Licensing . . . 107
3.14.1 Download Your License . . . 107
3.15 Backing Up Retain . . . 109
3.15.1 Where Data Is Stored In Retain . . . 109
3.15.2 Backing Up Retain . . . 111
3.16 To uninstall Retain . . . 113

4 Modules and Jobs 115

4.1 Modules and Job Overview . . . 115
4.1.1 Modules . . . 115
4.1.2 Jobs . . . 117
4.2 GroupWise Module . . . 118
4.2.1 Prerequisites . . . 118
4.2.2 Core Settings . . . 119
4.2.3 SOAP . . . 121
4.2.4 LDAP . . . 122
4.2.5 Proxy . . . 123
4.2.6 GroupWise Profile . . . 124
4.2.7 Profile Functions . . . 125
4.2.8 Messaging System Deletion . . . 126
4.2.9 Stubbing . . . 126
4.2.10 Archive Mark . . . 127
4.2.11 Scope . . . 128
4.2.12 Date Range to Scan . . . 129
4.2.13 About GroupWise's Purge and Retention Flags . . . 130
4.2.14 Setting up Purge and Retention flags in GroupWise . . . 131
4.2.15 Miscellaneous . . . 131
4.2.16 Advanced . . . 132
4.2.17 Folder Scope . . . 133
4.3 Exchange Module Office365 . . . 134
4.3.1 Prerequisites . . . 134
4.3.2 *Test Autodiscover* . . . 134
4.3.3 Get the SMTP URL . . . 134
4.3.4 Set up DNS SRV Record . . . 134
4.3.5 Set Rolling In-Place Hold . . . 135
4.3.6 Setup Users . . . 135
4.3.7 Retain Impersonation User . . . 135
4.3.8 Retain Administrator User . . . 136
4.3.9 PowerShell Sync Script 1.0 . . . 137

4.3.10 Setup Process Overview . . . 137
4.3.11 Setup . . . 137
4.3.12 PowerShell Sync Script 4.2 . . . 138
4.3.13 Configure SharePoint Online/OneDrive for Archiving . . . 141
4.3.14 Configure Exchange Module for Office365 . . . 144
4.3.15 Core Settings Tab . . . 145
4.3.16 Impersonation Tab . . . 147
4.3.17 Hosted Services Tab . . . 147
4.3.18 Refresh Address Book . . . 148
4.4 Exchange Profile . . . 148
4.4.1 Core Settings . . . 148
4.4.2 Messaging System Deletion . . . 149
4.4.3 Message Settings . . . 149
4.4.4 Scope . . . 150
4.4.5 Miscellaneous . . . 152
4.4.6 Advanced . . . 152
4.4.7 Folder Scope . . . 153
4.5 Exchange Module On-Premise Exchange . . . 154
4.6 Exchange Prerequisites required for Retain . . . 154
4.6.1 Create a Retain User . . . 155
4.6.2 Grant Impersonation Permissions to the Retain user . . . 155
4.6.3 Room and Equipment Resources . . . 156
4.6.4 Basic Authentication . . . 156
4.7 Configure Exchange Module for On-Premise Exchange . . . 159
4.7.1 Core Settings . . . 160
4.7.2 Impersonation . . . 162
4.7.3 Hosted Services . . . 162
4.7.4 Exchange Forest . . . 162
4.8 Finishing On-Premise Exchange . . . 165
4.8.1 Refresh Address Book . . . 166
4.9 Exchange Profile . . . 166
4.9.1 Core Settings . . . 166
4.9.2 Messaging System Deletion . . . 167
4.9.3 Message Settings . . . 167
4.9.4 Scope . . . 168
4.9.5 Miscellaneous . . . 170
4.9.6 Advanced . . . 170
4.9.7 Folder Scope . . . 171
4.9.8 Distributions Lists . . . 172
4.9.9 Exchange Message Dredging Process Overview . . . 172
4.10 BlackBerry Module . . . 174
4.10.1 Prerequisites . . . 174
4.10.2 BlackBerry Module Setup . . . 174
4.10.3 Configuring the Module . . . 174
4.10.4 BES Servers . . . 177
4.10.5 Caching . . . 178
4.10.6 BlackBerry Profile . . . 178
4.10.7 Scope . . . 179
4.10.8 Advanced . . . 180
4.10.9 Blackberry Jobs . . . 181
4.10.10 Core Settings . . . 181
4.10.11 Mailboxes . . . 182
4.10.12 Notification . . . 183
4.10.13 Status . . . 184
4.10.14 Device Management . . . 184
4.10.15 BlackBerry . . . 184
4.11 BlackBerry Auditing and Archiving Service for BBM . . . 186
4.11.1 Device Management . . . 192
4.11.2 BlackBerry . . . 192

4.12 Google Apps Module . . . 193
4.12.1 Prerequisite . . . 193
4.13 Google Apps Module Setup . . . 204
4.13.1 Core Settings Tab . . . 204
4.13.2 Settings Tab . . . 205
4.13.3 OpenID Tab . . . 206
4.13.4 Google Apps Profile . . . 207
4.13.5 Core Settings . . . 207
4.13.6 Message Settings . . . 208
4.13.7 Scope tab . . . 208
4.13.8 Miscellaneous tab . . . 210
4.13.9 Advanced tab . . . 210
4.13.10 Folder Scope . . . 211
4.14 Google Apps Jobs . . . 212
4.14.1 Core Settings tab . . . 213
4.14.2 Mailboxes tab . . . 214
4.14.3 Notification tab . . . 214
4.14.4 Status tab . . . 215
4.15 Bloomberg Module . . . 216
4.15.1 Prerequisites . . . 216
4.16 Module Configuration . . . 216
4.16.1 Core Settings tab . . . 216
4.16.2 Bloomberg Servers tab . . . 217
4.16.3 Bloomberg Profile . . . 218
4.16.4 Profile Functions . . . 219
4.16.5 Scope tab . . . 220
4.16.6 Date Range to Scan . . . 221
4.16.7 Advanced . . . 221
4.17 Mobile Module . . . 223
4.18 CellTrust SecureLine . . . 235
4.19 Generation 1 . . . 236
4.19.1 Core Settings . . . 236
4.19.2 SecureLine Servers . . . 237
4.19.3 Archiving/Data Collection – About Workers, Profiles, Schedules, and Jobs . . . 238
4.19.4 Retain Job components . . . 238
4.19.5 Secureline Server Configuration . . . 239
4.20 Generation 2 . . . 240
4.20.1 Configure Retain Server Connection . . . 240
4.20.2 Install REST Licenses . . . 240
4.20.3 Create REST Collector . . . 241
4.20.4 Configure the Module . . . 242
4.21 GBS Notes Module . . . 243
4.21.1 Install Module License . . . 243
4.21.2 Configure Retain Server Connection . . . 243
4.21.3 Install REST Licenses . . . 244
4.21.4 Create REST Collector . . . 245
4.21.5 Configure the Module . . . 246
4.22 Social Messaging Module . . . 247
4.22.1 Install Module License . . . 247
4.22.2 Configure Retain Server Connection . . . 247
4.22.3 Install REST Licenses . . . 247
4.22.4 Create REST Collector . . . 248
4.22.5 Configure the Module . . . 249
4.23 Schedules . . . 253
4.24 Workers . . . 254
4.24.1 Worker Configuration Options . . . 255
4.24.2 Worker Configuration . . . 256
4.25 Jobs . . . 262

4.25.1 Core settings . . . 263
4.25.2 Notification . . . 264
4.25.3 Mailboxes . . . 266
4.25.4 GroupWise Job option - Reload . . . 266
4.25.5 Exchange Job Option - Journaling . . . 267
4.26 Litigation Hold . . . 268

5 Configuration 271

5.1 Configuration . . . 271
5.1.1 Server Configuration . . . 271
5.1.2 Communications tab . . . 272
5.1.3 Database Connectivity . . . 272
5.1.4 Notification . . . 272
5.1.5 Retain Server Connection . . . 273
5.1.6 Index tab . . . 273
5.1.7 Storage tab . . . 276
5.1.8 Accounts tab . . . 277
5.1.9 Maintenance tab . . . 279
5.1.10 Ignore Domains . . . 281
5.1.11 Logging tab . . . 281
5.1.12 Auditing Records . . . 281
5.1.13 Miscellaneous tab . . . 282
5.1.14 Export XML . . . 283
5.1.15 Module Forwarding . . . 283
5.1.16 REST API . . . 285
5.2 Rules Configuration . . . 287
5.2.1 Core Settings tab . . . 288
5.2.2 Conditions tab . . . 288
5.2.3 Actions tab . . . 289
5.3 Router Configuration . . . 290
5.3.1 Communications tab . . . 290
5.3.2 Storage tab . . . 291
5.3.3 Logging tab . . . 292
5.4 Storage Manager . . . 292
5.4.1 Partitions . . . 294
5.5 Stubbing Server . . . 295
5.5.1 Overview . . . 295
5.5.2 Prerequisites . . . 296
5.5.3 Stubbing Install . . . 297
5.5.4 Stubbing Configuration . . . 298
5.5.5 Retain . . . 301
5.5.6 Configure Stubbing Server . . . 304
5.5.7 Stubbing Server . . . 304
5.6 Performing Retain Searches Within The GroupWise Client . . . 306
5.6.1 Prerequisites . . . 306
5.7 Configure Stubbing To Remove Messages from GroupWise and Leave Only Stubs . . . 308
5.7.1 GroupWise Client . . . 310
5.8 Mailbox Mapping . . . 311
5.8.1 Mailbox Mapping Options . . . 311
5.8.2 Mailbox Mapping Enabled . . . 312
5.8.3 Never . . . 312
5.8.4 Once, Address Book Cache only . . . 312
5.8.5 Always, all processes . . . 312
5.8.6 XML File . . . 313
5.8.7 Mailbox Mapping View From The User Mailbox . . . 314
5.9 Reload Integration . . . 314
5.9.1 A Brief Review on How Reload Works . . . 315
5.9.2 How Retain Takes Advantage of Reload’s features . . . 316

5.9.3 Multiple Workers on One Server . . . 316
5.9.4 Timing . . . 317
5.9.5 Retain Settings . . . 317
5.10 How to Setup Reload to work with Retain . . . 318
5.10.1 IMPORTANT Notes for the Integration . . . 324

6 Management 327

6.1 Users . . . 327
6.1.1 Creating Users . . . 327
6.1.2 Offline Password . . . 327
6.1.3 Creating a new user . . . 328
6.1.4 Core Settings . . . 329
6.1.5 User Rights . . . 330
6.1.6 User Rights Summary . . . 332
6.1.7 User-level rights . . . 333
6.1.8 Address Book Selector . . . 334
6.1.9 GroupWise Proxy Support . . . 337
6.2 Groups . . . 337
6.2.1 Creating a New Group . . . 338
6.2.2 TIPS . . . 340
6.2.3 Groups – Miscellaneous . . . 341
6.3 Reporting and Monitoring Configuration . . . 342
6.3.1 Setup . . . 342
6.4 Reporting and Monitoring Home Page . . . 345
6.5 Reports . . . 347
6.5.1 Report and Schedule Report . . . 347
6.5.2 Reports . . . 348
6.5.3 Audit User Activity . . . 349
6.5.4 Export Audit Table . . . 351
6.5.5 Job Run Averages . . . 352
6.5.6 Job Run Summary . . . 353
6.5.7 Job Summary . . . 354
6.5.8 Archive Summary by User . . . 356
6.5.9 Server Utilization . . . 358
6.5.10 Job Run Mailbox Summary . . . 359
6.5.11 Job Mailbox Summary . . . 360
6.5.12 Schedule Report . . . 362
6.5.13 System Monitor . . . 364
6.5.14 Mailbox Error Monitoring . . . 365
6.5.15 Maintenance . . . 367
6.6 Mailbox Error Monitoring . . . 367
6.7 Audit and System Logs . . . 369
6.7.1 Audit Log . . . 369
6.7.2 System Log . . . 370
6.7.3 Location of log files . . . 370
6.8 Data Removal . . . 371
6.8.1 Mailbox Deletion . . . 371
6.8.2 Item Deletion . . . 374
6.8.3 Profiles . . . 375
6.8.4 Job . . . 377
6.8.5 Delete Messages . . . 380
6.8.6 Deletion Logging . . . 381

7 Retain User Guide 383

7.1 Retain Web Interface . . . 383
7.2 Browse . . . 384

7.2.1 Toolbar . . . 386
7.3 Search . . . 395
7.3.1 Wildcard Search . . . 396
7.3.2 Search In . . . 397
7.3.3 Item Type . . . 398
7.3.4 Item Source . . . 399
7.3.5 Date Range . . . 399
7.3.6 Mailboxes . . . 400
7.3.7 Tag . . . 402
7.3.8 Misc. . . . 402
7.4 Advanced Search . . . 402
7.4.1 RegEx and Wildcards . . . 405
7.5 Legacy Search . . . 406
7.6 Exported Items . . . 407
7.7 Tag Definitions . . . 407
7.8 Options . . . 408
7.8.1 User Rights Tab . . . 408
7.8.2 Mailboxes Tab . . . 409
7.8.3 Confidential Exceptions . . . 409
7.8.4 Miscellaneous Tab . . . 410
7.8.5 Tool Compatibility with Retain 3.x (Lucene) Indexes . . . 412

8 Tools 413

8.1 Tools Overview . . . 413
8.1.1 Tool Compatibility with Retain 3.x (Lucene) Indexes . . . 413
8.2 Retain Publisher and Viewer . . . 413
8.2.1 Retain Publisher Install . . . 414
8.3 Retain Publisher . . . 420
8.3.1 Retain Server Information Page . . . 422
8.4 Retain Viewer . . . 428
8.4.1 Forward . . . 433
8.4.2 Settings . . . 434
8.4.3 Redaction . . . 435
8.4.4 Strikeout/Unstrikeout and Blackout . . . 437
8.4.5 Searching the Archive . . . 437
8.4.6 PDF and PST Export . . . 439
8.4.7 Enabling Debug Logging . . . 441
8.5 M+ and Netmail to Retain Migrator . . . 442
8.5.1 System Requirements . . . 442
8.5.2 Directories and Logs . . . 442
8.5.3 Install . . . 443
8.6 Prerequisites . . . 447
8.6.1 Download Your License . . . 447
8.6.2 Install the REST license . . . 447
8.6.3 Create a REST Collector . . . 448
8.7 Quick Start Guide . . . 449
8.7.1 Interface . . . 449
8.7.2 Configure Migrator . . . 449
8.8 Migration Process . . . 452
8.8.1 Users File Generation . . . 452
8.8.2 Select Users to Migrate . . . 453
8.9 Outlook Plugin Administrators Guide . . . 455
8.9.1 Overview . . . 455
8.9.2 Supported Software . . . 455
8.9.3 System Requirements . . . 456
8.9.4 Installation . . . 456
8.9.5 Features and Use . . . 459

8.9.6 Download Settings . . . 462
8.9.7 Account Settings . . . 463
8.9.8 Clear Local Archive . . . 463
8.9.9 Retain Download . . . 463
8.9.10 Retain Search . . . 463
8.9.11 Easy Search . . . 464
8.9.12 Advanced Search . . . 465
8.9.13 Quick Search . . . 466
8.9.14 Retain Outlook Plugin Settings . . . 467
8.9.15 Launch Retain . . . 467
8.9.16 Data File . . . 468
8.9.17 Log File Location . . . 468
8.10 Outlook Plugin Silent Install . . . 468
8.11 Outlook Plugin User Guide . . . 469
8.11.1 Outlook Account Setup . . . 469
8.11.2 Using Retain Outlook plugin . . . 471
8.11.3 Retain Search . . . 472
8.11.4 Outlook Search . . . 474
8.11.5 Other buttons under the Retain tab . . . 476
8.11.6 Uninstall Retain Outlook client Plugin . . . 478
8.12 Outlook Web App Plugin . . . 478
8.12.1 Installation . . . 479
8.12.2 Enable Plugin for the Organization . . . 485
8.12.3 The Plugin in the Outlook Web App . . . 486
8.12.4 Change Features, Repair or Remove the Plugin . . . 487
8.13 Retain PST Importer . . . 488
8.13.1 System Requirements . . . 488
8.13.2 Installation . . . 489
8.13.3 Use . . . 495
8.13.4 Logs . . . 497
8.14 GroupWise WebAccess Plugin . . . 497
8.14.1 Overview . . . 497
8.14.2 Supported Software . . . 498
8.14.3 System Requirements . . . 498
8.14.4 Installation . . . 498
8.14.5 Post Install Tasks . . . 503
8.14.6 Confirmation . . . 503
8.14.7 Uninstalling the GroupWise WebAccess Plugin . . . 504
8.15 Retain for Social Messaging Gateway . . . 504
8.15.1 Setting up the VM & Installing the OS . . . 504
8.15.2 Configure the Virtual Machine . . . 504
8.16 RSM Configuration . . . 507
8.16.1 Proxy Configuration . . . 507
8.16.2 RSM CA Certificate Installation . . . 507
8.16.3 Social Information Governance Configuration . . . 508
8.16.4 Enabling Social Information Governance . . . 508
8.16.5 Creating Pre-Defined Social Information Governance Rules and Actions . . . 508
8.16.6 Creating Custom Social Information Governance Rules . . . 509
8.16.7 Creating Custom Social Information Governance Actions . . . 509
8.16.8 Social Information Governance Moderation . . . 510
8.16.9 Creating a Moderation Queue . . . 510
8.16.10 Creating a Moderation Action . . . 510
8.16.11 Configuring Secure Social Media Authentication . . . 510
8.16.12 Configuring RSM to authenticate to a company’s social media site . . . 511
8.17 GroupWise Client Plugin . . . 511
8.17.1 System Requirements . . . 512
8.17.2 Install . . . 512
8.18 GroupWise WebAccess . . . 519
8.18.1 Requirements . . . 519

8.18.2 Install . . . 520
8.19 GroupWise Archive Migration Tool . . . 529
8.19.1 GroupWise Personal Archive Migration Tool . . . 529
8.19.2 System Requirements . . . 529
8.19.3 Overview; Modes of Operation . . . 529
8.19.4 How it Works . . . 529
8.19.5 Prerequisites . . . 530
8.19.6 Setting It All Up . . . 530
8.19.7 Archiving Behavior . . . 531
8.19.8 User Interface . . . 532
8.19.9 Distribution Lists . . . 533
8.19.10 User List . . . 534
8.19.11 Errors . . . 535
8.19.12 Logging . . . 537
8.19.13 Connections . . . 538
8.19.14 Running the Migration . . . 540
8.19.15 Limit of 5 Concurrent Migrations . . . 540
8.19.16 Other Notes . . . 541

9 About Retain 543

9.1 About and Diagnostics Page . . . 543
9.1.1 Diagnostic Buttons . . . 544
9.1.2 Core info, JAR List, Jobs, Schedules, Profiles, Work-ers . . . 544
9.1.3 Quartz . . . 544
9.1.4 Retention Date Utility . . . 544
9.1.5 Indexer Status . . . 544
9.1.6 Exchange Test . . . 544
9.1.7 Mailbox Merger / Splitter . . . 545
9.1.8 Indexer Test . . . 546
9.1.9 System Properties . . . 546
9.2 Retain Component Licensing . . . 546
9.2.1 Apache 2.0 License . . . 547
9.3 LGPL – GNU Lesser General Public License . . . 550
9.4 BSD License – Individual Component Licenses . . . 552
9.4.1 ANTLR 3 License . . . 552
9.4.2 ASM . . . 553
9.4.3 DOM4J . . . 553
9.4.4 JDOM . . . 554
9.4.5 XStream . . . 554
9.4.6 PAL Artistic License . . . 555
9.4.7 Bouncy Castle APIs . . . 557
9.4.8 PDF Box . . . 557
9.4.9 WSDL4J . . . 558
9.4.10 BSD License . . . 562
9.4.11 Open Software License v. 3.0 . . . 562
9.4.12 Mozilla Public License Version 1.1 . . . 565
9.4.13 GNU GENERAL PUBLIC LICENSE . . . 571
9.4.14 Bitstream Vera Fonts Copyright . . . 576
9.4.15 Arev Fonts Copyright . . . 577
9.4.16 Java Util Concurrent Public Domain . . . 578
9.4.17 SLF4J MIT License, below . . . 580

Preface

Retain

About This Guide

This Retain Unified Archiving Administrator's Guide helps you integrate this software into your existing email system.

Audience

This manual is intended for IT administrators in their use of Retain or anyone wanting to learn more about Retain . It includes installation instructions and feature descriptions.

Feedback

We want to hear your comments and suggestions about this manual and the other documentation included with this product. Please use the User Comment feature at the bottom of each page of the online documentation.

Additional Documentation

Online documentation can be found on the Micro Focus (https://www.microfocus.com/products/) website.

Knowledge Base articles can be found on the Micro Focus Knowledge Base (https://www.microfocus.com/support-and-services/knowledge-base/) website.

Technical Support

If you have a technical support question, please consult the Micro Focus Technical Support at http://support.gwava.com (http://support.gwava.com)

Sales

Micro Focus contact information and office locations: www.microfocus.com (http://www.microfocus.com)

To contact a Micro Focus sales team member, please e-mail info@gwava.com (mailto:info@gwava.com) or call 866-GO-GWAVA ((866) 464-9282), or +1 (514) 639-4850 in North America.

Copyright Notice

The content of this manual is for informational use only and may change without notice. Micro Focus assumes no responsibility or liability for any errors or inaccuracies that may appear in this documentation.

© 2017 GWAVA Inc., a Micro Focus company. All rights reserved.

Micro Focus, Retain, the Retain logo, GWAVA, and GroupWise, among others, are trademarks or registered trademarks of Micro Focus or its subsidiaries or affiliated companies in the United Kingdom, United States and other countries. All other marks are the property of their respective owners.

1 Retain 4.2.0.0

1.1 Overview

1.1.1 Intended Audience

This manual is intended for IT administrators in their use of Retain or anyone wanting to learn more about Retain. It includes installation instructions and feature descriptions.

1.1.2 What Retain does

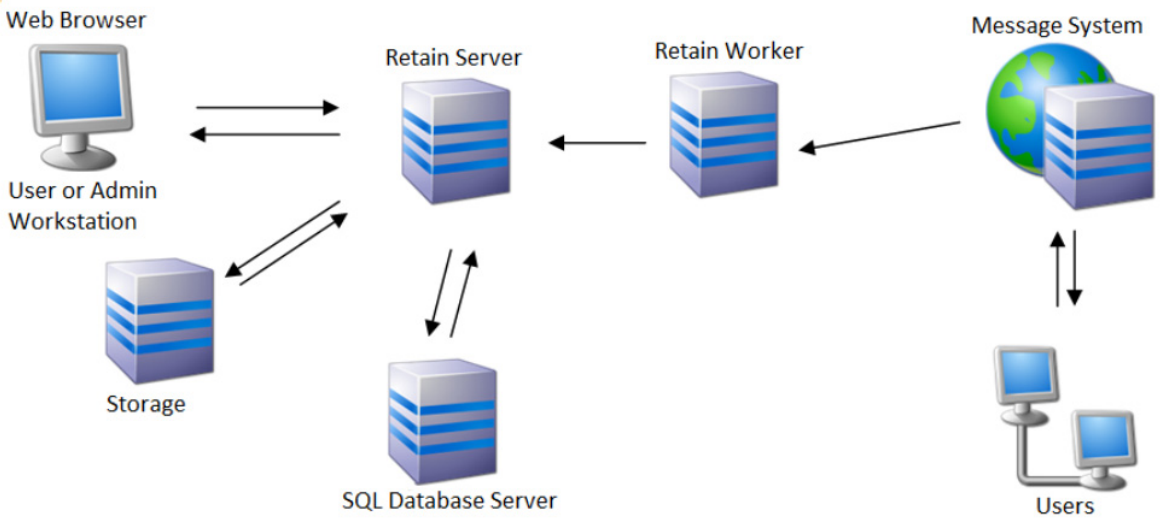
Retain provides a service of long-term storage of data as well as providing search, retrieval, and review services for archived messages. Retain is NOT a backup or emergency restoration system. Retain archives messages and data from messaging systems, phones, and social websites and stores the data for long-term reference. Users may log in and review their personal archived data and search through it. This not only provides legal compliance and litigation protection, but may be used to free up space on messaging systems and enhances the data management.

1.1.3 How Retain works

The Retain Worker process connects to the appropriate message server to collect data, using the mes-sage system's defined APIs, (For example, SOAP for GroupWise and Exchange). This data is transferred to the Retain Server which stores the collected data in a defined storage location and indexes the data in the SQL server. Users log in to the Retain Server's web interface to search through and access their archived messages. User rights are managed by the administrator.

1.1.4 Architecture

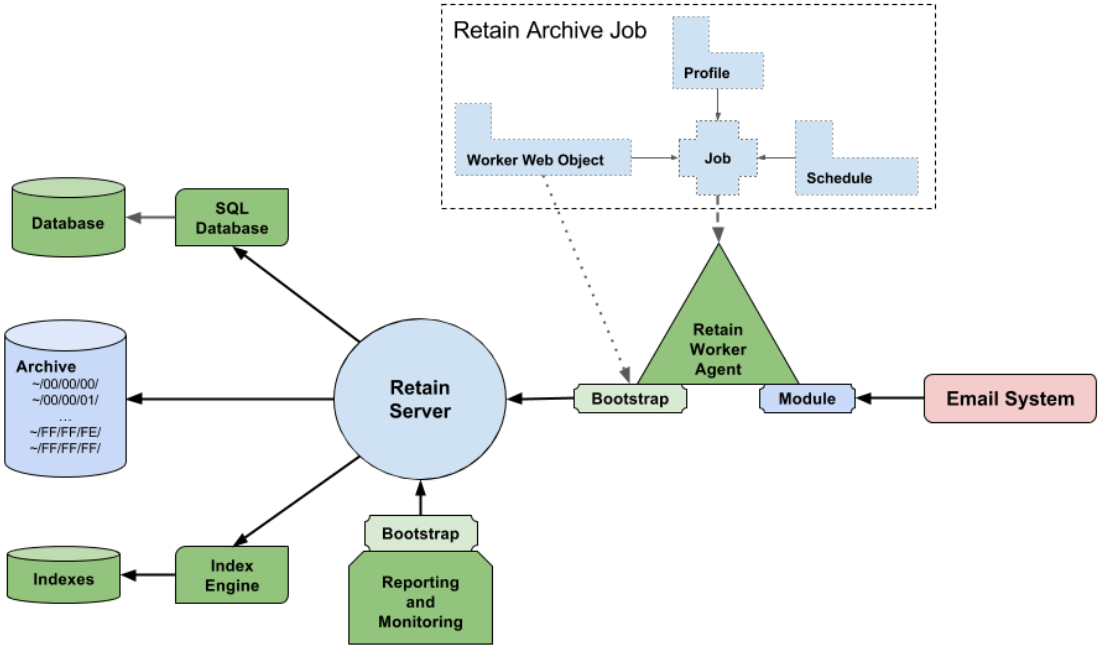
Retain consists of several main parts which can be installed on the same server or they can be spread out across different servers to allow flexibility in where data gets stored and which servers will be used to perform Retain functions.



- **Retain Server:** This is the core piece of Retain. All functions are controlled from the Retain Server. The archive is stored here. The server also manages the Retain Worker and stores data in the database server. There is only one Retain Server per system.
- **Retain Worker:** There is at least one per system, and often one per messaging server. The Retain Worker performs the data collection and transfers the collected data is handed to the server. The server stores the data in a database. This can be installed on the Retain server, the mail server or a standalone server.
- **SQL Database Server:** This is an SQL server where Retain stores the message header data, user data and links to the stored messages in the archive. It is not actually part of Retain. Retain was designed to support many different databases.

NOTE: Installation, maintenance, tuning, and backup of the database is the customer's responsibility. This can be installed on the Retain server or a standalone server.

- **Reporting and Monitoring Server:** This component keeps job and server statistics and handles mailbox error monitoring. This can be installed on the Retain server or on a standalone server.
- **Indexing Engine:** This component keeps indexes all the data. The standard High Performance Indexer is installed on the Retain server or on a separate server cluster as the High Availability Indexer.
- **Stubbing Server:** The stubbing server works to remove large messages from the GroupWise system and creates a 'stub', or link, to the message which is stored in the SQL database. See the Stubbing Server section in the Main Guide to decide if a stubbing server is correct for your system. From a User's point of view, there is no change to the behavior of their GroupWise mailbox. Currently stubbing is only supported for GroupWise 8.0.1 or later. This must be installed on the Retain Server itself.
- **Retain Router:** The Retain Router gathers messages data from Android and BlackBerry phones using REST, and is installed and controlled by the local administrator. Phone data is sent to the Retain Router or Server where the device ID has been registered. The data is stored locally until the Retain Router forwards it on to the Retain Server. Afterwards, the data is accessed the same as all other message data in the Retain Server. This is installed in the network DMZ.



1.1.5 Best Practices for Component Placement

Retain components communicate via TCP/IP. Though placing all the components in the same server would yield the best communication speed, such placement is impractical for larger systems. There must be balance the performance of the components on the physical servers with the speed of the network links joining them. Retain Workers may exist on the same physical machine as a messaging server.

1.1.6 How Retain Stores the Archives

Retain uses a hybrid data storage approach. The database contains all the meta data, folder structure and attachment information but does not hold the actual message text or actual attachments. These are stored on the file system in a single instance storage scheme using a hash for each individual message. Data is hashed using the SHA-256 algorithm which can be used to detect tam-pering.

1.1.7 Other components that Retain depends on

The following items must be ready before you install Retain's core components.

- Supported messaging system. (For social media capture the RSM gateway must also be installed)
- Apache Web Server.
- Apache Modules mod proxy and mod_rewrite. (Installed and enabled)
- SQL Database for storage.
- Java JDK. (This is installed automatically by the installer for Retain use only.)

1.1.8 Design Considerations

Retain is designed to be as flexible as possible, giving you choices as to where to install its components. Here are some points to keep in mind when deciding where to put everything.

1.1.9 SQL Database Server

Where should the SQL database server be placed in the network? The faster the network connection the better. Local installation gives the best communication speed, but it's usually unrealistic to do so. In a large system, you might have the database on a server by itself for performance or security reasons. Then, network speed and reliability become key considerations.

- Network link between the Retain Server and SQL Database Server must be speedy and reliable.
- The Retain database may have to be manually created by the administrator and a user account must be assigned with full rights. See the Database section.
- Storage requirements: Roughly equivalent to the cumulative size of the message data store. See the 'Estimating Storage Requirements' section.

NOTE: Installation, maintenance, tuning, and backup of the database is the customer's responsibility. Tuning an SQL Database Server can result in significant performance gains.

1.1.10 Retain Server

The Retain Server is the heart of the Retain system. All archive processes, search queries, user activities, and auditing is funneled through the Server. When planning the Server, consider the following:

- CPU requirements are high. The bigger and faster the better.
- Storage requirements: Storage sizes may change over time. An expandable storage scheme ensures options down the road.
- Other web applications such as GroupWise WebAccess or iManager should not be installed on the same server. The Retain Server should be a dedicated machine.
- Do not install Retain components on the same machine as iFolder.

1.1.11 Retain Worker

The Retain Worker is the piece that receives data on a scheduled basis from the messaging systems. It then passes this data to Retain Server. Things to consider when placing a Retain Worker are:

- A reliable, speedy network connection between Retain Server and Retain Worker.
- A reliable, speedy network connection to CAS and Mailbox Servers or Post Office Agents being accessed.
- If desired, one Retain Worker can be placed on the same box as the Server for communication performance reasons.

1.1.12 Retain Router

The Retain Router needs to be visible and accessible to the Internet to receive mobile data.

- The Retain router should be installed in the DMZ.

- A dependable network connection to both the INTERNET and the Retain Server are required.
- The Retain Router must have an active and constant connection to a data holding location, called the ‘data path’. The data path is simply the directory where phone data is kept while waiting to be sent to the Retain Server.
- Because data path storage requirements are low, it is highly recommended that the router and data path all be located on the same machine.

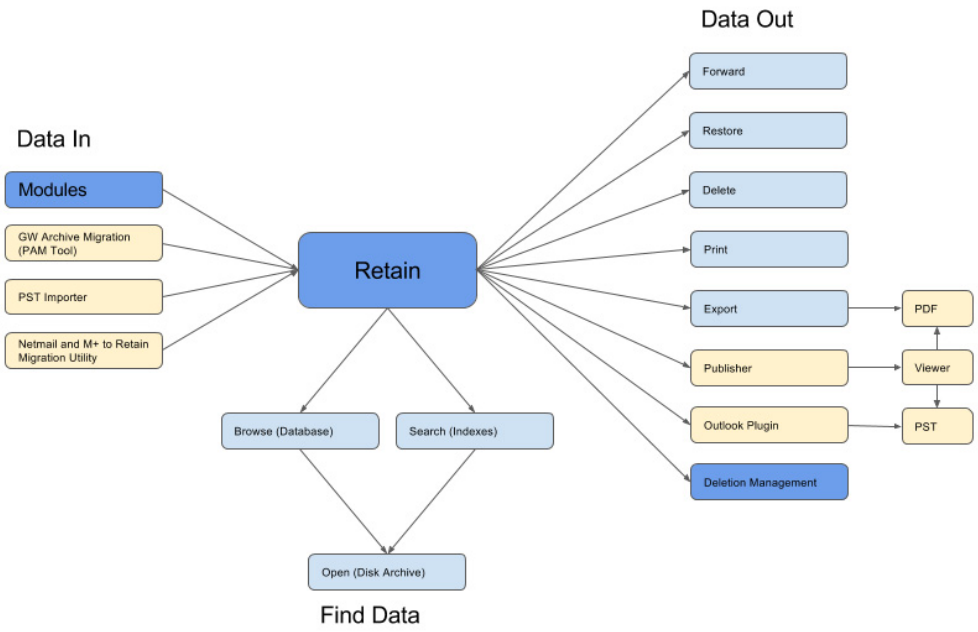
1.1.13 Messaging System Address Book

Retain gathers mail from known users. Users known to the messaging system are stored in the System Address Book. Retain caches this information locally. The address book needs to be updated as users are added. The Retain never deletes a user from the cached address book unless there is no mail archived for the user. Over time, Retain will know about all users in the messaging system, both current and past. Because Retain stores past users in its cached address book, it can distinguish between two users of the same name. For example, “John Smith” added today will be recognized as a different user from “John Smith” who worked at the company six months ago.

1.1.14 Retain Functions

Retain has three fundamental functions:

- Import data into Retain
- Find data in Retain
- Export data from Retain



Importing Data into Retain

After getting the Retain server up and running the next thing to do is to import data into Retain.

The Modules allow data to be dredged from live messaging systems.

There are also workstation tools for migrating offline data into Retain. These can be found under the Tools menu at the top-right of the Retain web console. These include the GroupWise Archive Migration Tool, the PST Importer and the Netmail and M+ to Retain Migration Utility.

Finding Data in Retain

- Once data is in Retain you may need to find the data that has been archived. When you are browsing a mailbox, you are seeing the metadata of each message stored in the database. When you search in Retain, you are viewing the indexes that Retain has built to find things quickly. When you open a particular message, you are viewing the data as stored on disk.

Exporting Data from Retain

Within the web interface, you can use the checkboxes next to each message to select additional messages. Now you can select actions to take on the message or messages. A default user has rights to View and Save attachments, as well as Forward and Print Messages. As admin, you can also Restore messages (if you are on Exchange), Export to PDF, Delete messages. Other actions that can be taken are: set Litigation Hold, Confidential State (if the mailbox has multiple users with access) and Tags. Users and Groups can be given additional rights as needed.

If you are using Outlook, the Outlook plugin allows the user to search and download messages from Outlook itself. Retain becomes another folder in Outlook.

The above methods are good for dealing with small amounts of messages, but for dealing with large amounts of messages you will want to use Publisher, a workstation tool found under the Tools menu. You can save a search query in the Retain web console and then use it again when you log into Publisher or you can create the search query in Publisher. Publisher will download messages to the workstation, once the download is complete you use Viewer to read, search, redact and export the messages to PDF or PST.

Finally, when the message has reached the end of the data retention policy you would use Deletion Management to permanently remove messages from the archive.

1.2 Retain Services

Retain utilizes configured modules to archive different messaging systems and to allow tight configura-tion of desired archive activity. From Email Systems, Retain archives all desired data, including email, notes, appointments, meetings, reminders, and tasks. From Social Platforms, Retain archives all postings and comments made to monitored pages and all chats. For Smart Phones, Retain can archive PIN, SMS, and phone call data. For Search Engines, Retain can archive search criteria used. All data streams will be organized by the user who created or used them, they time they were created, and the data source. All this information is indexed and can be quickly managed through search, export, and publishing of desired items from the archive.

The systems and data streams which Retain can archive are listed below:

1.2.1 Email Systems

- Exchange
- Office 365
- GroupWise
- Gmail

- Bloomberg
- GBS Notes

1.2.2 Smart Phone Communications

- BlackBerry (BES Server)
- BBM Protected (on all platforms)
- Android
- IOS and Android (via CellTrust Secureline)

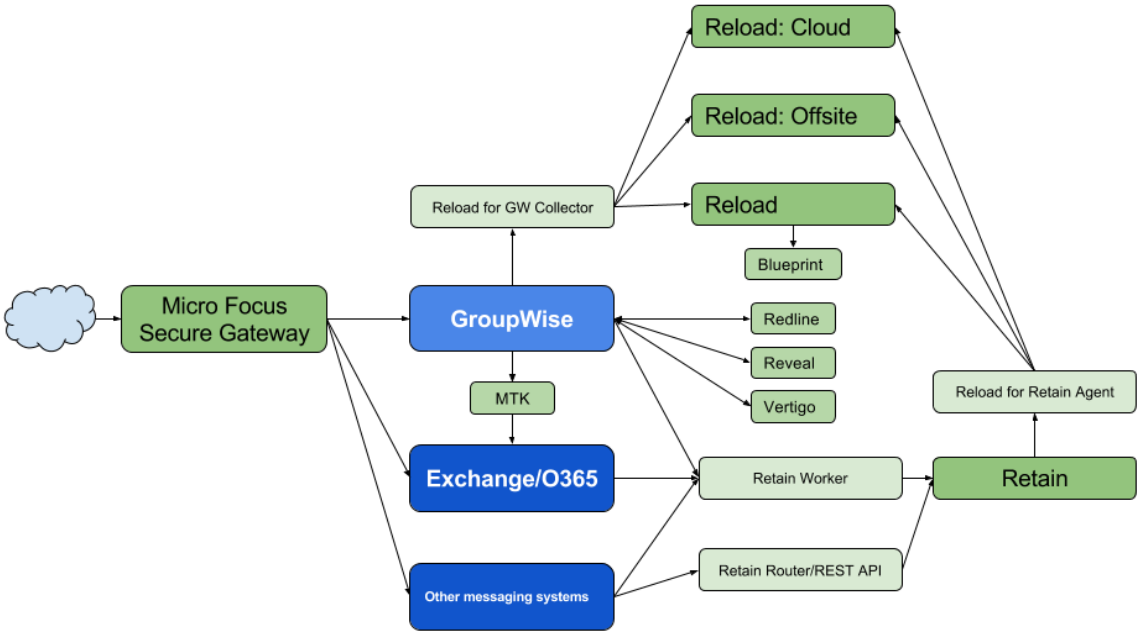
1.2.3 Social Platforms

- Facebook
- Flicker
- Google+
- LinkedIn
- Instagram
- Pinterest
- Twitter
- Vimeo
- YouTube
- Skype for Business with O365

1.2.4 Micro Focus GWAVA Family of Message Handling Products

- *Micro Focus Secure Gateway* is a message scanning product that protects your system from malware and spam.
- *Retain* is an archive storage product that is designed to keep messages from GroupWise, Exchange/O365, GMail, BlackBerry, Bloomberg, Notes, mobile, social and other messaging platforms for the long term to meet data retention legal requirements and has powerful search capabilities for eDiscovery.
- *GroupWise Disaster Recovery powered by Reload for GroupWise* is a hot-backup and disaster recovery product for GroupWise. It keeps a few weeks of data and can easily restore messages, calendar items, address books, and even whole users. It can also act as a fully functional Post Office in times when the GroupWise POA is down.
- *Reload for Retain* is a hot-backup and disaster recovery product for Retain. The Retain archive keeps items for the long term and often becomes the only repository of an item. Without a backup you can leave yourself open to legal liability if the items are lost before the retention period has passed. Reload for Retain provides high speed backups and can temporarily take over duties of the Retain server should an incident occur.
- *Blueprint for Reload* extracts important business intelligence data from your GroupWise message store by performing in-depth analysis on your Reload backups.
- *Migration Tool Kit* for migrating users and data from GroupWise to Exchange/O365.
- *GroupWise Reporting & Monitoring powered by Redline* is a comprehensive, customizable, monitoring and reporting tool for GroupWise.

- *GroupWise Forensics powered by Reveal* provides essential auditing and oversight capabilities that legal, human resources, and auditing personnel need within GroupWise.
- *GroupWise Mailbox Management powered by Vertigo* is the Enterprise Mailbox Management tool for GroupWise.



1.2.5 Technical Support

If you have a technical support question, please consult the Micro Focus Technical Support at http://support.gwava.com (http://support.gwava.com)

1.2.6 Sales

Micro Focus contact information and office locations: www.microfocus.com (http://www.microfocus.com)

To contact a Micro Focus sales team member, please e-mail info@gwava.com (mailto:info@gwava.com) or call 866-GO-GWAVA ((866) 464-9282), or +1 (514) 639-4850 in North America.

1.2.7 Copyright Notice

The content of this manual is for informational use only and may change without notice. Micro Focus assumes no responsibility or liability for any errors or inaccuracies that may appear in this documentation.

© 2017 GWAVA Inc., a Micro Focus company. All rights reserved.

Micro Focus, Retain, the Retain logo, GWAVA, and GroupWise, among others, are trademarks or registered trademarks of Micro Focus or its subsidiaries or affiliated companies in the United Kingdom, United States and other countries. All other marks are the property of their respective owners.

2 Upgrading to Retain 4

2.1 Upgrading to Retain 4

To upgrade to Retain 4 requires that Retain be upgraded to 3.5+ first. Please see the Section 3.1, “Retain Planning and Design Best Practices,” on page 37 page for system design recommendations.

Retain 4 requires new licenses to function after the upgrade. An upgraded Retain 4 system will not func-tion until the licenses are applied.

For Linux:

Retain 2.x to 3.x

Retain 3.x to 3.5 or 3.5.1

For Windows:

Retain 2.x to 3.x

Retain 3.x to 3.5.1

Once Retain is updated to 3.5+, the upgrade process to Retain 4 may be started.

2.2 Upgrading from Retain 3.5.1 or 4.x to 4.y

2.2.1 Upgrade changes

Retain 4 contains significant improvements over previous versions. The User Interface in Retain 4 is changed and a completely new search has been implemented. The new search contains a new, high pow-ered indexer and back end which allows searching in multiple new ways on each item to maximize rele-vant and intuitive results. Due to the increased ability, some system requirements have been increased.

2.2.2 Requirements

- Storage Space

  While the migration is performed, a new index will be created. The current index storage space will be temporarily increased while the two indexes exist.

  There is a check for at least 4 times the available free space compared to the existing 3.5.1 index. A warning will appear if there is less disk space available, this is only a warning, proceed at your own risk.

  Once migration is complete, the old index may be removed and space reclaimed.

  See Estimating Storage Requirements.
- Time and Access

The upgrade to Retain 4 will take some time, how much depends on how large the Retain archive is. The upgrade includes an update to the configuration database, which is completed at first start-up and should be completed in less than ten minutes, and a migration to a new index. While the index migration may take significant time, the existing index and system will still be available and accessible during the entire process.

- System Requirements Changes

  The system Requirements have been increased for Retain 4 to 12 GB RAM on the minimum side. Check the requirements in the Technical Support section to verify changes required for the sys-tem.

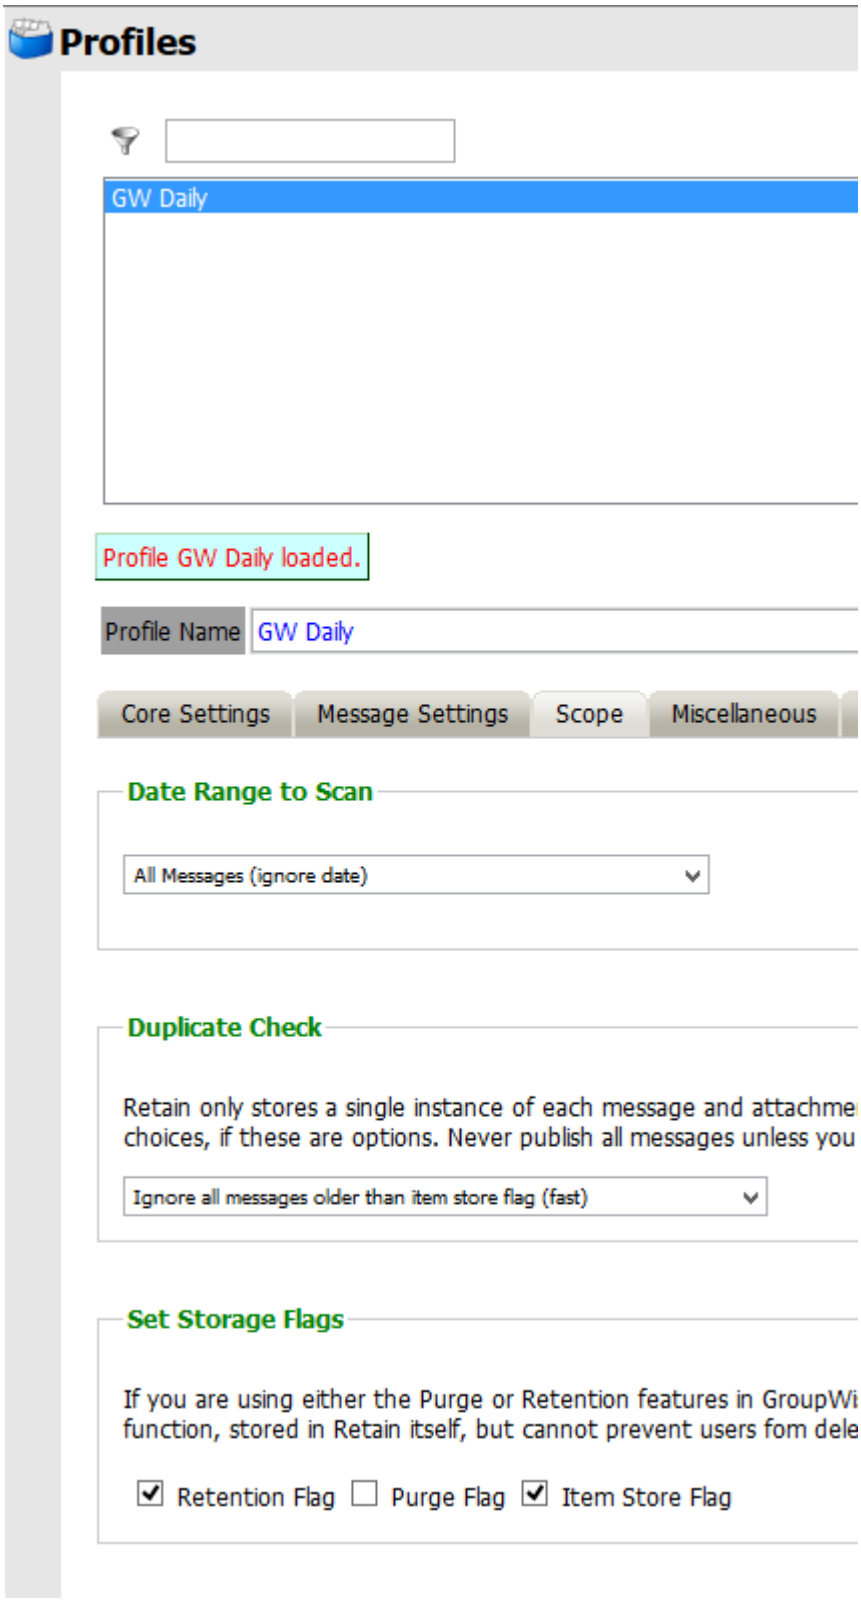
- Licenses

  Retain 4 will require new licenses to function after the upgrade. Make sure to have the required licenses on hand for the install. An upgraded Retain 4 system will not function until the licenses are applied.

  Make sure to run the install on every component in the installed Retain system to make sure all components are upgraded together. Any Workers and Servers not hosted on the same machine must each have the installer run to be upgraded.

2.2.3 GroupWise Module

If you are archiving with the GroupWise Module, log into Retain and determine which storage flag Retain is currently using. This is the flag you want to choose during the upgrade process. Retain will query GroupWise and update the Item Store Flag in the Retain database for each user.



If you should choose incorrectly, it can be changed in the module. If not set correctly, it will mean that the next job will, at most, do a full dredge.

Please note that after the upgrade to 4.2 there will be a new default Scope setting in Profile(s).

The option "New items" will become the default selection if previously the Date Range to Scan was set to "All Messages (ignore date)" and Duplicate Check was set to anything other than "Try to publish all messages (SLOW)."

Please note that in 4.2 upgrade the controls for the GroupWise Storage Flags have been moved from the Profile to the Module Configuration page.

See the GroupWise module page for more details.

2.2.4 Upgrade Process

If upgrading from Retain 4.0.x the upgrade process should take between 15 and 30 minutes, depending on the size of the system, very large systems may take longer.

1. Download the latest version of Retain.

   On Linux, you may use the command line "wget http://download.gwava.com/download.php?product=Retain&version=current"

2. Extract the archive. This will extract into a new folder.

   On Linux, in the GUI right-click and select "Extract Here" or on the command line "unzip Retain<version>.zip"

   On Windows, right-click and select "Extract All..."

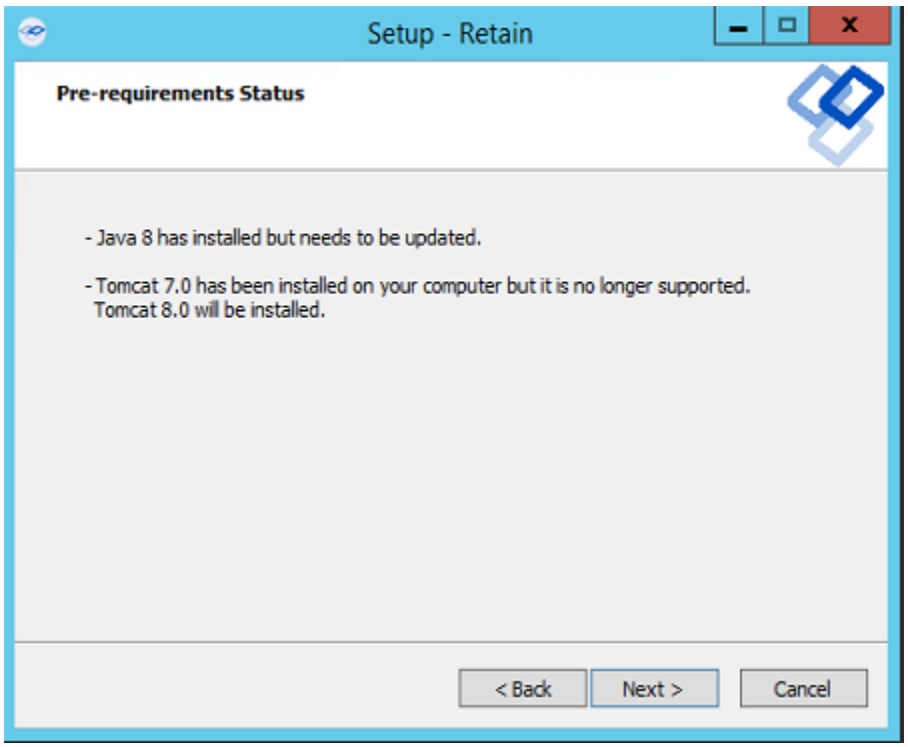
3. If on Linux, open a terminal window in the Retain install folder and make the bash script file executable by running the command: "chmod +x *.sh"

4. Run the installer:

   On Linux: "./RetainInstall.sh"

   On Windows: "RetainInstall.exe"

5. The installer will check the tomcat version that is installed and will upgrade it to Tomcat 8.0



6. Retain will detect if a previous version of Retain exists on the server. It will report what components have been installed and will ask if you wish to upgrade the current installation or if you want to add new components as well.

```
Retain Installer 4.1.0.0


Retain has already been installed on this system with the following components:

Retain Server
Retain Reporting & Monitoring Server
Retain Worker

Choose one of these options:
1) Upgrade current installation of Retain
2) Upgrade current installation of Retain and add new components
3) Exit this installer

Choose (1-3): 
```

7. If you choose to install additional components it will ask what you wish to install.
8. Then the installer will ask if you want to Upgrade Retain, preserving the existing configuration or Overwrite Retain which will replace the configuration file but will not erase data. This is useful in certain troubleshooting circumstances.

```
================================================================================

                        Apache/Tomcat Configuration

================================================================================

Apache 2 appears to be installed
Apache Base is /etc/apache2
Apache Service is /usr/lib/systemd/system/apache2.service
Apache Conf is /etc/apache2/httpd.conf
Apache Version is 2.4
Apache Group is www
Will assume that mod_proxy, mod_rewrite, and mod_proxy_ajp
are installed and enabled. These are standard in Apache 2.2+.
You may check in the Http Server control Panel in Yast

/opt/beginfinite/retain exists. Choose one:
1) Upgrade Retain, preserving configuration (recommended)
2) Overwrite Retain, losing configuration
3) Quit

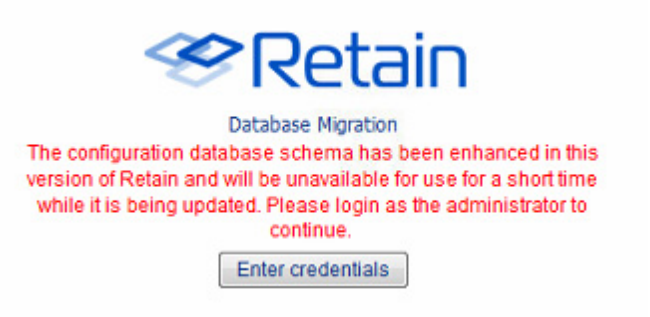
Choose (1-3): 
```

9. After choosing to upgrade or overwrite Retain, the installation will begin. Wait until the installation is finished.

2.2.5 Post-Install Tasks

1. Browse to Retain Web Console.

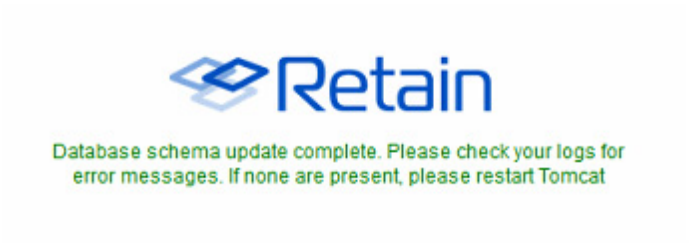
2. The upgrade may require a configuration database update, depending on the version of Retain you upgraded from. Select the ‘Enter Credentials’ button and enter the administrator account username and password. Select the button to start the migration.



   a. During the database Migration, do not interrupt the process. Wait until the process completes.



   b. Once the migration is complete, the system will instruct you to restart tomcat.



   c. Once the update has completed, restart tomcat to initiate the new database.
   d. The tomcat restart will take a little bit longer than a normal restart, as the database needs to be initial-ized and set up the new interface.

e. After a few minutes, browse again to the Retain Management interface.



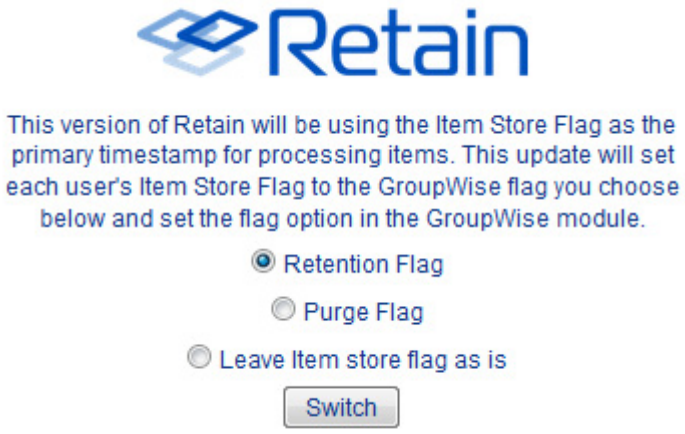
3. If you have the GroupWise module configured how the GroupWise timestamps are used by Retain have changed.
   a. You will receive a warning page describing what will happen.



   b. You will need to enter the admin credentials.



   c. Choose which timestamp flag Retain will use in the Module.

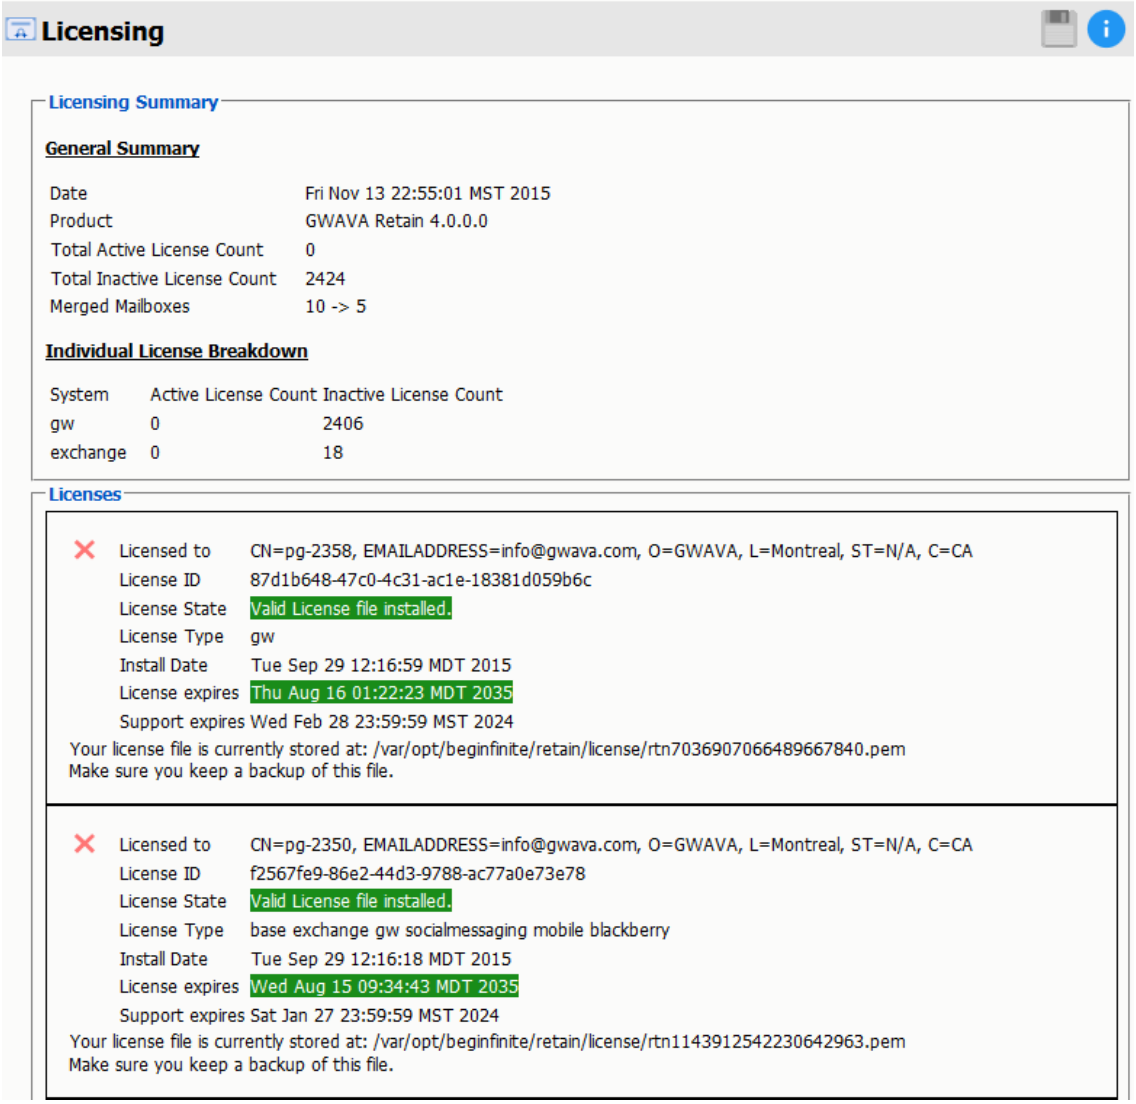


d. Retain will update, which will take some time. The Retain interface will not be available until this process completes. In Retain 4.2, this process is done in the background so the interface is available immediately, however jobs will be disabled until the process is complete.

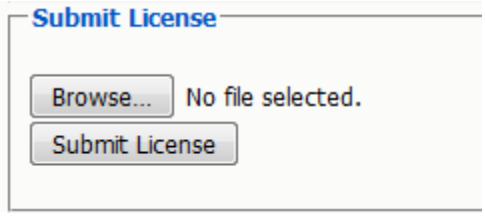


e. 3e. When the migration is complete you will need to restart tomcat. In Retain 4.2, no restart of tomcat is needed, jobs will begin again on schedule once the process is complete.

4. Login to the Retain Management interface. If upgrading from a version of Retain previous to Retain 4.x, new licenses for Retain 4 must be immediately applied. These licenses should be provided by the sales representative. (Removing old licenses is not required, but may help to keep interface looking clean.)



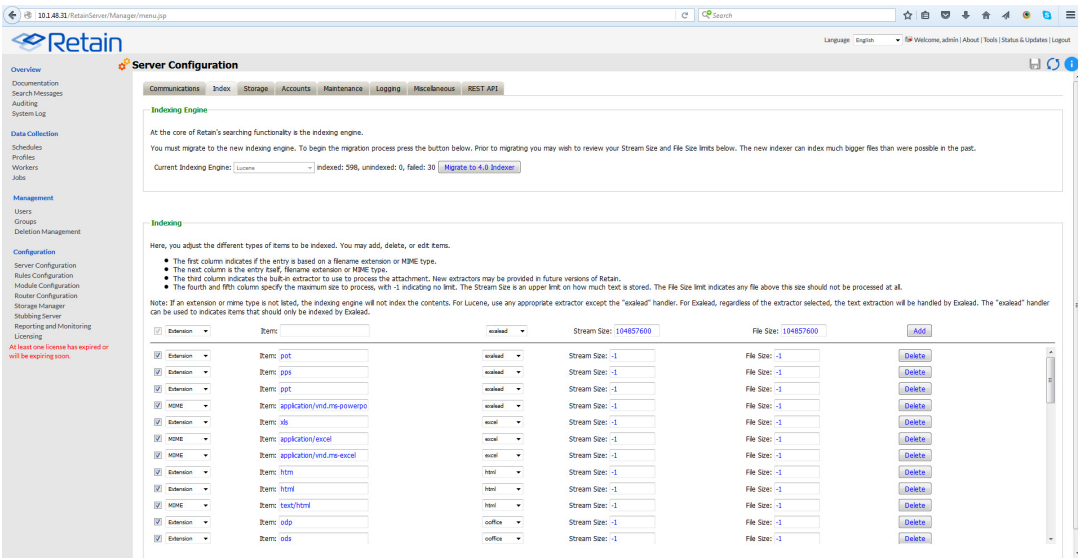
a. To install the new licenses, select the ‘Browse’ button from the ‘Submit License’ section at the bottom of the screen. Select the desired license and when back in the Retain interface, select ‘Submit License’ to upload the new license to the system. Repeat the process for each license required for the active and desired modules in the system.



b. It is a good practice to remove expired licenses from the Retain system to keep the interface clean and to easily see any licensing issue. However, removing expired licenses is not required.

5. If upgrading from Retain 3.5.1.1:

a. Immediately go to the Server Configuration page and select the Index tab.

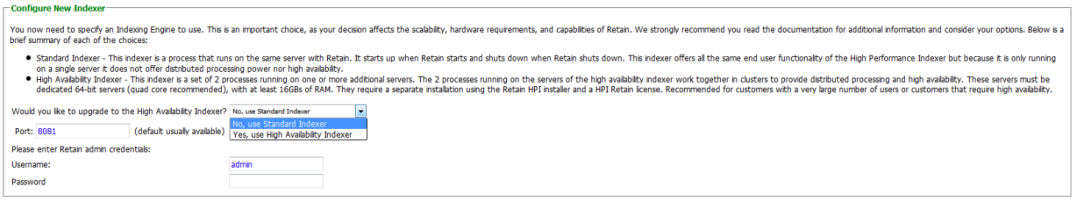


b. Select the Migrate to 4.0 indexer button.

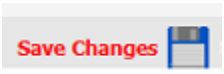
c. There are two index options: High Availability Indexer and the standard index. Both options utilize the same high performance index engine and will both provide new searching functionality; however, the 'standard' engine is embedded while the 'High Availability Indexer' is a scalable, external, cluster system for large systems. The default is set to the standard index. The standard index engine is very powerful and sufficient for everything up to extremely large or busy systems.

d. If you wish to use the High Availability Index engine, a separate license is required. Please contact support for guidance, system planning, and installation.

e. Once the migrate button has been pressed, the new indexer section is displayed.



f. The Standard engine only requires the Retain admin credentials. Enter the administrator credentials then save the changes by selecting the save-changes button to con-tinue.



g. The index migration begins automatically. During the migration of the index, the active index is displayed at the top of the page. The active index displayed at the top of the page will still be available and fully functional. New search options and features will not be available until the new index becomes active.

h. The index page is updated every 10 seconds to display the progress. With no archive jobs, migrations of the index will average 60 messages per second on a Linux server, and about 45 messages per second on Windows Server. Migration speed is limited by disk IO, (and network speed with external indexes). Large messages with large attachments will take longer. If archival jobs are running, the speed of the migration will depend wholly on the size and activity of the archival job.

i. When the index page updates showing the new index is running, the new index has automatically become active.

j. Once the index migration is complete, the indexer will need to optimize the index files. This is done during the nightly maintenance. The larger or more segmented the indexes the longer this will take. This may temporarily triple the disk space the indexes use as the files are consolidated. It is a recommended to go to the *Server Configuration* | *Maintenance Enable Index Optimization* tab and set the *Enable Index Optimization* to Saturday weekly. The first optimization is the longest and may take many hours, after that it usually only takes minutes and can be done nightly.

6. The upgrade is complete.

2.2.6 Tool Compatibility with Retain 3.x (Lucene) Indexes

Not all tools are compatible with the Retain 3.x (Lucene) indexes.

Tool	Compatibility
Retain web console	Compatible
Publisher	Not compatible
Outlook Plugin	Not compatible
Retain App for Android Phones	Not compatible
Retain App for iOS Phones	Not compatible
Retain Single Sign-on plugin for GroupWise Client	Does not use indexes
Retain Single Sign-on plugin for GroupWise WebAccess	Does not use indexes
GroupWise Archive Migrator/PAM Tool	Does not use indexes
PST Importer	Does not use indexes
Netmail and M+ to Retain Migration Utility	Does not use indexes
PowerShell Sync Script	Does not use indexes

3 Install

3.1 Retain Planning and Design Best Practices

There are so many variables that impact performance and design that no single document could adequately cover them. Instead, this document mostly discusses concepts so that the reader can apply them their needs. This document has four major sections:

- Retain Architecture
- Hardware Design
- Retain Configuration
- Backing Up Retain

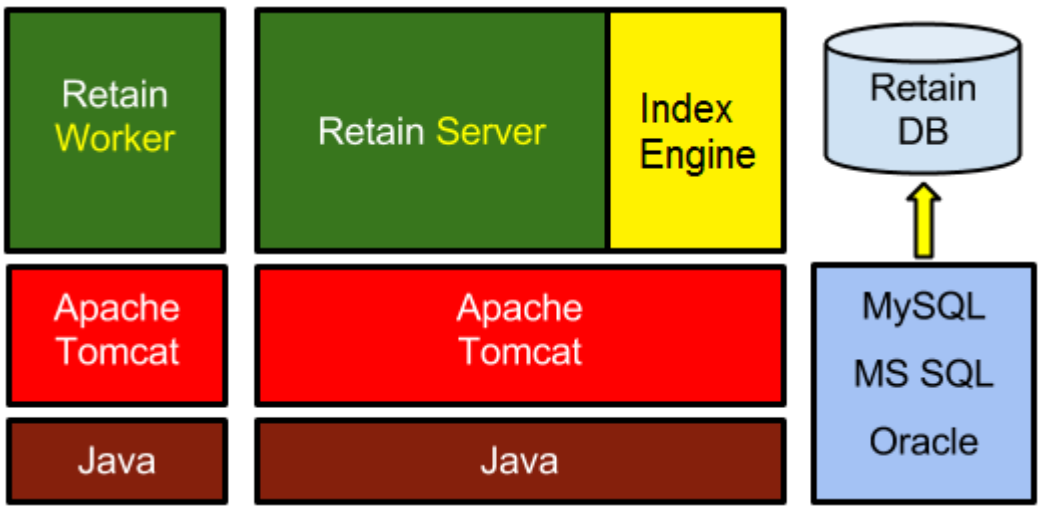
This document's intent is to provide insight into how Retain works so you know how to design your system. It will provide some guidelines and best practices; however, it does not attempt to make hardware recommendations or provide specifications of such. It is anticipated that the reader will understand how Retain works, how the various hardware components play a role, and be able to make hardware decisions based on that understanding. Again, no two systems are alike, so one size does not fit all. And, keep in mind, Retain's minimum system requirements are MINIMUMS. Retain will work with those minimums; but, if performance is a priority, then read on and apply these concepts to your specific situation.

3.1.1 Retain Architecture

Retain can run on a bare metal server or on a VM running Windows Server or SuSE Linux. For backup purposes and flexibility, we recommend running it on a VM.

Retain consists of four major components:

- Server
- Worker
- Indexer
- Database



Server

The Server is where the archive system is configured and maintained. It is responsible for storing, indexing, searching, and reading archived items.

Worker

The Worker is the component that interfaces with the messaging host/mail server containing the messages you are archiving. It retrieves the messages and passes them onto the Retain Server.

Indexer

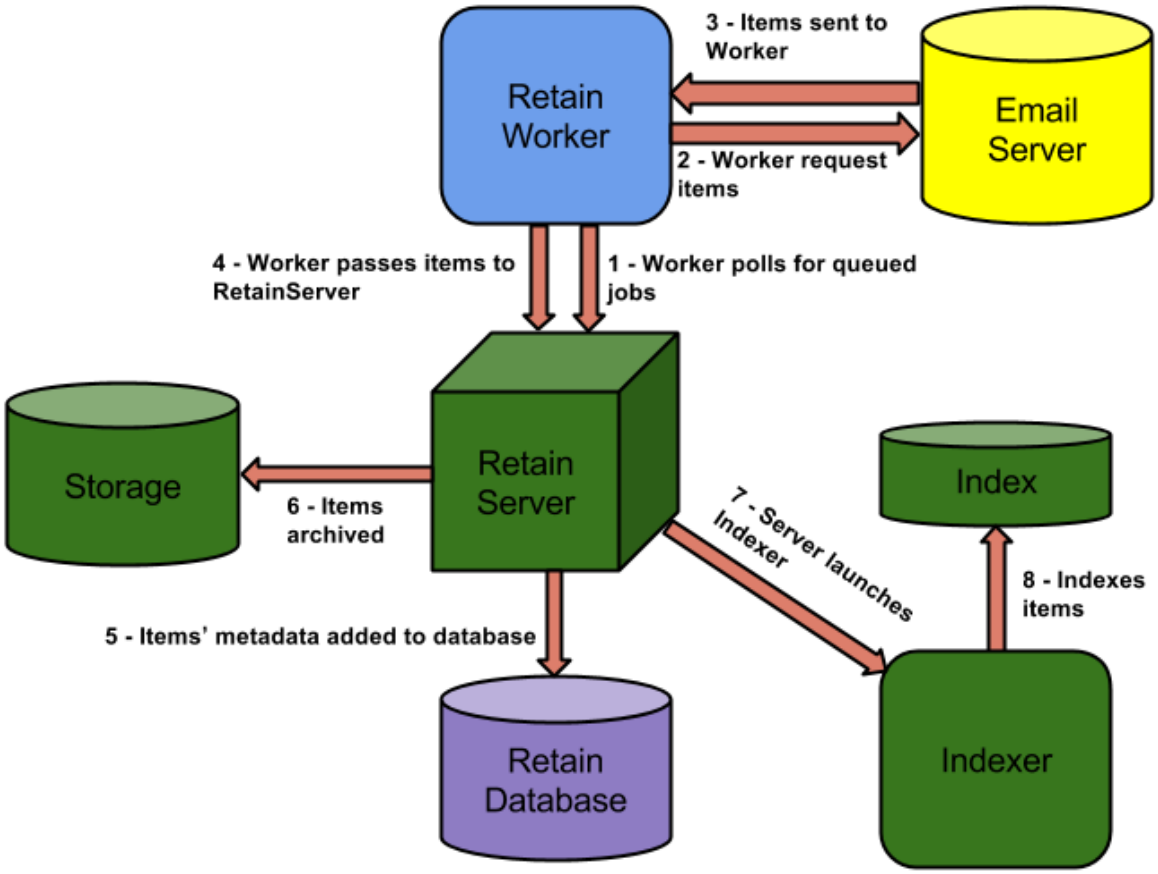
The indexer scans the messages and attachments to make them searchable, indexing each word and some phrases. When a user performs a search in his/her Retain mailbox, the list of messages returned is coming from the indexer, not the database. Understanding that difference helps you make decisions on memory configuration choices for Tomcat (which powers the indexer) and for the database. The index process is the most memory intensive (I.E. uses the most memory) part of Retain.

Database

The database stores most of the Retain configuration as well as the message metadata, which is information about the messages being store (subject, sender, recipients, links to attachments, indexed state of messages, folder location of the message, etc). When a user logs in to his/her Retain mailbox, the list of folders and their messages are being retrieved from the database.

3.1.2 Archive Job Flow

When an archive dredge job is running the work flow follows this pattern.



1. The Worker polls the Server every 10 minutes (default) for new queued jobs it may need to run. It launches the job.
2. The Worker connects to the mail system and does a login to each mailbox. For every mailbox, it requests the items in that mailbox based on the settings in the profile.
3. The mail server responds by sending the items to the Worker.
4. The Worker sends smaller items to the Server. For large items, it sends the items' metadata and awaits instructions from the Server as to whether the item already exists.

   If it already exists, the Server notifies the Worker that it does not need to send the item over.

   If it does not exist, the Server notifies the Worker to send the item.
5. The Server updates the Retain database with the item's metadata if it does not already exist.
6. A good example of this is when User A sends message 1 to User B. When User A's mailbox is archived, Message 1 gets stored on disk with a record getting written in the database pointing to the file on disk and associating the message to User A's mailbox. When User B's mailbox gets archived, the Server notices that message 1 is already stored on disk and that a record in the database exists for that message. Now it updates that record and associates it with User B's mailbox.
7. The Server adds the item to the storage area on disk.
8. The Server launches the indexing process (if it is not already running) to begin the indexing process.
9. The Indexer indexes any items that need to be indexed.

3.1.3 Hardware Design

There are four major considerations you need to take into account when designing the hardware for a Retain system:

- Bandwidth
- CPU
- Storage/Disks
- RAM
- VM Configuration

Bandwidth

The Worker will be querying your messaging system for messages to be sent to it and receiving every one of those items; however, not all of those items will be sent over to the Server. If the link between the Worker and the messaging system is slow, you should consider placing the Worker on the messaging system's server or on a server that has a fast link to the messaging system.

The only downside to this strategy is software updates. When upgrading Retain software, you will have to go to each server hosting a Worker and upgrade its software. If the Worker is running on the Retain Server itself, then the installer automatically upgrades every Worker on the Server.

CPU

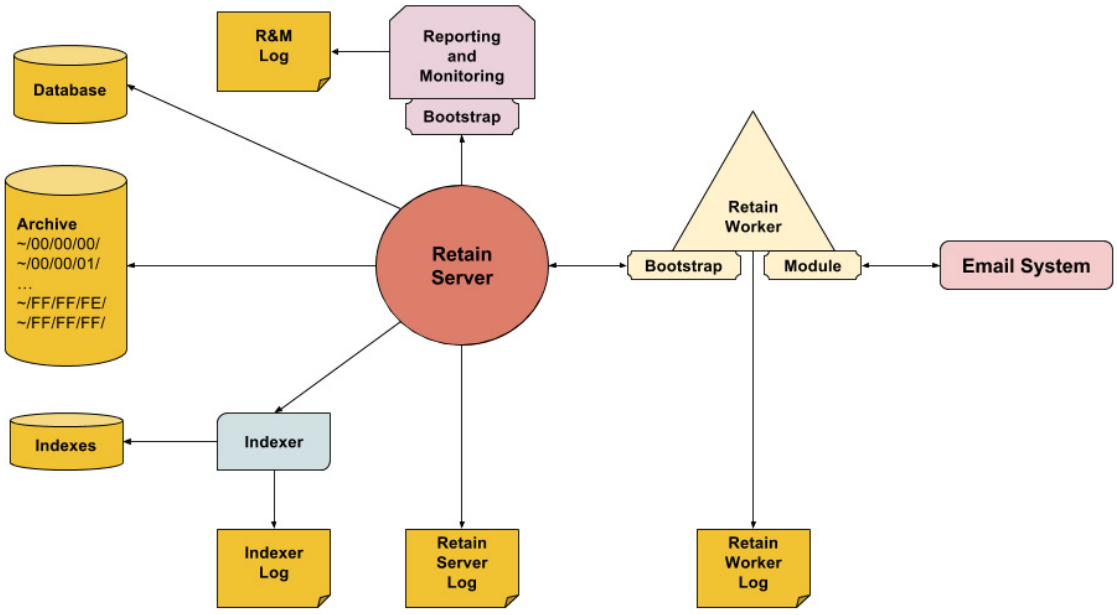
Retain is multithreaded and able to make use of multiple CPU cores. The base server uses 4 threads, and the Indexer starts with 3 threads. If more than 7 CPU cores are available, additional Indexer threads are spawned. The basic formula is [cores - 4 (minimum 3)].

Cores	Retain Server Threads	Indexer Threads
2	1	3
3	2	3
4	3	3
5	4	3
6	4	3
7	4	3
8	4	4
9	4	5
10	4	6

We have found that 8 CPU cores provides a sweet spot for performance gains, allowing Retain 4 threads and the Indexer 4 threads.

Storage/Disks

Since Retain is an archiving solution, it will happily fill up your storage system. While it will warn you of impending doom, it is up to you to keep bad things from happening like the hard drive filling up completely. It is difficult to recover from that since not even the OS works very effectively under those conditions. You need to design in the ability to add extra storage easily. Retain stores data in the archive, database and indexes the data for search. Retain also logs what is happening for auditing and troubleshooting purposes.



Since the point of Retain is data safety, we will design around that priority. It is fine in a lab setting to just leave everything on a single drive, but in production, you should desire something more robust. At the very least put the OS on its own partition. It is much easier to recover if a data drive is full rather than a boot drive. It is also important to consider what needs to be done to backup Retain.

Size

Archive Files (BLOBs)

The trend with disk storage is that it is increasingly becoming less expensive. Micro Focus cannot tell a customer how much storage space will be needed - it is almost impossible to predict with any degree of accuracy. A general rule of thumb is current size + rate of yearly growth. Understanding the trends of your email growth, the current space consumed by it, and how Retain stores those messages can help you in your planning. If anything, you'll want at least enough storage space to handle your initial archive job of all your email to date and up to the next year or two. You can always add more disk space and Retain storage partitions after that.

Retain's storage design is efficient. Only one instance of any message or attachment is stored on disk, regardless of how many users received the same items and regardless of whether those users are on the same mail server.

Retain also allows you to expire and delete messages stored in the system after a specified time period, saving you disk space as well.

According to some research we found on the Internet, the average number of emails per user is currently 120 per day or 32,000 emails per year given an average workweek and increasing over the years. The average size of an email is 75kb. Do some math and the average is 2.4GB per user per year.

Database

There are a variety of factors that determine how much disk space Retain will consume. The biggest factors are number of messages per day and the length of time you intend to archive your data. It may be difficult to figure out either of those factors. If you have a virtualized environment, you can allocate more space than you think you will use and thin provision the disks. For our cloud customers, we typically set the db partition to 500 GB and go from there. If a partition runs low on disk space at any point, support can direct you on the proper steps to move the data to another partition if necessary.

The numbers provided in the following table are simply representations of three different systems. Remember, no two systems are alike. Two customers with the same number of messages in their system may have vastly different database sizes due to the difference in the message metadata. For example, Customer A may have short distribution lists while Customer B has a lot of emails with hundreds if not thousands of recipients associated with the messages; therefore, this sample data may or may not represent what you may experience with your system. The purpose of providing sample data is to get you into a possible ballpark range. If anything, err on the side of going higher on the storage need.

Example Systems	Server A	Server B	Server C
Message Count	104,976,966	18,261,383	2,699,654
Archive Size	5.5 TB	1 TB	115 GB
File Size per Message (archive)	4.54 KB	4.71 KB	6.21 KB
Database Size	455 GB	82 GB	16 GB
File Size per Message (database)	56.27 KB	64.02 KB	45.06 KB

File System

When selecting the type of file system on Linux, we recommend going with XFS because it dynamically creates iNodes whereas ext3 forces you to configure those up front. Once you run out of iNodes, nothing can be written to that disk, even if you have plenty of disk space left. XFS is also popularly regarding as performing faster as well.

It is NOT recommended to use ReiserFS, performance on that file system is much worse than XFS.

Other Disk Partitions

Let's talk concepts first, so you can determine if the following recommendation makes sense for you.

Retain archive jobs are disk I/O intensive. You have message content being stored in the archive directory, each message is getting indexed, you have the database getting updated with message metadata, and you have logs being written to continuously. That is a lot of disk access.

Now, let's apply that information to various disk configurations:

Physical ("bare metal") Server

This is your old-school standalone server not running as a VM. In this situation, Retain is installed on a physical machine, not a virtual one. The server has its own locally attached disk. If there is just one disk, you have a lot of disk I/O contention during a job and that will negatively impact performance. In such a case, the recommendations we will make in this section would apply.

VM Guest on Host With Local Disks

Your Retain server is a VM guest running on a physical server that has only local disks. Nothing is on a NAS or SAN. In such a case, the recommendations we will make in this section would apply. You would simply create multiple disks and just ensure that each one gets created on a different datastore.

NAS or SAN

This could be physical box where the storage is mounted/mapped to a NAS or SAN; or, this could be a VM guest where:

- The VM guest itself is stored on a NAS/SAN; thus, the VM guest's "local disks" are also sitting on a NAS/SAN; or,
- The VM guest itself is stored on the hosts local disks but the "local disks" of the VM guest are on datastores residing on a NAS/SAN; or,
- The VM guest is mounting volumes stored on a NAS/SAN.

If the Retain storage will ultimately be on a NAS/SAN and if the volumes are expandable on the fly, then the recommendations we make don't necessarily apply. As you can see, there are so many configurations that a discussion on what exactly to do becomes impossible. The reader really needs to understand what Retain is trying to do and then see what can be done on the hardware end to facilitate best performance.

If it is a NAS/SAN, that doesn't necessarily translate into great disk performance. Many have a NAS but call it a "SAN". So then you have to consider the pipe speed to the storage: 1 gigabit/sec is considered very slow. On top of that, you have to consider how many disks are in the array, their RAID configuration, and the speed of the disks themselves.

For those situation where it makes sense to follow our recommendations:

There is a logical reason for separating your archive files from everything else and a performance reason.

If all the Retain storage is located on the same volume and you run out of space, Retain provides the ability to create additional storage volumes for the archive files. After an additional logical storage volume is created within Retain, all archive files from that point forward as archive jobs run will go to the new location. That works great.

However, the indexes will continue to grow and Retain doesn't have the ability to partition indexes. Some customers have run out of disk space, created new logical storage partitions that point to another volume, but then run into problems with their archive jobs because they are still out of disk space for the INDEXES. Thus, for logical reasons, you want to have your archive files on a separate volume to begin with unless the volume containing the Retain storage is expandable on demand.

If it makes sense to do so (based on all the concepts previously discussed), you'll want to separate your archive files from your indexes and from your database, which means two to three other partitions on your Retain Server in addition to your OS partition. If your database is on a separate server from Retain, then only two other partitions are needed; otherwise, you'll want three additional partitions.

Data Partitioning

We recommend dividing up your storage directories onto separate disks, so beyond the OS disk there should be:

- Disk 1: Archive
- Disk 2: Index (250G start). For best search performance, consider making this a solid state drive.
- Disk 3: Logs, xml, ebdb, export, backup, and license (150 - 200G)
- Disk 4: Database (if on-board)

Disk 2 should be expandable and you'll want to give it room for the indexes to grow; but, if you cannot do that, then when it runs out of space, you'll simply need to move your index files to another volume with more disk space in the future. For disk 2 - as mentioned previously - you may want to consider an SSD, as that would increase the search performance.

If disk 1 and disk 2 can literally be on different "spindles", then you will also get some performance gains from that because an archive job writes simultaneously to the archive directory, the index directory, and to the database. If each of those are on different physical disks (a.k.a., "spindles"), then this eliminates disk contention bottlenecks. Smaller systems may not need to be concerned with performance while larger systems that have archive jobs running for hours may want the performance gains.

Using disk 3 for logs is especially helpful for larger systems. If you have 6 Workers averaging 5 - 10 messages per second, expect a RetainServer log of around 60G unzipped. Plan for 150 - 200G for your logs directory. For the initial archive job, the rule of thumb is 10G per day per Worker. If you do not use a third disk, then the logs will go to the OS partition and that could spell trouble. Also, if users access their archives often and perform PDF exports, that can grow as well. The xml, ebdb, and license directories are pretty much static with minimal to no growth. The backup directory is a backup of the index directory and other important items. However, if the disk begins to run out of room, you can easily copy this data over to a larger disk at some future time and point retain to that new disk.

Finally, if your database is on the Retain Server, you'll want a third or fourth disk for it (depending on whether you decided to dedicate a disk for your Retain logs).

If performance is an issue, not only will you want all three of these partitions/disks to be on physically different spindles (or at least written to a NAS/SAN with many disks that it can swipe across), but you'll benefit from putting indexes and the database on high speed drives. Your archive directory does not need the performance and can be on less expensive disk media.

Make sure to set the permissions of the new disks correctly in Linux or the installation will fail.

Disk Performance

This comes up a lot; thus, it bears repeating what was mentioned previously. A customer will claim that they are on a fast SAN. Often, it is not a SAN but a NAS and there are many considerations that go into performance. So, knowing that disk I/O is the top issue with archive job performance, it is best to plan out your disk storage accordingly.

Storage design and disk I/O has everything to do with Retain performance as archive jobs are I/O intensive. You have the following processes writing to disk simultaneously:

- The indexer to the [storage path]/index
- The database (if on the Retain server)
- The Retain Server to [storage path]/archive
- The Retain Server to the logs directory:
  - Linux: /var/logs/retain-tomcat8
  - Windows: [drive]:\Program Files\Beginfinite\Retain\Tomcat8\logs

With all of that disk activity, if a single spindle (drive) is having to handle all of it, then you can see that the performance bottleneck would be disk I/O. However, many disk systems these days involve multiple disks using (i.e., RAID 5 or RAID 10) that write the data across multiple disks. The more disks involved, the more you are spreading the load and typically the faster the disk performance will be. You also have a difference in drives (SATA/SAS/SSD). In those cases, you now are looking at whether the disks are local to the server or in a SAN/NAS.

RAID Considerations

Let's say your server employs RAID 5, which provides better redundancy than, say, RAID 10. If there were 4 disks. As you know, RAID 5 uses an extra parity bit that consumes an entire disk, which leaves it with 3 drives on which to stripe across. If one of those drives becomes unavailable, that leaves you with 2. Striping across 2 or 3 drives doesn't lend for great speed, especially if the disks are lower end SATA drives.

SAN / NAS Considerations

If on a SAN/NAS, now you are looking at the network link speed as well. You could have very fast drives, but if your link speed is 1 Gb/s, your bottleneck is going to be your link.

The 1 GB/s network link is slower than a SATA 2 or 3 connection (AKA SATA 3 Gb/s and SATA 6 Gb/s.) Your SATA 2 connection (which is now getting to be a pretty old standard) is 3x faster than a 1000 Mb/s network link (or 1 Gb/s network connection). A fast single HDD can saturate a 1 Gb/s connection but not quite a 3 Gb/s connection (SATA 2.0, or SATA 3 Gb/s) with a sequential read/write. 7,200 RPM platter drives usually top out around 160-170 MB/s (or 1.28-1.36 Gb/s).

Measuring Disk Performance

It really comes down to IOPS. Here is a very simple IOPS calculator: http://www.thecloudcalculator.com/calculators/disk-raid-and-iops.html or you can find one of your own.

So, it really comes down to you understanding your underlying disk storage. This article just gives food for thought. If you are running Retain on a VM guest server like most customers do, then you need to also understand your VM host and VM infrastructure. Is the Retain storage viewed by the

server OS running on the VM guest as "local" storage? If so, what type of disk system is holding your VM's datastore? If it is not local storage but the server is connecting to external storage, then you need to take a look at the external system's configuration.

Bottom line: Disk I/O performance is key to Retain's performance and there are several areas to investigate where the bottlenecks could be.

In addition to partition considerations, make sure that your storage is reliable. NFS mounts can be problematic, so you may want to shy away from those. NSS volumes are not supported, so do not use them.

RAM

The amount of memory depends on the number of active mailboxes you are archiving, the mail volume, your underlying hardware, and how your Retain system will be used.

We'll first discuss the concepts so that you can apply them to your system and we'll give you some general guidelines. In most instances, you will need to experiment with various memory configurations until you find what works best in your environment.

Concepts

Retain runs under Tomcat as shown at the beginning of this article and Tomcat runs on Java. The Retain Server will use the Java "heap" for its memory and the indexer will use the OS memory as well as virtual memory (see the Virtual Memory subsection below). For this reason, you should configure Tomcat/Java with the bare minimum to have it run in an acceptable fashion for you. If logins or Retain in general seems sluggish when in the mailbox or using the web admin tool, you may need more heap. The sweet spot for most systems with a single Worker installed on the local Retain server is 8 GB minimum (xms) and maximum (xmx). You want to leave as much RAM as possible for the Indexer, which uses non-heap RAM.

The amount of Java heap you set will depend on total RAM on your system and the number of Workers you install in addition to the default single Worker. As we grow in customer experience with Retain 4, we adjust this article's memory recommendations accordingly.

Right now, development has suggested 1 - 2GB per additional Worker beyond the 8 GB you normally would give to the Java heap for a system with a single Worker local to the Retain server; however, we've had a customer with 110 million messages with 7 Workers local to the Retain server get away with 8 - 10 GB of RAM, but that is really pushing it. They didn't run under that configuration for more than 24 hours, so we cannot tell whether it would have been successful in the long run.

The installer for Retain 4.0.1 and later tunes Tomcat/Java memory based on total RAM and which Retain components are installed. See the online manual's topic, "Tomcat Memory tuning" (note: that link goes to the 4.0.1 documentation, so if the link doesn't exist in the future, go to the online manual and find that topic). Again, as we learn more from customer experience, the installer's default RAM configuration is subject to change.

If you really want the fastest search performance, load it up with RAM, like 64GB or more. Systems with large numbers of messages (100 million or more) seem to be needing 64 GB of RAM or more. If you have a database system running on your Retain Server as well and multiple local Workers, then those will cut away at available RAM for the indexer, so you need to take that into account. The indexer wants to cache indexing data into RAM and memory access is much quicker than disk.

General Guidelines

All of this really depends on the priority you place on Retain performance. If a customer is only interested in getting data into Retain and it doesn't matter how long the archive jobs take (as long as they finish within a 24-hour timeframe) nor does the customer care how long it takes to search for messages (because they do not do it that often), then none of this matters.

The key test is how quickly tomcat shuts down and how much memory the OS is sending to swap. If tomcat is shutting down slowly, that's probably an indication that it has code in swap memory that it is having to call off of disk in order to close out. Reserving more memory for the OS should alleviate that problem; thus, reserve a minimum of 4G for the server OS right up front. On some systems, we have had to allocate more, on others, less. So, the key is to try different configurations on your system to see what makes the difference.

Once you have subtracted the OS memory from your total memory, give 2 - 4G of RAM to the database (if the database is on the same server; otherwise, the remainder can go to Tomcat). Note that Tomcat will need a minimum of 2G.

For small systems (1 - 250 mailboxes), 8G of RAM might be OK if that's all you can afford to allocate that. It is true that a small Retain system can run on 4G, but performance will be awful in most cases. If even decent performance matters, you really should not go lower that 8G unless you are a very small business and have 0 - 50 mailboxes. You might even want to consider trying 12 - 16G and see what difference that makes for you. For some, it will make a big difference. For others, it may make no difference as the bottleneck is elsewhere.

For medium sized systems (250 - 750 mailboxes), 12 - 16G of RAM should be considered.

For larger systems, 16G should be considered a minimum. Many large systems range from 24 - 48G of RAM. The more mailboxes and mail volume, the more RAM you might consider giving your Retain server. But, again, we have to emphasize that every system is unique and RAM may not be the biggest performance factor for them.

Case in point: We have a customer with 700 users that found allocating 24G of RAM made a big difference. In another case, a customer that had 1,500 users needed only 12G. We have systems with thousands of mailboxes and those systems do benefit from increased memory allocation, but their needs vary.

Tomcat Memory Configuration

Tomcat memory is manually configured. The latest version of Retain sets it to 8G by default. It is an industry best practice to set the minimum and maximum memory values to the same value.

In Linux:

You set the Tomcat memory parameters in a file called j2ee found at /etc/opt/beginfinite/retain/tomcat8. See Tomcat Memory tuning for more detail. Tomcat must be restarted after configuring it.

In Windows:

You can set Tomcat parameters by running Programs | Tomcat 8.0 | Configure Tomcat. Go to the "Java" tab to set them. Note, we also recommend setting the stack size to 256k (it defaults to 160k in Windows).

Database Memory Configuration

Since most organizations employing Oracle or MS SQL have someone designated as a database administrator (DBA), they typically understand memory configuration. What they need to know is that archiving speed and user mailbox browsing performance is affected by the amount of memory given to the Retain database.

Virtual Memory

If you have the available disk space, we recommend increasing the virtual memory to at least 50GB. In Linux, this is known as swap. In Windows, this is called the page file. Ideally, this swap or page file should be placed on a fast storage for performance reasons.

VM Configuration

VM (Virtual Machine) NIC Settings

We have found that using VMXNET3 for the network adapter in VMs helps performance.

Virtual Machine SnapShots

We have found that VM snapshots can reduce performance of the Retain Server. Keeping the number of snapshots to a minimum is highly recommended.

Resource Restrictions

See "Retain tomcat becomes unresponsive after several mailbox searches", which addresses a VM resource limitation setting that - if set - can affect the performance of Retain and cause it to freeze.

Retain Configuration

If your Retain system was installed with version 3.0 or older, we recommend checking some settings and consider making some adjustments:

Logging Level = Diagnostic

Both your Server and Worker logging levels should be set to "diagnostic" unless you do not have the available disk space on your root partition.

The Worker logging level is found under *Data Collection* | *Workers* | *[select the Worker object]* | *Logging:*.

The Server logging level is found under *Configuration* | *Server Configuration* | *Logging:*.

Indexing All Content

Indexing all content is now the default.

Workers

Under the *Connection* tab of the Worker object (*Data Collection* | *Workers*), be sure the port is set to *48080*. For systems created with Retain 3.2 and later, it defaults to port 48080; however, if the Worker object was created with an earlier version of Retain than 3.2, the port default was 80. This causes the Worker traffic to the Server to pass through the web server, which slows the job down. Avoid SSL unless required by security requirements. It will cut the job speed in half.

For larger systems, you can divide up your jobs and assign each job to its own Worker. Only one job can run at a time for a given Worker; thus, if you install multiple Workers, you now can have those jobs run simultaneously in parallel. We recommend that each Worker object (created in the Retain Server administration interface) be given a name that describes its location and - if on the Retain server with other Workers - the directory to which it was installed. For multiple Workers on the same server, each Worker program directory will be named RetainWorker (the first installed Worker) and RetainWorkerX (where "X" is the Worker number assigned during the installation).

For installing Section 3.9.5, “Installing Additional Workers,” on page 81 .

For installing Section 3.9.5, “Installing Additional Workers,” on page 81.

For Office 365 and Gmail, you will definitely want multiple Workers as archiving over the Internet from those systems is very slow - especially with Gmail. In those cases, create multiple distribution lists within those systems and create a separate job for each distribution list.

Archiving from Reload

Several years ago when GroupWise first supported the SOAP protocol, it was unstable. When Retain would run a job against a post office, the POA would often crash. Because this issue had such a big impact on customers, Micro Focus developed a method for archiving from a Reload backup designed to be a temporary solution until all of the SOAP issues in GroupWise were addressed. Novell eventually addressed the issues in their product and it became reliable and stable.

It is a recommended best practice to archive directly from the live post office and not from Reload. The archive job has a very light footprint on the post office - it only logs into one mailbox at a time and requests items. It would be like running the GroupWise client in caching mode and downloading mail on a daily basis. In addition to that, the performance and reliability of the archive jobs is much higher going straight to the production post office. There can be issues when archiving from Reload backups, especially if you don't stay right on top of the jobs to ensure they are completing properly.

We do not recommend running archive jobs against Reload unless there is a stability issue with your production post office. In such cases, that should be a temporary solution.

Profiles

GroupWise

Scope

One of the main points of having Retain is to have a complete archive of messages in your system because of data retention policies under which your organization must operate. If your email system's retention functionality is not understood or properly configured, users could purge items from their trash folder before Retain gets a chance to archive it. See the KnowledgeBase article: "How Retention Services and Item Store Flags Work".

Message Placement

Also, users make take a few days to organize their mail into folders. When the Retain archive job runs, it will take a "snapshot" of the user's mailbox as it exists at the time of archiving that mailbox.

If an Item A is in the main mailbox folder at the time of the job, it gets placed into the main mailbox folder in the Retain mailbox. If the user subsequently moves the item to another folder, Retain will not remove the item from the main Retain mailbox folder; instead, Retain simply creates an additional pointer to the item in the folder to which the user moved the message. This means the item will be found in the main mailbox folder and in the folder the user eventually moved it to.

To avoid this, some customers set the Date Range to Scan to Number of days from job start (older) and set the number of days to a value between 7 - 21. This gives the users that extra time to organize their new email before Retain archives it.

Miscellaneous

To save disk space, enable the option Don't store MIME.822 attachments and then activating the settings:

- Store/index Internet Headers
- Include Routing Properties

These settings will grab the data that is in the MIME.822 attachments and store and index them.

We also recommend only including shared folders that are owned by the mailbox. The other setting can cause errors in the job due to circular folder references.

Under Disabled/Expired Users, set "When Retain encounters a disabled or expired mailbox" setting to "Ignore".

Exchange/O365

Scope

One of the main points of having Retain is to have a complete archive of messages in your system because of data retention policies under which your organization must operate. If your email system's retention functionality is not understood or properly configured, users could purge items from their trash folder before Retain gets a chance to archive it. See the KnowledgeBase article: "How Retention Services and Item Store Flags Work".

Message Placement

Also, users make take a few days to organize their mail into folders. When the Retain archive job runs, it will take a "snapshot" of the user's mailbox as it exists at the time of archiving that mailbox.

If an Item A is in the main mailbox folder at the time of the job, it gets placed into the main mailbox folder in the Retain mailbox. If the user subsequently moves the item to another folder, Retain will not remove the item from the main Retain mailbox folder; instead, Retain simply creates an additional pointer to the item in the folder to which the user moved the message. This means the item will be found in the main mailbox folder and in the folder the user eventually moved it to.

To avoid this, some customers set the Date Range to Scan to Number of days from job start (older) and set the number of days to a value between 7 - 21. This gives the users that extra time to organize their new email before Retain archives it.

Miscellaneous

Enable these settings:

- Store/index Internet Headers
- Include user's recoverable items.

We strongly recommend abandoning the practice of using journal mailboxes to meet retention requirements (ensuring nothing gets deleted before being archived). Journaling doubles the traffic on your Exchange server and can adversely impact performance. Instead, have Retain archive the Recoverable Items folder in each mailbox and configure Rolling In-Place Holds. See the KnowledgeBase article: "How to Transition from Journaling to Rolling In-Place Hold for Exchange Archiving" on how to configure this. It also provides links to articles on other recommended settings in Exchange.

Reporting & Monitoring Server

This tool provides archive stats as well as archive job stats. In Retain 4, it becomes a very critical tool for identifying errors in your jobs because it not only shows the error counts, but through drill downs on the error numbers, an administrator can see the exact message on which the error occurred, the folder in which it is stored, its delivered date, and the error description.

If you use this tool, make sure that when configuring it, you check the box under Core Settings that reads, Disable the disk usage statistic. This utility is too much of a drain on system resources and will be re-designed in the future. In the meantime, disable it. In Retain 4, this utility is disabled by default.

And, like the Worker, the port settings under the Connection tab should all be set to 48080 if possible so that the Server and the R&M Server are talking directly to one another's tomcat port.

3.2 System Requirements

First decide where everything should be installed. Once the support and core components are installed, using Retain is the same no matter what platform it is installed on. You may spread out Retain’s compo-nents on different servers or run them all on one. VMWare ESX and ESXi is supported, but requirements are still in full effect and must be observed.

Retain may be configured in separate ways which heavily impact the requirements of the Retain server, indexing engine, database server, and locations dictate resource needs. Ultimately the number of active users in the system determine the resource needs of the system. See Planning and Best Practices for details.

3.2.1 Minimum System Requirements

- Dual-Core 3 GHz or better Processor

  NOTE: 8 cores provides the best performance during initial archiving and re-indexing processes.

- Server requires 12 GB RAM, installation will abort if less is detected

  NOTE: 16 GB is recommended.

- Worker requires 4GB RAM, if on a separate server

- Storage Space: As much storage space as messaging system requires. See Estimating Storage Requirements.
- Apache 2.2.x or IIS, installed by system administrator
- Supported SQL Database, installed by system administrator

RAM

Memory needs will vary depending on the load and use of the Retain system, however, Java will automatically tune for the memory which is available on the server. Tuning depends on what is selected for installation. The installation will detect the amount of ram and automatically assign memory values as shown

Less than 16 GB system memory:

8 GB: Server and Worker

6 GB: Server only

Less than 20 GB system memory:

12 GB: Server and Worker

10 GB: Server only

More than 20 GB system memory:

16 GB: Server and Worker

14 GB: Server only

In all cases where only the Worker is installed, the memory is tuned to 3 GB.

3.2.2 Single Server

(Retain Server, Retain Worker, Database, and Index engine housed on same machine and archiving less than 500 active mailboxes. The Retain Message Router should be installed on dedicated server in the DMZ.)

- 16-24 GB RAM

3.2.3 Remote Database

(Retain Server, Retain Worker, and Indexing engine on one server, separate Database server)

500-5,000 active mailboxes

- **12-16 GB RAM**

5,000-30,000 active mailboxes

- **16-32 GB RAM**

Dedicated Index servers*, Dedicated Database server, Dedicated Retain Server

50,000 Active mailboxes

- **16+ GB RAM for Retain Server (depending on environment)**
- **Database on appropriate server**

- **Retain Worker is installed remotely on each post office in the message system**

Disk Space is dependent on the amount of email being archived and length of retention time. It is recom-mended to estimate two years of storage and adding additional space as needed. Storage speed is a determining factor for performance. SAN and local storage preferred. NAS is not supported.

3.2.4 *High Availability Indexer

The High Performance Indexer can be installed on multiple external servers at the High Availability Indexer. This requires at least three (3) dedicated indexing servers. Please call support for more details.

3.2.5 Software installed by Retain

Retain will download and install Java Development Kit (JDK) 1.8 64-bit and Tomcat8, the Java webserver that sits on top of Apache or IIS. Instal-lation on an existing Retain system running a standard Tomcat installation will remove the current Tom-cat connection and install the embedded Tomcat that comes with Retain.

3.2.6 Supported Linux Versions

- SUSE Linux Enterprise Server 12 SP2 (64-bit)
- SUSE Linux Enterprise Server 12 SP1 (64-bit)
- SUSE Linux Enterprise Server 12 (64-bit)
- SUSE Linux Enterprise Server 11 (64-bit)
- Open Enterprise Server 11, 2015 (64-bit)

NOTE: The Retain system requires more than simply space. Due to the storage of many small files, especially on systems archiving mobile message data, sufficient inodes are also an issue. SLES 11 changed the way the EXT3 file system defaults, and unless modifications are made, the EXT3 file system on SLES 11 will become insufficient for large storage systems. For information on the limitations and how to fix it, see How to setup an EXT3 partition with more iNodes. If a substitute file system is desired, the XFS file system is flexible enough to mitigate the risk.

3.2.7 Supported Windows Versions

- MS Windows 2016
- MS Windows 2012 R2 64-bit
- MS Windows 2012 64-bit
- MS Windows 2008 (R2) 64-bit

  PowerShell 3.0 or higher

3.2.8 Microsoft Exchange Requirements

- Exchange Server 2013
- Exchange Server 2016

3.2.9 Novell GroupWise Requirements

- GroupWise 8.01 or above. (8.01, 2012, 2014, 2014R2)
- SOAP port must be enabled on all POAs to be scanned for messages.
- SSL is supported, but significantly impacts performance.

3.2.10 Android Requirements

- 4.4.x KitKat or higher

3.2.11 CellTrust Secureline

- CellTrust Gateway 4.2.0 +

3.2.12 BlackBerry Requirements

- BES 5.04, 10, 10.2, 12, 12.2
- BBM Protected

3.2.13 Supported Social platforms

- Facebook
- Twitter
- Yammer
- LinkedIn
- Yahoo Messenger
- Google Talk
- GMail
- Lync & Skype for Business
- AOL Instant Messenger

3.2.14 Supported Search

- Google
- YouTube
- Yahoo
- Bing
- Wikipedia

3.2.15 Supported SQL Database Management Software

- MySQL 5.6x or 5.7x (If using MySQL, the MySQL Connector/J (JDBC 5.1.21 or higher driver) must be provided for the Retain Server and / or Reporting and Monitoring Server.)
- Oracle 10/11g/12c

- MS SQL Server 2008 R2, 2012, 2014
- PostgreSQL 9

3.2.16 Supported Browsers

- Internet Explorer 10+

NOTE: Retain's support of IE 10 will be limited after Microsoft ends support for IE10 on Jan. 12, 2016

- Firefox 34+
- Chrome 39+

3.3 Ports in use

Retain utilizes many ports to facilitate communication between the different components and for communication with the different archived systems. By default, Retain will use the following ports for the fol-lowing services:

3.3.1 Retain Server

The Retain Server is the key component in the Retain system. Other Retain processes communicate with the Server through port 48080 by default. The Server is always listening on that port regardless of how other components might be configured to communicate with it (i.e., SSL port 443).

- 48080 (TCP) Requires incoming access if any Retain processes are running on a server external to the server hosting the Retain Server.
- The Server Requires outgoing access if the Reporting & Monitoring Server component has been installed on a server external the server hosting the Retain Server. See also the Reporting & Mon-itoring component in this article.
- 48009 (TCP) The AJP (Apache JServ Protocol) port is used by for communication between the web server and Tomcat. Since both should reside on the same server, there are no external port access requirements.
- 80 / 443 (TCP - HTTP / HTTPS) Requires incoming access to reach the Server web interface.
- 25 (TCP) Requires outgoing access so that the Retain Server can send email notifications on server errors, job statuses, and job errors. Outgoing access to Database Management System (DBMS) port. This depends on the database system you are using. See the "Database Management Systems" section of this page.

Other ports will also need to be opened on the server hosting the Retain Server depending on the mod-ules being used.

3.3.2 Retain Worker

The Retain Worker is the component that pulls the data from the messaging source, whether that be an email system, social media application, or mobile device.

- 48080 (TCP) Requires outgoing access if on a server external to the Retain Server.
- 80 or 443 (TCP HTTP or HTTPS) Requires incoming access to reach the Worker web interface

Other ports will also need to be opened on the server hosting the Retain Worker depending on the mod-ules being used; and, in some cases, on servers hosting the messaging system Retain will be archiving.

3.3.3 Retain Message Router

The Message Router is for customers of the Retain for Mobile module. The Message Router connects with mobile devices to handle SMS message log forwarding, BBMP device configuration. It typically would sit inside a DMZ.

- 443 (TCP) Requires both incoming from and outgoing access to the Internet as well as incoming from and outgoing access to the Retain Server.
- 111/2049 (UDP / TCP) for NFS Client services. (Only required if not using REST) Requires outgoing access if the Worker is not on the Message Router server so that it can place the logs on the Worker's server. Samba can be used, but NFS seems to be more reliable.

Other ports may need to be considered for NFS to work properly.

It is up to the customer to do this research to get NFS services to work properly.

3.3.4 Retain Reporting & Monitoring Server

New to Retain v3.4, this process provides archive job and server reporting and monitoring services. It is not installed by default unless specifically selected during the installation process.

- 48080 (TCP) Requires both incoming and outgoing access if on a server external to the Retain Server.
- 80 / 443 (TCP) Requires Incoming access to reach the R&M Server's web interface.
- 25 (TCP) If running on a server external to the Retain Server, then it requires outgoing access.

3.3.5 Retain Stubbing Server

The Retain Stubbing Server is the component that provides stubbing services to the Retain Server. It is rarely used or installed by customers. See the Administration and Users Guide for a listing of its advan-tages and disadvantages.

- 48080 (TCP) Requires outgoing access if on a server external to the Retain Server.
- 80 / 443 (TCP) Requires incoming access to reach the Stubbing Server web interface.

3.3.6 Retain for Social Media (RSM) Proxy Server

This is a VM appliance running as a proxy server for social media traffic going out to and coming in from the Internet. It logs this traffic so that Retain can archive that data. Port use will differ depending on net-work setup and internet access policies.

3.3.7 From RSM WAN IP to Untrusted, all TCP/UDP ports.

If the RSM WAN IP is a private IP, it needs be NATed to an appropriate routable IP address. The LAN IP address does not need a corresponding inbound NAT rule.

3.3.8 Database Management Systems

The following are the default ports these database management systems use, but they are configurable within those systems. Requires incoming access for the database server and outgoing access on the Retain Server.

- MySQL: 3306
- MS SQL: 1433
- Oracle: 1521
- Postgres: 5432

3.3.9 Blackberry

BES Web Services (SOAP) only supports the secure ports via TLS. Retain will initiate contact with BES Web Services if you are syncing the Address Book with the BES. Outgoing access is required for the Retain Server and incoming access for the BES server for the following ports:

- BES 12: 18084 (TCP)
- BES 10: 38443 (TCP)
- BES 5: 443 (TCP)
- 111/2049 (UDP / TCP) on the Blackberry Enterprise Server (BES) for NFS Server services.

The Blackberry module requires incoming access if the Worker is not on the BES server so that it can retrieve the BES logs. Samba can be used, but NFS seems to be more reliable.

Other ports may need to be considered for NFS to work properly.

It is up to the customer to do this research to get NFS services to work properly.

3.3.10 Exchange

- 80 / 443 (TCP HTTP / HTTPS) on the Worker server. Requires outgoing access. The autodiscover process will attempt https to the CAS server(s) for connecting to Exchange mailboxes. If that fails, it will use port 80 as a last resort. If Exchange / autodiscover / EWS are set up properly, only port 443 should be necessary.
- 3268 or 3269 (TCP). Requires outgoing access. It uses one of these ports for LDAP lookups to the global catalog host, which is the primary database server for Active Directory.
- Port 3268 for plaintext and 3269 for SSL (recommended).
- 53 (UPD). Requires outgoing access. It's the port used by DNS. Retain will do DNS lookups during its autodiscover process.

3.3.11 Google Apps Module

- 443 (TCP) Requires outgoing access for the Server (address book sync) and the Worker (it attempts to use the Gmail API for archiving; if that fails, it reverts to IMAP, thus the need for port 993 as described below).
- 993 (TCP). Requires outgoing access for the Worker(s) only. Under certain circumstances, the Worker may switch to using IMAP over SSL when requesting email from Gmail.

3.3.12 GroupWise

- 7191 (TCP) Requires outgoing access so that the Retain Server can download the Address Book. This is the default SOAP port the GroupWise POAs use, but this is configurable and is dependent upon the POA agent setting in GroupWise.

3.3.13 Mobile

- 80 / 443 (TCP) Requires both incoming and outgoing access so that the Retain Server and the Message Router can communicate device configuration information with each other. See also the Retain Message Router component in this article as well as the Mobile subsection under Retain Worker. See also the Retain Message Router component in this article.
- 111/2049 (UDP / TCP) for NFS Server services Requires incoming access if the Worker is not on the Message Router server so that the Message Router can place the logs on the Worker's server. Samba can be used, but NFS seems to be more reliable.

Other ports may need to be considered for NFS to work properly.

It is up to the customer to do this research to get NFS services to work properly.

3.3.14 Office 365

- 443 (TCP) Requires outgoing access. Retain uses SSL to connect with Office 365 to authenticate users logging in to Retain.
- 80/443 (TCP) Requires outgoing access to the Retain for Social Media proxy server appliance. Retain will make an http connection and request the "bundles". See also the Retain for Social Media (RSM) Proxy Server component in this article

3.3.15 CellTrust Secureline

- 111/2049 (UDP / TCP) on the CellTrust Secureline server for NFS Server services.

Requires incoming access if the Worker is not on the CellTrust Secureline server so that it can retrieve the CSV logs. Samba can be used, but NFS seems to be more reliable.

Other ports may need to be considered for NFS to work properly.

It is up to the customer to do this research to get NFS services to work properly.

3.4 Estimating Storage Requirements

It must be understood that no system storage requirement estimation can be expected to maintain or have any kind of reliable accuracy. Future mail use, litigation requirements, and compliance standards all may change and are unpredictable at best. Micro Focus bears no responsibility to accurately define or recommend storage needs for various messaging systems. Different messaging systems have different storage characteristics, and individual implementation renders general calculations invalid.

Keep in mind, however, that storage system performance will dictate mail view, indexing, and data performance. If the storage system is housed on slow hardware employing a file system designed for any other consideration than speed, (e.g. compact data storage), performance may be impacted. Consider what types of files are to be archived and access frequency when choosing the file system.

For example, a Retain system dealing with mainly or only archived mobile data (SMS, pin, MMS, &etc.) will perform best with different file system settings than a Retain system archiving a message system with large attachments.

For best results; pair current mail storage needs against projected future needs with the ability to easily add extra storage to the Retain system as needed. The ability to freely add additional storage space grants control and freedom over the messaging system and should be of paramount consideration. This practice is the only course which can be relied on with any confidence. Due to the challenges and circum-stances involved with each different system, (and even certain versions of different systems), only indi-vidual consideration will provide a reliable baseline for storage needs.

The simplest way to check current disk usage and storage requirement size is to monitor disk space usage on the mail servers and create a projection for the near future for needs.

However, the different options and variables between messaging systems make disk storage estimations so unpredictable that anything other than specific system monitoring cannot determine real disk usage. It is best to create a storage system where additional space may be added as required when existing space is consumed.

In addition, Exchange 2010 has abandoned single instance storage in favor of highly available perfor-mance, possibly causing multiple Exchange servers in the system to all have copies of the same data. Retain utilizes single instance storage and may vastly decrease the storage size of a system that heavily utilizes this feature. Due to the differences between the storage and main messaging system, it is nearly impossible to establish a baseline for Retain storage needs. Retain may tremendously decrease the needed size to archive an Exchange 2010 system, or, depending on system size and implementation, it may not significantly decrease the needs of the current system. Though Retain will require additional space to continue archiving mail, the first initial archive job will not exceed the size of the current mes-saging system.

Consulting with the Retain Sales representative will offer the best tailored information for each system and each implementation of the different platforms available.

3.4.1 Considerations for Storage Requirements

- **Operating System** The OS will see a number of updates over the life of a Retain server so that should be provisioned for.
- **Retain Program** The Retain program is updated often as well. This is often placed on the same volume as the OS.
- **Archive Storage Area** The Archive contains the item bodies and attachments. This is generally the largest percentage of the storage requirements on a Retain server. The bodies and attachments are stored as BLOB (Binary Large OBject) files in a folder structure that starts with /00/00/00 and grows to /FF/FF/FF for 16.7 million directories. BLOB files are stored only once, Retain implements a single-instance storage system, only a single copy of identical items is stored. BLOB files never change, they are only created or removed.

  The files are accessed when a message is opened in the Retain Search Message interface or downloaded by Retain Publisher.
- **Indexesa** The indexes allow for fast search of the data in Retain. When searching for items in Retain the indexes are used to return the results.

  For best search performance, the indexes should be placed on a fast disk and optimized regularly.

  The Index requires periodic optimization. This is set under *Server Configuration* | *Maintenance*. The Indexer requires as much free disk space as currently used index space for optimization. It will require three times as much if optimizing during an archive job.

During index migration from Retain 3.x to Retain 4.x, storage requirements for the indexes are the most complex and are described below.

- **Database** The database contained the header information for each item. For example: Sender, recipient, date received, and so on as well as pointers to body and attached files kept in the archive. Each item may be quite small but may be replicated many times in the database if there are many recipients.

  When browsing messages, the items in the database is being viewed. On larger systems, the database is often placed on a dedicated database server.

- **Logs** The logs track the actions that Retain takes. They can become quite large but are compressed at the end of each day and removed after 10 days, by default. This can be changed under *Server Configuration* | *Logging:*. Logs can be moved to another volume as described below.

3.4.2 Storage Minimums

Retain is very disk intensive and running out of disk space can very challenging to recover from and can potentially result in data loss. To minimize the chances of this happening certain safeties have been implemented.

The Retain system requires a minimum of 5GB of free disk space or it will enter maintenance mode.

Retain will send warning messages when there is less than 10GB free space for the storage, index, and system volumes.

These minimums are configurable, but it is not recommended to permanently change these settings:

1. Change to the configuration file directory, by default in:

Linux: */opt/beginfinite/retain/RetainServer/WEB-INF/classes/config/*

Windows: *C:\Program Files\Beginfinite\Retain\RetainServer\WEB-INF\classes\config*

2. Edit the *misc.properties* file

3. Change the following settings (in gigabytes) as desired, defaults being:

discspace.warn.gb=10

diskspace.error.gb=5

3.4.3 Indexes During Migration

While the migration is performed, a new index will be created, so the current index size will be doubled temporarily while two indexes exist. Once migration is complete, the old index may be removed and space reclaimed.

Also, the new index may be larger than the current index, due to the increased power and abilities of the new indexer. How much change there is, depends on the composition of the current archive, any limits set in indexing attachments, and how many attach-ments there are in the archive. If there are a lot of attachments, or very large attachments, then the increase in size will be significantly larger than if there are small or a limited number of attachments in the archive. With no limits on indexing attachments and a lot of large attach-ments in the system, an increase of up to 4 times the current index size is possible. With limits included, the new index may be smaller after upgrading. For the upgrade and a worst-case sce-nario, ensure that up to 5 times the current index size is available

before starting the index migra-tion. (With a current index size of 5 GB, ensure that the volume has a minimum of 20 GB free before starting the migration.) The average increase is expected to be around 20%.

If there is insufficient space to perform the migration, Retain will be placed into maintenance mode until more space is provided. Retain 4 checks for free disk space on the storage, index, and system locations. If there is less than 20 GB free space, Retain will send a warning message every 6 hours. At 10 GB of free space left, Retain will enter maintenance mode and all jobs are disabled. To exit maintenance mode, provide more free space on the volume.

NOTE: It is important to recognize that migration time and performance will largely depend on the performance of the storage system. External storage systems, such as appliances, may 'prioritize' seldom-used data to low performing storage and have a negative effect on migration performance. In addition, file systems of existing storage will largely dictate how fast indexing, migration, and message view can be performed. To ensure high performance, house the storage on performance hardware with a performance file system.

3.4.4 Change the Log File Location

Linux

To change the location of the log files, create a symbolic link (somewhat synonymous with Windows shortcuts). Retain, by default, gives you the option during installation to store the logs at /var/log/retain-tomcat8 or at /opt/beginfinite/retain/tomcat8. If stored at /var/log/retain-tomcat8, then Retain creates a symbolic link at /opt/beginfinite/retain/tomcat8 directory called 'logs" that points to the /var/... location.

1. Create the directory in the location you wish to use.
2. Stop tomcat.
3. Move the current logs to the new location: *mv /var/log/retain-tomcat8/* /*[path to new directory]
4. Make tomcat the owner of the new directory path: *chown -R tomcat:tomcat /*[path to new directory]
5. Set the appropriate file *permissions for the tomcat user and group: *chmod -R 664 /*[path to new directory]

   * If you are moving them to a separate volume, the file permissions must be 774.
6. Create a symbolic link in the parent directory of the default logs directory and point it to the new location:
   a. Change to the /var/log directory
   b. Remove the current log directory: *rm -r retain-tomcat8*
   c. Create a new symbolic link called retain-tomcat8 that points to your new log location: *ln -s* [path to new directory] *./retain-tomcat8* (and press ENTER)
7. Start tomcat.

Windows

1. Ensure no archive jobs are running and stop Tomcat.
2. Configure Tomcat's default log location.
   a. Click on Start.
   b. In the "Search programs and files" box, type: configure tomcat

c. Click on Configure Tomcat

d. Click on the Logging tab.

e. Type in the new log path using the standard Windows path utilizing backslashes "\" (i.e., d:\retain\logs) or browse to it by clicking on the button with "..." on it.

f. Click OK.

3. Make a backup copy of the existing *log4j.properties* file.
4. Edit the log4j.properties located at *[drive]:\Program Files\Beginfinite\Retain\[RetainServer, RetainWorker, RetainWorker1, &etc]\WEB-INF\classes*.
5. Do a search on *${catalina.base}/logs/* and replace it with *[desired path using forward slashes "/"]/logs/ (i.e., D:/retain/logs/)*. An easy way to do this is to use the text editor's (i.e., Notepad") "Replace" function.
6. Repeat steps 2 - 4 for every log4j.properties file (server, worker, stub server, stats server). The only log that will be created at the default log location will be the localhost.[date].log, but it is a very small log.

3.5 Starting and Stopping Tomcat

To start, stop, or restart either Retain; you start, stop, or restart Tomcat.

Here are instructions on starting, stopping, and restarting Tomcat on Windows and Linux:

Windows

- In Windows, it runs as a service (Retain Tomcat n), so simply start, stop, or restart that service.
- Retain 2/3: Apache Tomcat [n] (where "n" is the Tomcat version, 5/7)
- Retain 4.0-4.1.0.1: Retain Tomcat 7
- Retain 4.2 and higher: Retain Tomcat 8

Linux

The following commands are typed at the command prompt:

SLES 11

- **Start:** rcretain-tomcatx start
- **Stop:** rcretain-tomcatx stop
- **Restart:** rcretain-tomcatx restart
- **Status:** rcretain-tomcatx status
- Where "x" is the version of tomcat (i.e., Retain 2.x uses tomcat5, Retain 3.x and 4.0-4.1.0.1 uses tomcat7, and Retain 4.2 or higher uses tomcat8)

SLES 12

- **Start:** systemctl start retain-tomcatX.service
- **Stop:** systemctl stop retain-tomcatX.service
- **Restart:** systemctl restart retain-tomcatX.service
- **Status:** systemctl status retain-tomcatX.service
- Where "x" is the version of tomcat (i.e., Retain 2.x uses tomcat5, Retain 3.x and 4.0-4.1.0.1 uses tomcat7, and Retain 4.2 or higher uses tomcat8)

3.6 Tomckat Memory Tuning

3.6.1 General Recommendations

Retain automatically tunes on installation, depending on available system memory. Further tuning should not be necessary except in extreme conditions.

By default, Retain detects installed system memory and tunes the maximum memory value according to the following levels:

- Less than 16 GB system memory:
- 8 GB: Server and Worker
- 6 GB: Server only
- Less than 20 GB system memory:
- 12 GB: Server and Worker
- 10 GB: Server only
- More than 20 GB system memory:
- 16 GB: Server and Worker
- 14 GB: Server only

In all cases, when only the worker is installed the memory is automatically tuned to 3GB.

To manually tune the memory, carefully consider the following information.

3.6.2 Tomcat Configuration Parameters

There are three significant parameters:

1. Memory allocated upon Tomcat startup
    - Indicates how much memory is immediately allocated and reserved to Tomcat upon startup. This memory will be in use for the entire lifetime of Tomcat and never available to the other server processes.
    - Typically, this is set to 50%-100% of the maximum memory parameter discusses below. It can be lower, but pre-allocating a sizable percentage of memory enhances performance and reduces memory fragmentation.
2. Maximum Memory available to Tomcat
    - If the memory usage grows beyond the startup allocation, Tomcat will allocate additional blocks of memory in chunks as needed up to this limit. It will never return this memory to the general server memory pool. (Although memory internally will be freed and reused for Tomcat applications).
    - Tomcat guarantees that it will never exceed this memory allocation parameter. If Tomcat runs out of memory, it will try to reclaim unused memory via garbage collection. If this is insufficient, the web application will be denied the memory allocation. Unpredictable (but invariably unpleasant) results will then occur.
    - This is the most critical parameter to tune, and is the value which is automatically tuned by the installer.
3. Stack Size
    - For each thread (which includes each and every concurrent user request), Tomcat will allocate stack space.

- This value is typically measured in KB, and defaults to 512KB.
- It is allocated per concurrent users, and is in fact far too generous a number in general. 1000 users for example would take 500 MB of RAM just for stack space, before the program even allocates memory to run!
- Generally, we recommend reducing this number to 160k. Increase conservatively, in 64 KB chunks, if you see out-of-stack-space errors in the logs.

The parameters are abbreviated in the Catalina Opts line as follows:

Stack Size="-Xss<value>k"

Memory Heap="-Xms<value>g"

Maximum Memory=" -Xmx<value>g"

On Linux: edit the following configuration file:

- Retain 4.0-4.1.0.1

  /etc/opt/beginfinite/retain/tomcat7/j2ee
- Retain 4.2 or higher

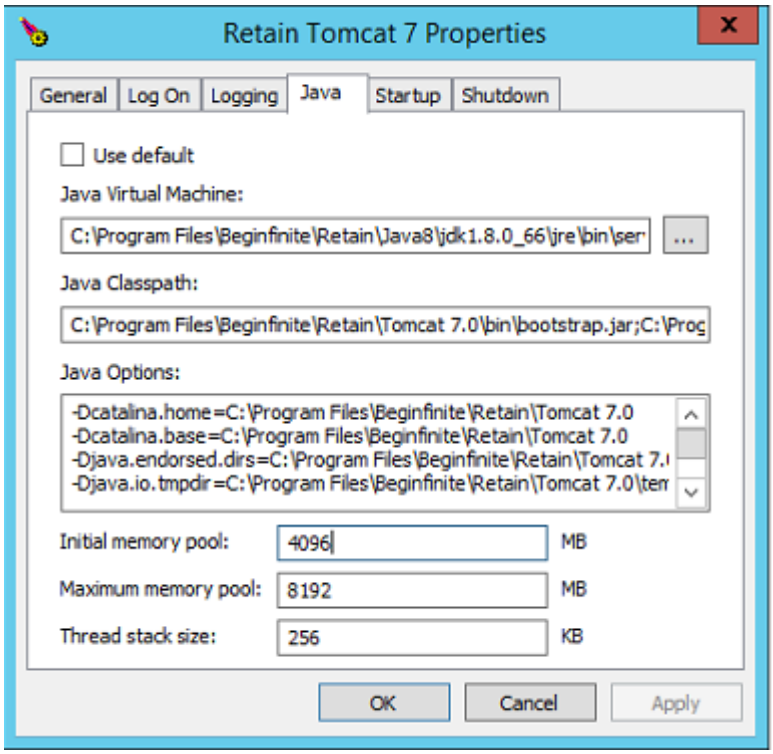
  /etc/opt/beginfinite/retain/tomcat8/j2ee
- Retain-Tomcat is automatically tuned depending on system memory. To tune the system manually, edit the following lines with the appropriate parameters for your system. For example:

CATALINA_OPTS="-Xms4g -Xmx12g -Xss-256k""

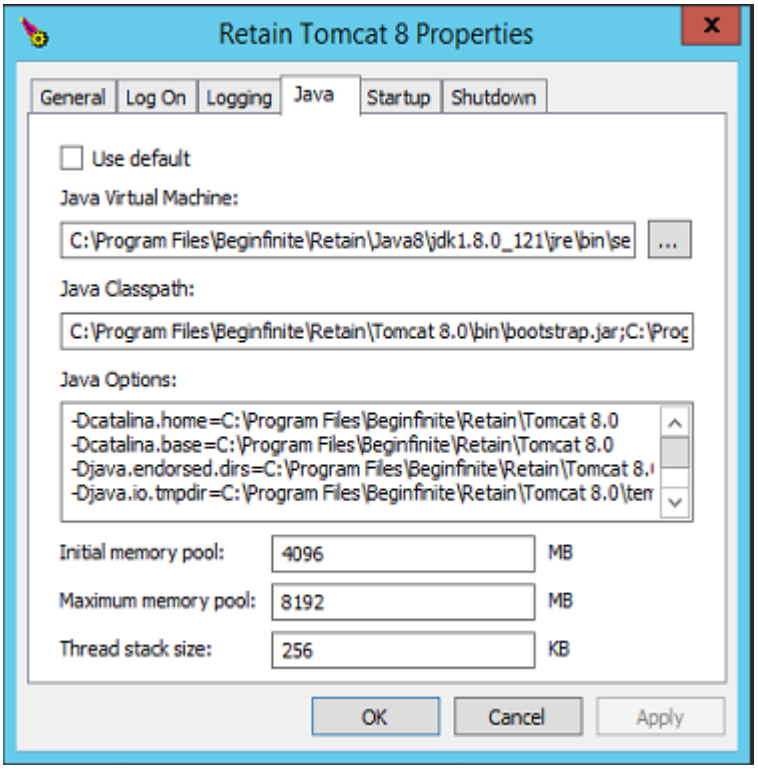
JAVA_OPTS="$JAVA_OPTS -XX:MaxmetaspaceSize=1024m"

On Windows: run the utility to configure Tomcat and modify the appropriate settings.

- Retain 4.0-4.1.0.1

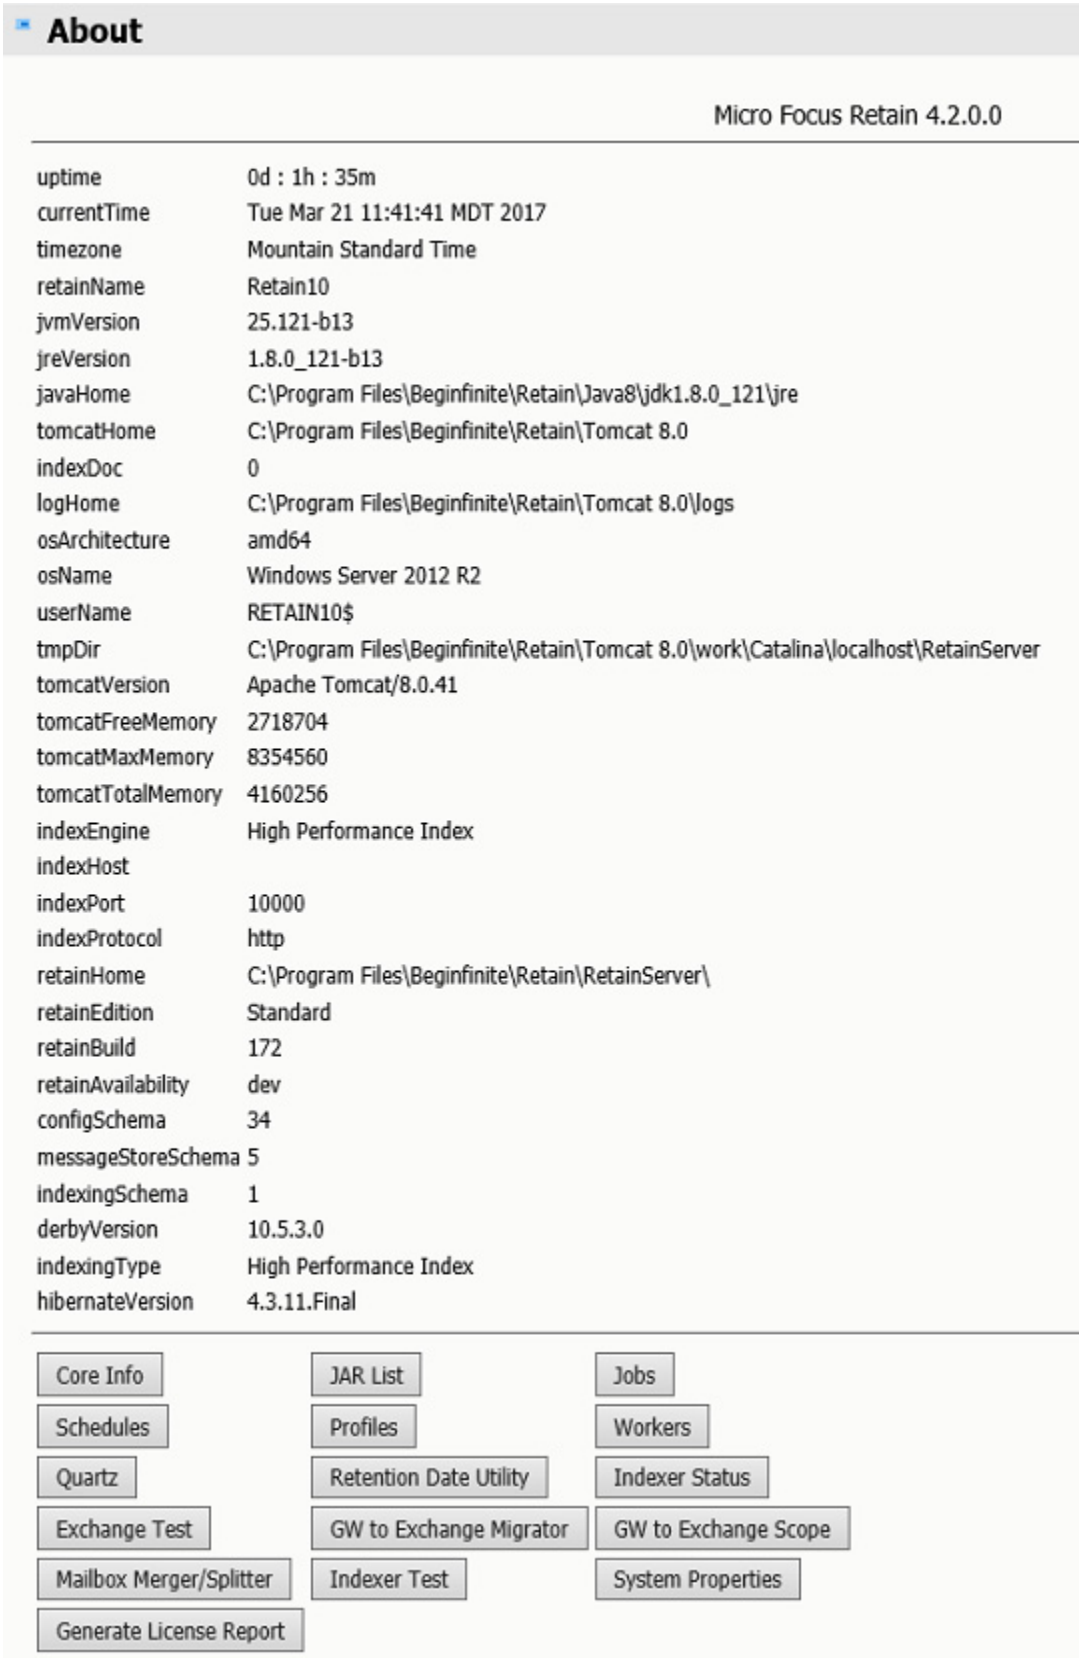


- Retain 4.2 or higher



3.6.3 Checking memory

After logging into Retain Server, click about link in the upper right of the interface. Among other stats, you'll see the tomcat Memory statistics.


About
Micro Focus Retain 4.2.0.0
uptime 0d : 1h : 35m
currentTime Tue Mar 21 11:41:41 MDT 2017
timezone Mountain Standard Time
retainName Retain10
jvmVersion 25.121-b13
jreVersion 1.8.0_121-b13
javaHome C:\Program Files\Beginfinite\Retain\Java8\jdk1.8.0_121\jre
tomcatHome C:\Program Files\Beginfinite\Retain\Tomcat 8.0
indexDoc 0
logHome C:\Program Files\Beginfinite\Retain\Tomcat 8.0\logs
osArchitecture amd64
osName Windows Server 2012 R2
userName RETAIN10$
tmpDir C:\Program Files\Beginfinite\Retain\Tomcat 8.0\work\Catalina\localhost\RetainServer
tomcatVersion Apache Tomcat/8.0.41
tomcatFreeMemory 2718704
tomcatMaxMemory 8354560
tomcatTotalMemory 4160256
indexEngine High Performance Index
indexHost
indexPort 10000
indexProtocol http
retainHome C:\Program Files\Beginfinite\Retain\RetainServer\
retainEdition Standard
retainBuild 172
retainAvailability dev
configSchema 34
messageStoreSchema 5
indexingSchema 1
derbyVersion 10.5.3.0
indexingType High Performance Index
hibernateVersion 4.3.11.Final
Core Info
JAR List
Jobs
Schedules
Profiles
Workers
Quartz
Retention Date Utility
Indexer Status
Exchange Test
GW to Exchange Migrator
GW to Exchange Scope
Mailbox Merger/Splitter
Indexer Test
System Properties
Generate License Report

Page File

The page file guide is included here as a courtesy. It is your responsibility to find the proper - procedures and documentation from your chosen operating system vendor. The entire responsibility for installation, care, and maintenance of the operating system lies with the customer. We do not provide any kind of support for the operating system.

In some environments, tomcat will run out of resources and crash. Increasing the page file may solve this issue. In some cases doubling the page file was sufficient, in other cases increasing it to 64GB was sufficient.

3.7 SQL Database Guides

The SQL start guides are included here as a courtesy. It is your responsibility to find the proper installa-tion procedures and documentation from your chosen database vendor. The entire responsibility for installation, care, and maintenance of the database server lies with the customer. We do not provide any kind of support for the database server.

Disclaimer

Retain Technical Support cannot support your Database Management System (DBMS) for several reasons including:

- Database corruption issues are caused by bugs in the database system itself or by issues in the underlying subsystem.
- Retain can be configured to use one of several database systems and there are too many combinations for our staff to become expert in all of them.
- Legally, we cannot provide support for the DBMS.

Section 3.7.1, “SQL Quick Start Guide,” on page 68

Section 3.8.1, “MySQL,” on page 69 and Section 3.8.2, “MySQL Connector/J (JDBC driver for MySQL) Installation,” on page 70

Section 3.8.3, “MS SQL Server 2008 R2, 2012, 2014,” on page 71

Section 3.8.4, “ORACLE 10/11g/12c,” on page 71

Section 3.8.10, “Postgres 9,” on page 73

Section 3.8.11, “ Further tuning,” on page 74

3.7.1 SQL Quick Start Guide

This guide is provided purely as a courtesy and will detail a very basic installation and database creation for Retain. This guide is provided with no warranty and no SQL database support is offered.

The database setup on all platforms proceeds as follows:

1. Download the SQL DBMS of your choice
2. Install it to the server of your choice, this may the Retain server itself, or on a dedicated database server
3. Log in and create a database for Retain to use

4. Create a user with rights for Retain
5. Configure Retain to connect to the database server and database with the user specified.

WARNING: NFS shares should not be utilized in any database system. NFS locking is insufficient for data-base requirements and will result in corruption.

3.8 Create a new database for Retain

3.8.1 MySQL

NOTE: Requires that Connector/J also be installed.

1. Download and install the latest 5.6x MySQL server and client and configure according to local system needs.

   Install the server: rpm -ivh MySQL-server-5.6.3.5-1.sles11.x86_64.rpm

   install the client: rpm -ivh MySQL-client-5.6.3.5-1.sles11.x86_64.rpm

   Modify /usr/my.cnf, to move the database to a data volume, by adding the line datadir=/path/to/where/you/want/the/MySQL/data

   Find the root password in /root/.mysql_secret

   cat /root/.mysql_secret

   Start MySQL:

   rcmysql start

   Change the root password:

   /usr/bin/mysql_secure_installation
2. Log in to MySQL using the client:

   mysql -u root -p
3. Create the database: * This is critical. Not storing the database in UTF-8 format guarantees irreparable corruption for non-US-ASCII characters. *

   CREATE DATABASE retain DEFAULT CHARACTER SET 'utf8' DEFAULT COLLATE 'utf8_bin';

   3a. If you make a mistake, enter this command:

   DROP DATABASE retain;

   …and repeat step 3 to create the database.
4. Retain needs a user account created that has full rights to the Retain database. It is strongly recom-mended that the ‘root’ account is not used. While logged in to MySQL, enter the following commands. Replace (UserName) and (Password) with the desired user name and password (Both commands are needed. The ‘%’ is a wildcard. If you want to restrict the connection to a specific IP address, put that address instead of the ‘%’ sign.):

   GRANT ALL PRIVILEGES ON retain.* TO '(UserName)'@'%' IDENTIFIED BY '(Pass-word)';

   GRANT ALL PRIVILEGES ON retain.* TO '(Username)'@'localhost' IDENTIFIED BY '(Password)';

   For example, if the database user's name is "retain" and the password is "Password1":

GRANT ALL PRIVILEGES ON retain.* TO 'retain'@'%' IDENTIFIED BY 'Password1';

GRANT ALL PRIVILEGES ON retain.* TO 'retain'@'localhost' IDENTIFIED BY 'Password1';

3.8.2 MySQL Connector/J (JDBC driver for MySQL) Installation

The MySQL Connector/J driver must be provided to the Retain Server and/or the Reporting and Monitoring Server. This must be installed *after* Retain has been installed but *before* Retain is configured.

1. Download the MySQL Connector/J (JDBC driver for MySQL) version 5.1.21 or later ZIP file.
2. Extract the mysql-connector-java-5.1.X-bin.jar file (*X being the current version number*) from the ZIP. For example, on Linux use the command:

   unzip mysql-connector-java-5.1.41.zip
3. Stop Tomcat service

SLES12:

service retain-tomcat8 stop

SLES11:

rcretain-tomcat8 stop

Windows:

Start | Services | Retain Tomcat | Stop the service

4. Copy the mysql-connector-java-5.1.X-bin.jar file (*X being the current version number*) to the following Retain directories or where Retain was installed:

   /opt/beginfinite/retain/RetainServer/WEB-INF/lib (for Retain Server)

   /opt/beginfinite/retain/RetainStatsServer/WEB-INF/lib (for the Reporting and Monitoring Server)

   /opt/beginfinite/retain/RetainStubbingServer/WEB-INF/lib (for the Stubbing Server)

   For example:

   cp mysql-connector-java-5.1.41-bin.jar /opt/beginfinite/retain/RetainServer/WEB-INF/lib/

   a. If on Linux, change ownership of the file to match the other files with the terminal command:

      chown tomcat:www mysql-connector-java-5.1.41-bin.jar

   b. If on Linux, change rights of the file to match the other files with the terminal command:

      chmod 774 mysql-connector-java-5.1.41-bin.jar
5. Start Tomcat service

SLES12:

service retain-tomcat8 start

SLES11:

rcretain-tomcat8 start

Windows:

Start | Services | Retain Tomcat | Start the service

Schema update note: On Linux, some systems have encountered an error after a schema update caused when MySQL cannot create a temporary file for the result of a Retain query.

The following appears in the Retain Server log:

2008-10-27 00:00:59,786 [TP-Processor2] ERROR com.maintainet.dao.HibernateUtil - SCHEMA UPDATE FAILURE: Something went wrong during Schema Update - contact tech support immediately

2008-10-27 00:00:59,786 [TP-Processor2] ERROR com.GWAVA.utils.ErrorHandle - reportError: SchemaUpdate :: EXCEPTION : org.hibernate.exception.GenericJDBC-Exception: could not get table metadata: Audit org.hibernate.exception.Gener-icJDBCException: could not get table metadata: Audit

This is usually caused when MySQL cannot create or has lost rights to the tmp working directory. Try the following from a system terminal:

1. Add rights to the tmp working directory:

   mkdir /var/lib/mysql/tmp

   chown mysql:mysql /var/lib/mysql/tmp
2. Add the following line into the [mysqld] section of /etc/my.cnf:

   tmpdir = /var/lib/mysql/tmp
3. Restart the Server

3.8.3 MS SQL Server 2008 R2, 2012, 2014

Install as default instance, with *Latin 1_ General_C1_AS* as standard encoding.

1. Enable SQL Server Authentication in addition to Windows Authentication
2. Once installed, verify the server is listening on port 1433 (on the command line run: telnet <IP address> 1433)
3. Ensure TCP/IP is enabled - Run SQL Server Configuration Manager
4. Expand SQL Server Network Configuration
5. Protocols for MSSQLSERVER
6. Enable TCP/IP
7. Run SQL Server Management Studio, connect to Server
8. Under Security/Logins, right click and select "create Login, (for example Retainuser) that uses SQL Server authentication. Assign a password.
9. Right click on Databases, and create a new database named Retain and assign Retainuser as the owner. (There are other ways to grant the permissions but this is easiest)

3.8.4 ORACLE 10/11g/12c

Retain supports the usage of Oracle 10, 11g, and 12c on all platforms. However, it is the customer’s responsibil-ity to employ an experienced Oracle DBA to install, maintain, and tune Oracle.

The instructions below serve as a guideline, but may result in poorly performing, insecure environments.

Oracle Installation instructions vary wildly depending the version, OS, storage, clustering, etc. so consult your DBA. There is no one-size-fits-all. This is also true of other SQL Server products, but Oracle even more so.

Using Oracle with Retain consists of:

- Installation of the Oracle Server
- Setting up a TCP IP listener
- Setting up a new database named Retain
- Creating a user to access the database, and granting sufficient rights.

3.8.5 Installation

Install your Oracle Server normally, according to your standards and practices. One option is to use the Universal Installer, specifying a custom installation, choose the type of Oracle Server to install, and to Install Files Only

3.8.6 TCPIP Listener

Next use the Net Configuration Assistant (or manually edit the Oracle Listener configuration) to create a TCP based listener on port 1521. This allows TCPIP clients to connect to Oracle and is required for Retain.

3.8.7 Database Creation

Next, either use the Database Configuration Assistant or manually create the database with your favorite tool. Regardless, you want to create a database and sid, named Retain, and create your SYSTEM account

An example of a create statement (which should not be used as is):

```
connect "SYS"/"&&sysPassword" as SYSDBA
set echo on
spool C:\oracle\product\10.2.0\admin\Retain\scripts\CreateDB.log
startup nomount pfile="C:\oracle\product\10.2.0\admin\Retain\scripts\init.ora";
CREATE DATABASE "Retain"
MAXINSTANCES 8
MAXLOGHISTORY 1
MAXLOGFILES 16
MAXLOGMEMBERS 3
MAXDATAFILES 100
DATAFILE 'C:\oracle\product\10.2.0\oradata\Retain\system01.dbf' SIZE 300M REUSE
AUTOEXTEND ON NEXT  10240K MAXSIZE UNLIMITED
EXTENT MANAGEMENT LOCAL
SYSAUX DATAFILE 'C:\oracle\product\10.2.0\oradata\Retain\sysaux01.dbf' SIZE 120M
REUSE AUTOEX-TEND ON NEXT  10240K MAXSIZE UNLIMITED
SMALLFILE DEFAULT TEMPORARY TABLESPACE TEMP TEMPFILE
```

```
'C:\oracle\product\10.2.0\ora-data\Retain\temp01.dbf' SIZE 20M REUSE AUTOEXTEND ON
NEXT 640K MAXSIZE UNLIMITED
SMALLFILE UNDO TABLESPACE "UNDOTBS1" DATAFILE 'C:\oracle\product\10.2.0\ora-
data\Retain\undotbs01.dbf' SIZE 200M REUSE AUTOEXTEND ON NEXT 5120K MAXSIZE
UNLIMITED
CHARACTER SET AL32UTF8
NATIONAL CHARACTER SET AL16UTF16
LOGFILE GROUP 1 ('C:\oracle\product\10.2.0\oradata\Retain\redo01.log') SIZE
51200K,
GROUP 2 ('C:\oracle\product\10.2.0\oradata\Retain\redo02.log') SIZE 51200K,
GROUP 3 ('C:\oracle\product\10.2.0\oradata\Retain\redo03.log') SIZE 51200K
USER SYS IDENTIFIED BY "&&sysPassword" USER SYSTEM IDENTIFIED BY
"&&systemPassword";
spool off
```

The most critical items to note are the CHARACTER SET (AL32UTF8 aka UTF8), and the NATIONAL CHAR-ACTER SET (AL16UTF16). If using a GUI based installed, be careful to set these - they are easily missed. (They are located on the Encoding tab of the 10th step or so of the wizard).

3.8.8 User Account Creation

Finally, using the SYSTEM account, connect to the database, create a user, and grant full system privileges to the account. An example is shown below:

```
CREATE USER user-name IDENTIFIED BY password;
grant dba to user-name;
```

Restart everything, and verify you can telnet on the TCP LISTENER's PORT (1521 if as above).

If a restricted user is desired, Retain requires a user with the following roles:

```
RESOURCE CONNECT
And the following database privileges to the Retain database:
CREATE PROCEDURE
CREATE SESSION
CREATE SYNONYM
CREATE TABLE
CREATE TYPE
CREATE VIEW
```

3.8.9 Tuning

Tuning Oracle is quite far beyond the scope of this manual. The manuals on Oracle's site, as well as an experienced DBA or extensive training is strongly recommended.

3.8.10 Postgres 9

When creating a PostgreSQL database for use by Retain, specify UTF8 encoding. If the Retain Server is on a different machine, you will need to configure PostgreSQL to accept connections from that machine's IP address.

1. Connect to PostgreSQL using its psql utility.
2. Create a new database and user for use by Retain. In the example below, we are creating a user named "retainuser":

   CREATE DATABASE retain WITH OWNER = retainuser ENCODING = 'UTF8' TABLESPACE = pg_default CONNECTION LIMIT = -1;

3.8.11 Further tuning

It is the customer's responsibility to learn more about and tune the SQL Database Server. Tuning data-base servers for performance often requires an experienced DBA. It is the customer's responsibility to back up, store, and make the database is fault tolerant.

3.9 Linux Installation

3.9.1 Prerequisites

Before running the Retain installer, you must first ensure that the following are installed and working:

- **Apache 2.2.x (SLES11) or Apache 2.4 (SLES12):**
- **SQL server with appropriate database created.:** For more informa-tion on database installation, or configuration and creation of an appropriate database for Retain, see the 'SQL Database Guides section'. The SQL server does not need to be on the same server, though for small systems having the SQL server local to the Retain Server simplifies the setup. To create the database, follow the instructions in the 'SQL Database Guide Section'.

The Retain installation requires an *internet connection Java Development Kit 1.8* to function properly. The Retain install will auto-matically download and install the appropriate *Java Development Kit 1.8* if the server has a connection to the internet.

NOTE: Be sure your server's clock is accurate before you begin! Using NTP is recommended to keep your clock accurate.

The install will complete the following tasks:

- Java Development Kit 1.8 (JDK) will be checked/installed.
- Retain_tomcat and any existing Retain installation will be stopped
- Retain will be unzipped into /opt/beginfinite/retain.
- The data will be stored in /var/opt/beginfinite/retain by default.
- Backups of modified files go in /opt/beginfinite/retain/backup
- The Apache integration file is in /opt/beginfinite/retain/ws
- Tomcat is made to restart and the runlevel is altered to run automatically on system start.

3.9.2 Retain Install

Pre-install tasks:

1. Download the latest version of Retain.

   You may use the command line "wget http://download.gwava.com/download.php?product=Retain&version=current"
2. Extract the archive. This will extract into a new folder.

   In the GUI right-click and select "Extract Here"

   On the command line use "unzip Retain<version>.zip"

3. Open a terminal window in the Retain install folder and make the installer script executable by running the command: "chmod +x *.sh"

   If a previous version of Java was installed via YaST, tomcat may fail to start. Uninstall the existing Java with YaST and unset the JRE_HOME path (unset JRE_HOME).

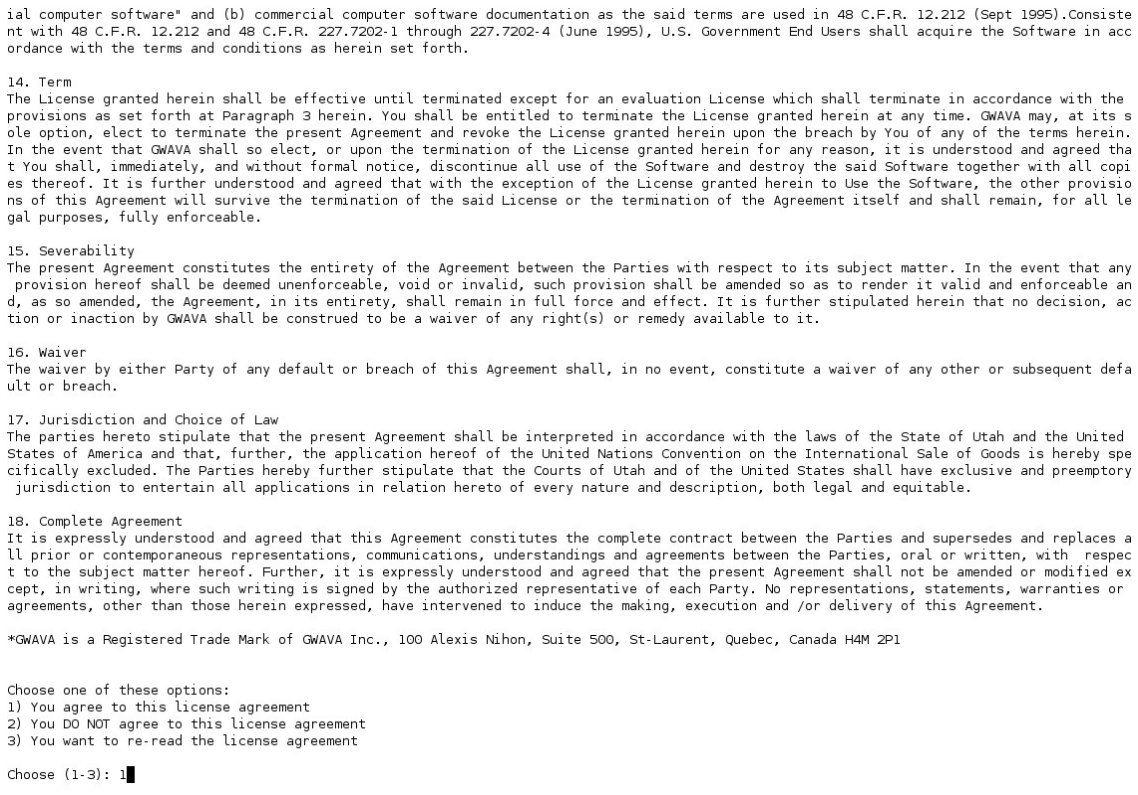
Retain Install:

(The installer installs Retain-Tomcat and initialize the Retain installation which will install the embedded Tomcat as well. Consider removing any existing Tomcat installation.)

1. Install Retain. Enter the following command:

   ./RetainInstall.sh

2. Read and accept the license agreement.



3. Read the software requirements and press Enter.

```
Retain requires the following:
        - Apache 2.2.x or 2.4.x
        - Tomcat 8.0 (Installer will install it).
        - Java 8 (Installer will install it).
        - Supported database (not necessarily on this server), with database cre
ated as per install guide.
        - MySQL Connector/J (JDBC driver for MySQL) version 5.1.21 or later if u
sing MySQL database


Minimum memory requirements for different install scenarios:
        1- Any combination of Server, Reporting and Monitoring Server, and Messa
ge Router: 8 GB RAM
        2- Worker: 4 GB RAM
        3- Worker and any other component: 12 GB RAM


After running Setup: Be sure to read the documentation on tuning memory usage in
 Tomcat.



Press Enter to continue.
```

4. You will be asked what components to install. Each of these components can be installed on a separate server or on the same server. However, the Message Router needs to be in the network DMZ and is usually installed separately. Options include:

   Retain Server

   Stubbing Server (GroupWise only)

   Reporting and Monitoring Server

   Retain Worker

   Retain Message Router (needs to be installed in the network DMZ)

   Generally, we recommend installing the Server, Reporting and Monitoring, and Worker on the Retain server.

```
================================================================================

                        Retain Installation Options

================================================================================

Choose components...
Install Retain Server (y/n)?y
Install Retain Stubbing Server (y/n)?n
Install Retain Reporting & Monitoring Server (y/n)?y
Install Retain Worker (y/n)?y
Install Retain Message Router (y/n)?n
```

a. If you are upgrading Retain, you are given a list of currently installed options and the ability to select upgrade options, if you wish to add components you can choose which ones to add here. Otherwise during a new installation, Java will be installed.

```
Retain has already been installed on this system with the following components:

Retain Server
Retain Reporting & Monitoring Server
Retain Worker

Choose one of these options:
1) Upgrade current installation of Retain
2) Upgrade current installation of Retain and add new components
3) Exit this installer

Choose (1-3):
```

b. During an upgrade the installer will check the Apache/Tomcat configuration.

NOTE: If components already exist on the system, but new features are being added, all new and existing features must be selected for install, AND the upgrade must be selected with the first option: “Upgrade Retain, Preserving configuration” must be selected.

```
================================================================================

                        Apache/Tomcat Configuration

================================================================================

Apache 2 appears to be installed
Apache Base is /etc/apache2
Apache Service is /etc/init.d/apache2
Apache Conf is /etc/apache2/httpd.conf
Apache Version is 2.2
Apache Group is www
Will assume that mod_proxy, mod_rewrite, and mod_proxy_ajp
are installed and enabled. These are standard in Apache 2.2+.
You may check in the Http Server control Panel in Yast

/opt/beginfinite/retain exists. Choose one:
1) Upgrade Retain, preserving configuration (recommended)
2) Overwrite Retain, losing configuration
3) Quit

Choose (1-3):
```

5. Java is checked, and automatically installed or upgraded, if necessary.

```
================================================================================

                          JAVA Configuration

================================================================================
Checking Java...

Downloading Java from http://download.gwava.com/download.php?product=jvm8&version=alin64server ...
--2016-11-04 09:05:13--  http://download.gwava.com/download.php?product=jvm8&version=alin64server
Resolving download.gwava.com... 209.90.108.66
Connecting to download.gwava.com|209.90.108.66|:80... connected.
HTTP request sent, awaiting response... 302 Found
Location: http://download.gwava.com/jvm/server-jre-8u60-linux-x64.tar.gz [following]
--2016-11-04 09:05:17--  http://download.gwava.com/jvm/server-jre-8u60-linux-x64.tar.gz
Connecting to download.gwava.com|209.90.108.66|:80... connected.
HTTP request sent, awaiting response... 200 OK
Length: 59798027 (57M) [application/x-gzip]
Saving to: `jvm.tar.gz'

39% [================================>                                      ] 23,608,670  10.1M/s
```

6. Installation completes, shows the URL to the RetainServer web console and returns to the prompt.

```
From your browser, if the Apache integration works, you should be
able to type http://127.0.0.1/RetainServer and/or
 http://127.0.0.1/RetainWorker and/or
 http://127.0.0.1/RetainRouter.

An additional diagnostic if that fails:
Try http://127.0.0.1:48080/RetainServer and
 http://127.0.0.1:48080/RetainWorker or
 http://127.0.0.1:48080/RetainRouter .
The direct Tomcat ports may be used to troubleshoot Jakarta and the
proxy connector. If TOMCAT's ports work but Apache's don't, double
check Jakarta and the proxy connector settings.

If the TOMCAT install fails, it is likely due to a missing or
incorrect JAVA_HOME variable. Verify or reset the JAVA_HOME variable.
See the installation guide for steps to set the JAVA_HOME variable.

IMPORTANT: Be sure to read the documentation on tuning Tomcat's
memory usage. The default memory may not be sufficient.
Backup /root/Desktop/Retain4100dev_218/RetainInstall.log to /opt/beginfinite/retain/backup/20161104/RetainInstall.log
sf105:~/Desktop/Retain4100dev_218 #
```

7. Check the status of Apache 2 and Retain Tomcat. (They should both be 'running'.)

Controlling Tomcat

SLES12

To check the status, start or stop apache or retain-tomcat, use the following runtime commands:

- service apache2 status|start|stop|restart
- service retain-tomcat8 status|start|stop|restart

SLES11

To check the status, start or stop apache or retain-tomcat, use the following runtime commands:

- rcapache2 status|start|stop|restart
- rcretain-tomcat8 status|start|stop|restart

8. If using MySQL, install the MySQL Connector/J (JDBC driver for MySQL) before configuring Retain.

You are now ready to continue to the configuration tasks. See the Initial Setup section.

If you wish to set a custom storage path or install additional workers, those instructions are continued below.

3.9.3 Custom Storage Path Rights

If the data is to be stored on a disk other than the default volume, follow these instructions to set the directory ownership and permissions.

During initial configuration, a storage path for the archived mail is selected. Retain requires read and write access and permissions to the storage path. If a storage path other than the default (/var/opt/beginfinite/retain) is desired, then rights and ownership need to be granted to Retain.

After running ./retainInstall.sh, find out which user Retain is using to access the file system. (For SLES 10.1, 10.2, 11 the user is called ‘tomcat’).

In a console, get a listing of the files and owners by typing:

```
l /opt/beginfinite/retain
```

This shows the User and the group that own the directory structure. In the system shown, Retain is using tomcat as the user.

Create the directory structure for the desired storage location, and then grant the permissions and own-ership to the Retain user.

```
(mkdir <desired directory structure>)
(chown -R <user> <storage directory>)
(chmod -R 744 <storage directory>)
```

(In this example, we have already mounted our SAN or physical drive to the /Retain directory and wish to store mail in the /Retain/archives directory.)

```
mkdir /retain/archives
chown -R tomcat:www /retain/archives
chmod -R 744 /retain/archives
```

Now check the directories for ownership:

```
l /retain
```

```
vm-sles10-retain:/retain # l /retain
total 1
drwxr--r--  3 tomcat root  72 Mar  7 16:06 ./
drwxr-xr-x 25 root   root 608 Mar  7 16:01 ../
drwxr--r--  2 tomcat root  48 Mar  7 16:06 archives/
vm-sles10-retain:/retain #
```

Tomcat now has rights to the storage location.

```
vm-sles10-retain:~ # l /opt/beginfinite/retain
total 1
drwxrwxr--  6 tomcat www  160 Mar  7 13:03 ./
drwxr-xr-x  3 root   root  72 Mar  7 13:03 ../
drwxr-xr-x 14 tomcat www  528 Mar  7 12:58 RetainServer/
drwxr-xr-x  7 tomcat www  416 Mar  7 12:58 RetainWorker/
drwxrwxr--  3 tomcat www   72 Mar  7 13:03 backup/
drwxrwxr--  2 tomcat www   88 Mar  7 13:03 ws/
vm-sles10-retain:~ #
```

3.9.4 Installing Just a Worker to a Server

The Worker may be installed on a server other than the Retain server itself. For example, the worker may be installed on the email server or on a dedicated worker server.

This is much simpler than installing the entire Retain server system. It only requires installing the webserver and the worker agent software. It uses the same installer as above.

Prerequisite:

1. Install the web server (apache)

Pre-install tasks:

1. Download the latest version of Retain.

   You may use the command line "wget http://download.gwava.com/download.php?product=Retain&version=current"

2. Extract the archive. This will extract into a new folder.

   In the GUI right-click and select "Extract Here"

   On the command line use "unzip Retain<version>.zip"

3. Open a terminal window in the Retain install folder and make the installer script executable by running the command: "chmod +x *.sh"

   3a. If a previous version of Java was installed via YaST, tomcat may fail to start. Uninstall the existing Java with YaST and unset the JRE_HOME path (unset JRE_HOME).

4. Run the installer. Enter the following command:

   ./RetainInstall.sh

5. Select *yes* to "Install Retain Worker" only
6. Finish the install as normal
7. Workers are accessed by browsing to *http://[IPAddress or HostName]/RetainWorker*

3.9.5 Installing Additional Workers

Overview

In larger email systems, it can take a long time to dredge all the users especially if the users are on multiple mail servers. You can install multiple workers that can run in parallel to speed up the archiving process. We generally recommend one worker per mail server. The Retain installer can install up to 10 workers on a single server. We recommend an additional 1GB for RAM per worker. Workers are accessed by browsing to http://[IPAddress or HostName]/RetainWorkerN

The installer allows you to install one additional worker at a time.

Procedure

1. From the terminal run "./retaininstall.sh" and install at least one Worker or upgrade the existing installation, if you haven't already done so.
2. Run "./retaininstall.sh addworker" to start the installer in Add Worker Mode.
3. Acknowledge requirements.
4. Accept the license agreement.
5. Enter the number (1-9) of the Worker agent you wish to install. You may install workers in any order.

```
================================================================================

                        Retain Multiple Worker Installation

================================================================================

Current RetainWorker instances appear to be
/opt/beginfinite/retain/RetainWorker

Below you'll get to select a number from 1 to 9
Which represents a Worker install. For example,
choosing 5 installs a Worker named RetainWorker5

Choose (1-9, or Q to quit): 
```

6. Confirm the install location.
7. Choose to keep or move the log location.
8. Choose to run retain-tomcat8 on start up or not.
9. After installing and acknowledging the location of the documentation the installer will show you how to access the Workers that were just installed.
10. Log into the Retain web console as admin
11. Create a new Worker and download the Bootstrap for each new Worker you installed. You may save it with a different name to organize it.
12. Browse to the new Worker at http://[ServerAddress]/RetainWorkerN where N is the number of the worker you just installed.
13. Upload the bootstrap.
14. Click Continue and login to the worker.
15. Check the friendly name that it is the one you expect.

    There is a separate log file for each worker, found in /var/log/retain-tomcat8/ RetainWorkerN.yyyy-mm-dd.log unless otherwise specified in Configure Tomcat.

3.10 Windows Installation

3.10.1 Prerequisites

Before running the Retain installer, you must first ensure that the following are installed and working:

- IIS or Apache 2.2.x
- PowerShell 3.0 or higher
- Section 3.7, “SQL Database Guides,” on page 68

The Retain installation requires an internet connection to function properly. The Retain install will auto-matically download and install the appropriate Java Development Kit 1.8 if the server has a connection to the internet.

NOTE: IIS has a size limit restriction which will limit uploads, or message and attachment archiving to 30MB, by default. For more information see http://support.gwava.com/kb/?View=entry&EntryID=2089

NOTE: Be sure your server's clock is accurate before you begin! Using NTP is recommended to keep your clock accurate.

3.10.2 Webserver Install

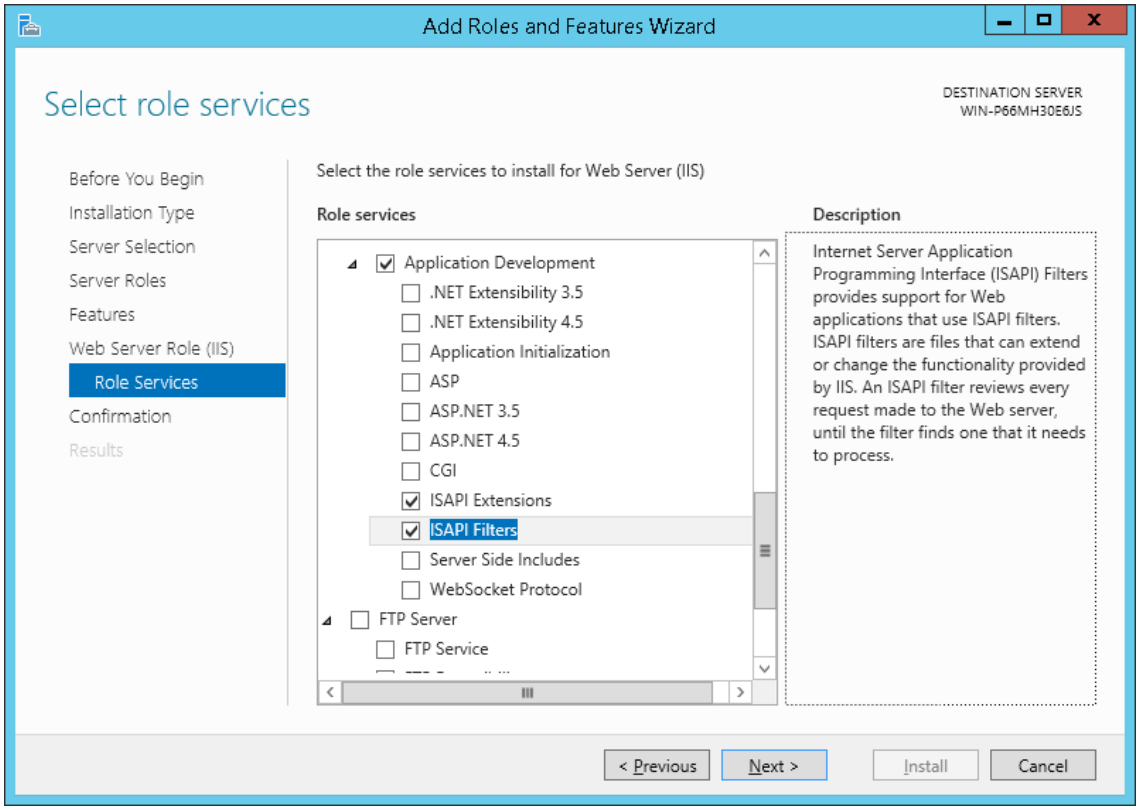
Install a webserver.

If using IIS, use the appropriate version for the platform.

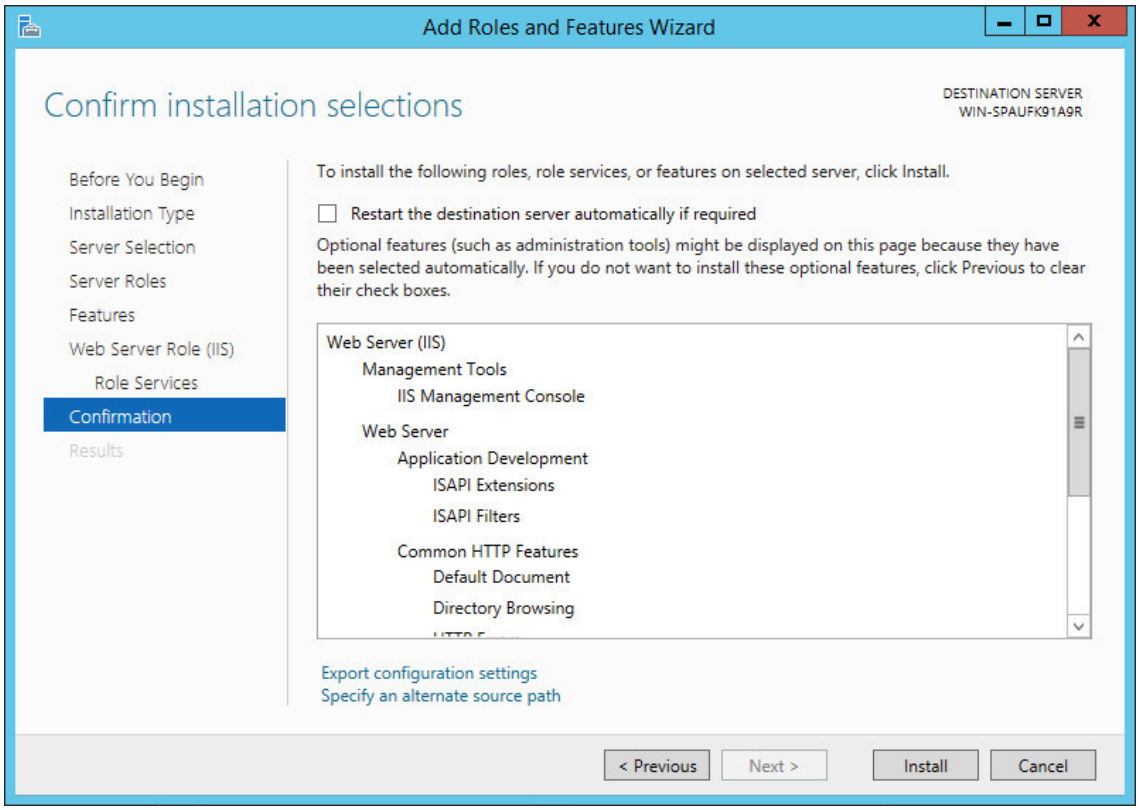
3.10.3 Installing ISAPI extensions and filters with IIS

ISAPI filters allow the gathering and manipulation of, and information about internet traffic on the system. The following are instructions to install ISAPI filters to an existing IIS installation. If IIS is not installed yet, the Web Server (IIS) should be installed as well.

1. Open the Server Manager.
2. In Server Manager, click the Manage menu, and then click Add Roles and Features.
3. In the Add Roles and Features wizard, click Next. Select the installation type and click Next. Select the destination server and click Next.
4. On the Server Roles page, expand Web Server (IIS), expand Web Server, expand Application Development, and then select ISAPI Filters and ISAPI Extensions. Click Next.



5. On the Select features page, click Next.



6. On the Confirm installation selections page, click Install.
7. On the Results page, click Close.

If using Apache, download the 2.2.x binary for Windows from apache.org, (choose HTTP Server from left side). Install normally, accepting all prompts. Retain only supports version 2.2.x.

Once Apache and a SQL server are installed, the system is prepared to install Retain's core components.

The install will complete the following tasks:

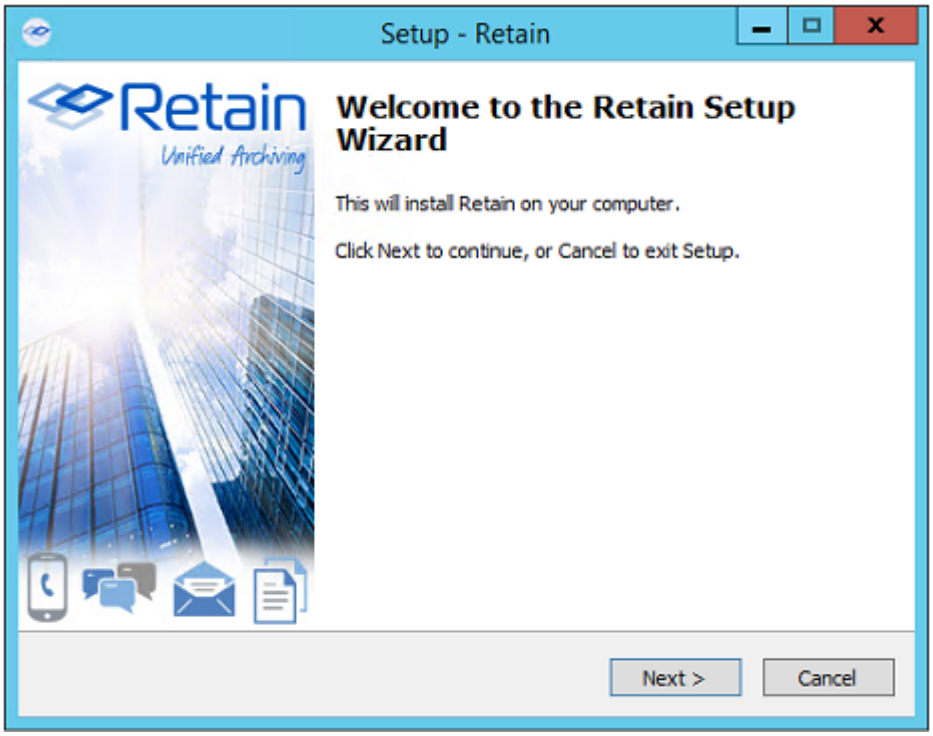
- Java Development Kit 1.8 (JDK) will be checked/installed.
- Retain Tomcat 7/8 and any existing Retain installation will be stopped.
- Retain will be installed.
- Retain Tomcat 8 is restarted and set to run as a system service.

3.10.4 Retain Install

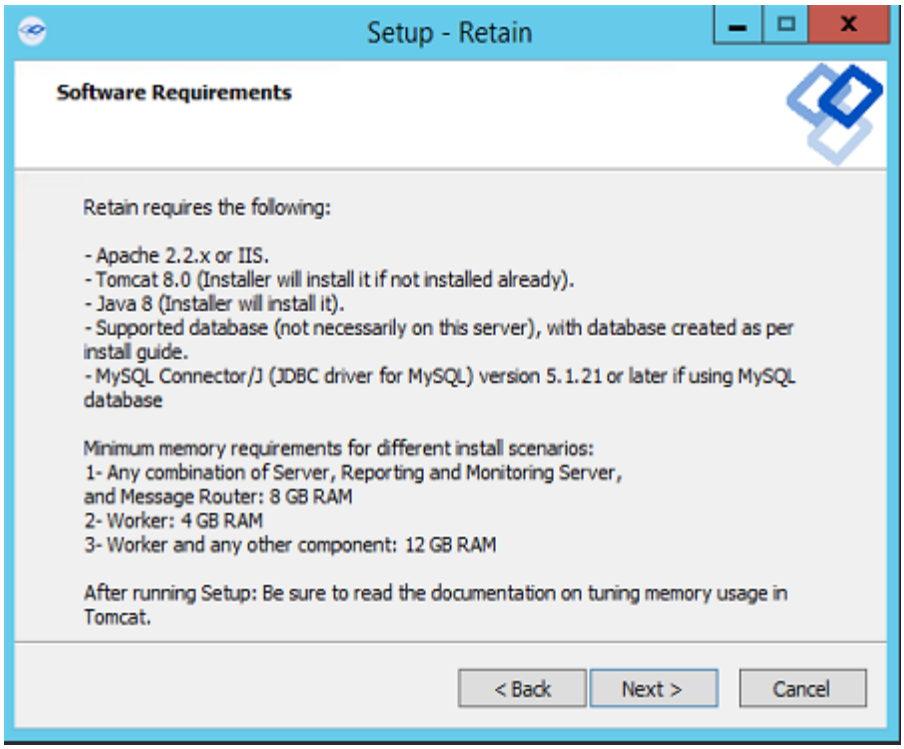
1. Download the latest version of Retain.
2. Extract the archive. This will extract into a new folder.

   Right-click and select "Extract All..."
3. Unzip the downloaded Retain install package to a work directory. Almost any name will do, as long as you remember it.

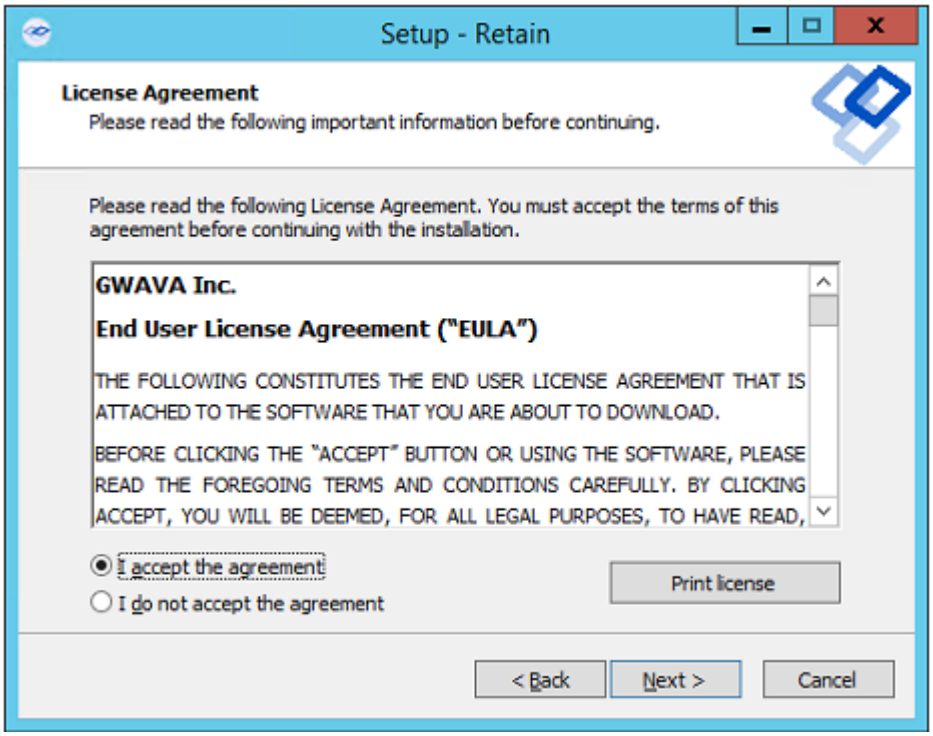
4. Run “RetainInstall.exe”. NOTE: If it already exists, Tomcat will be automatically shut down by the installer.
5. The Retain Setup Wizard welcomes you to the installer.



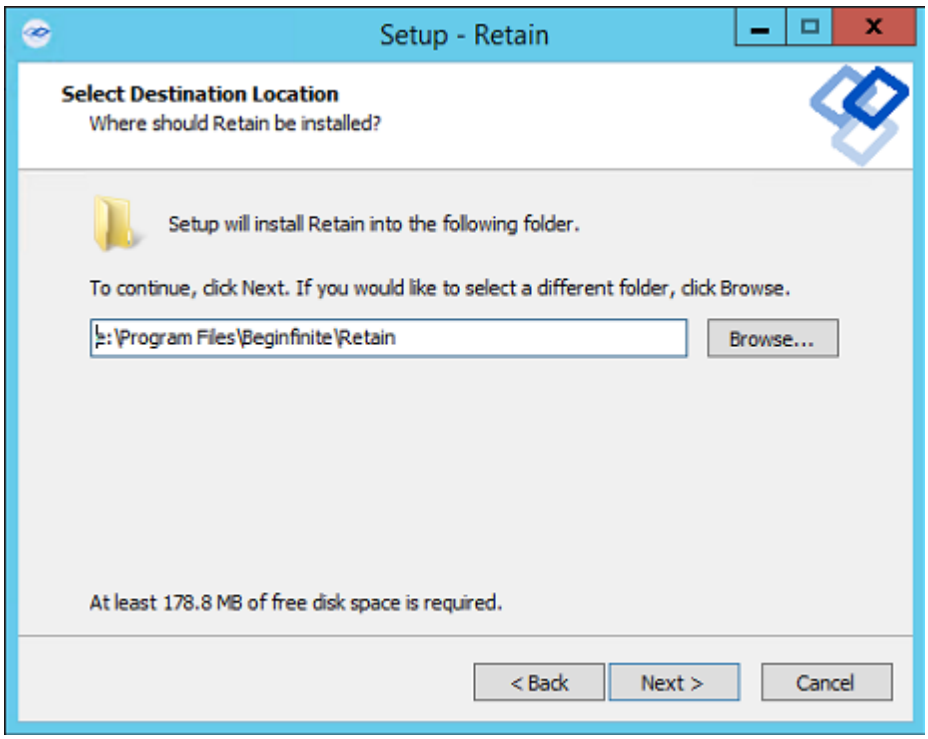
6. The Wizard reminds you of the Software Requirements. Click ‘Next’ to continue.



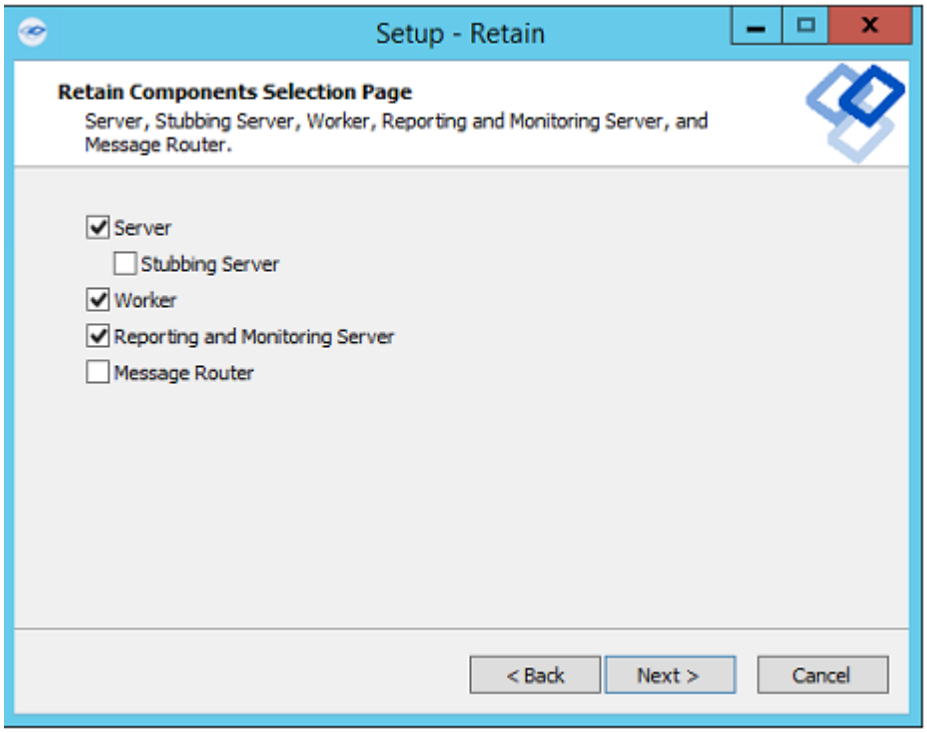
7. Agree to the license agreement.



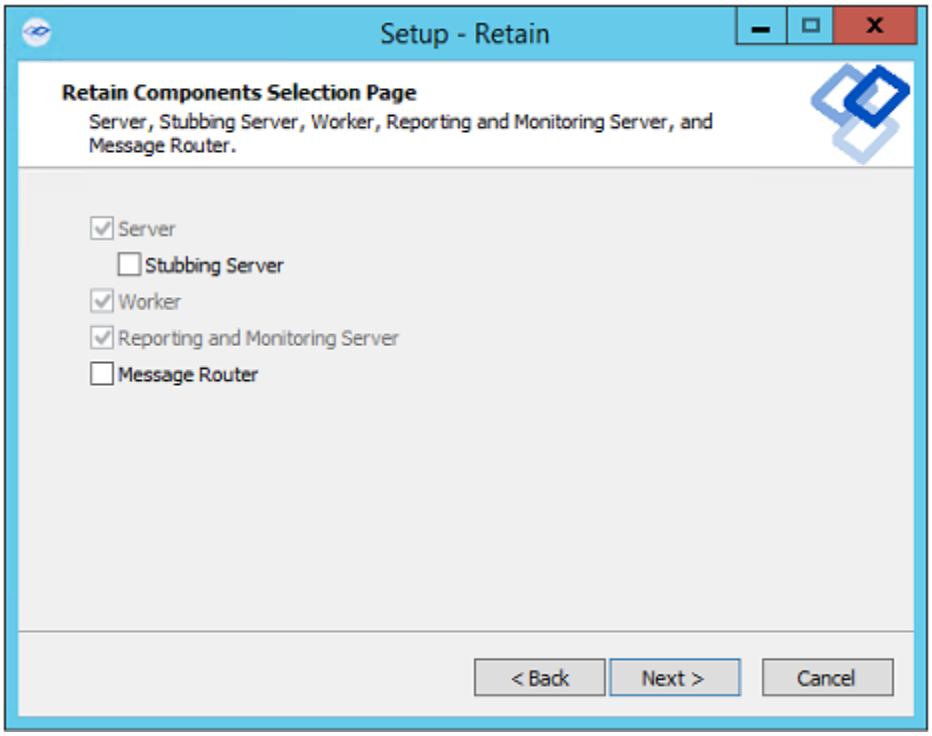
8. Select the install location.



9. Select the program features to install. (Retain Server, Stubbing Server, Retain Worker, Reporting and Monitoring Server, and the Message Router.)



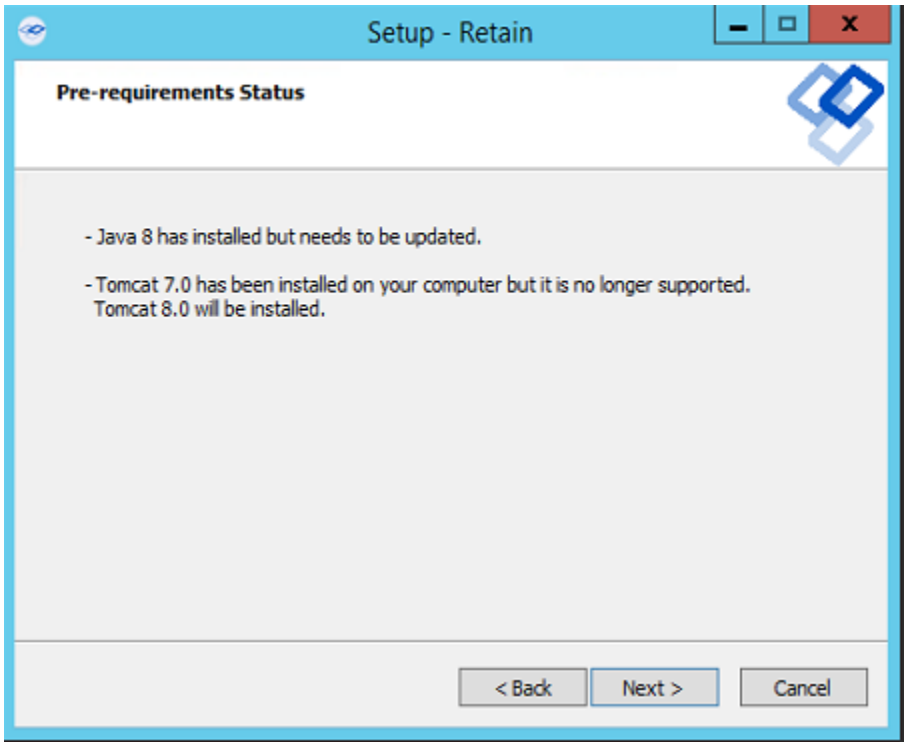
9a. If upgrading, previously installed components will be preselected and you will be able to select additional modules to install.



10. Pre-requirements status is displayed. Installer attempts to detect previous Tomcat installations.
    If none are found, it asks whether Tomcat has been installed. Indicate ‘yes’ or ‘no’.
    If upgrading, Tomcat 7.0 will be upgraded to Tomcat 8.0.

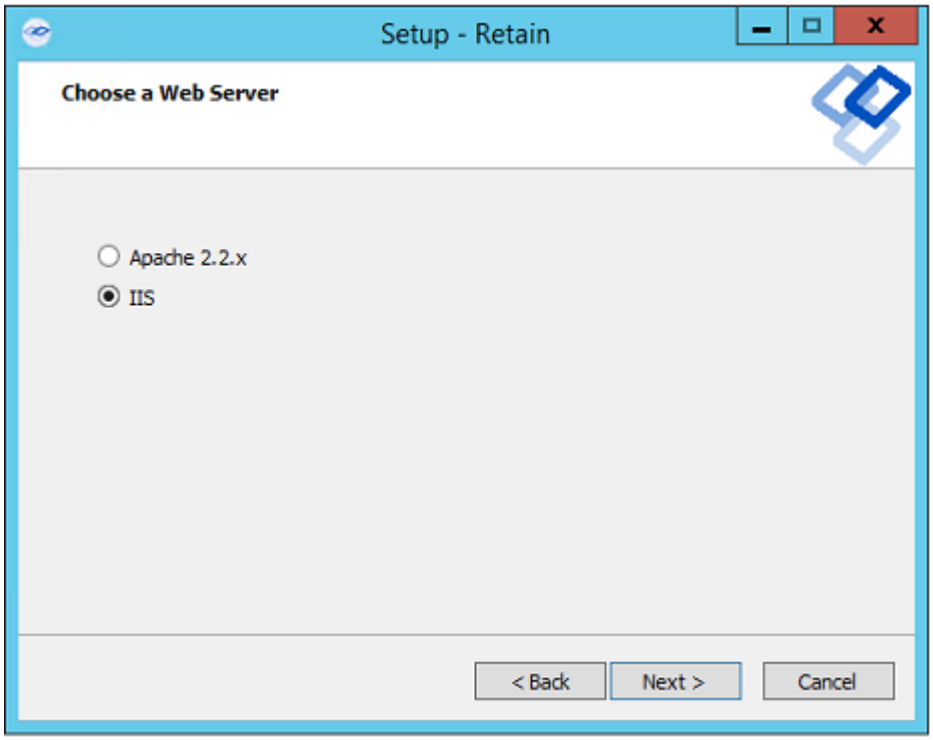
Please note that the environment variable ‘RETAIN_CATALINA_HOME’ is set after Tomcat installation.

Click ‘Next’ to continue.

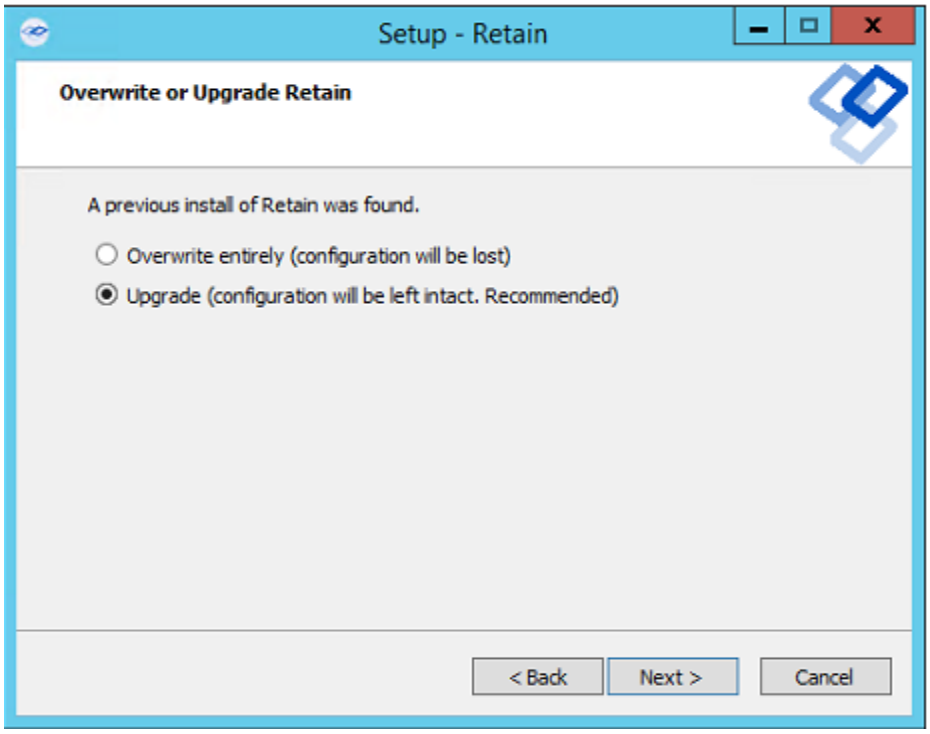


Java and Tomcat are installed or upgraded. (Note: If there are issues with Apache and Tomcat starting correctly, set the Java initial memory pool and the maximum memory pool to the same value, appropriate to the system and expected load.) Tomcat is run as a System Service and shows up in the system services list as “Retain Tomcat 8”. (Start | Run | services.msc)

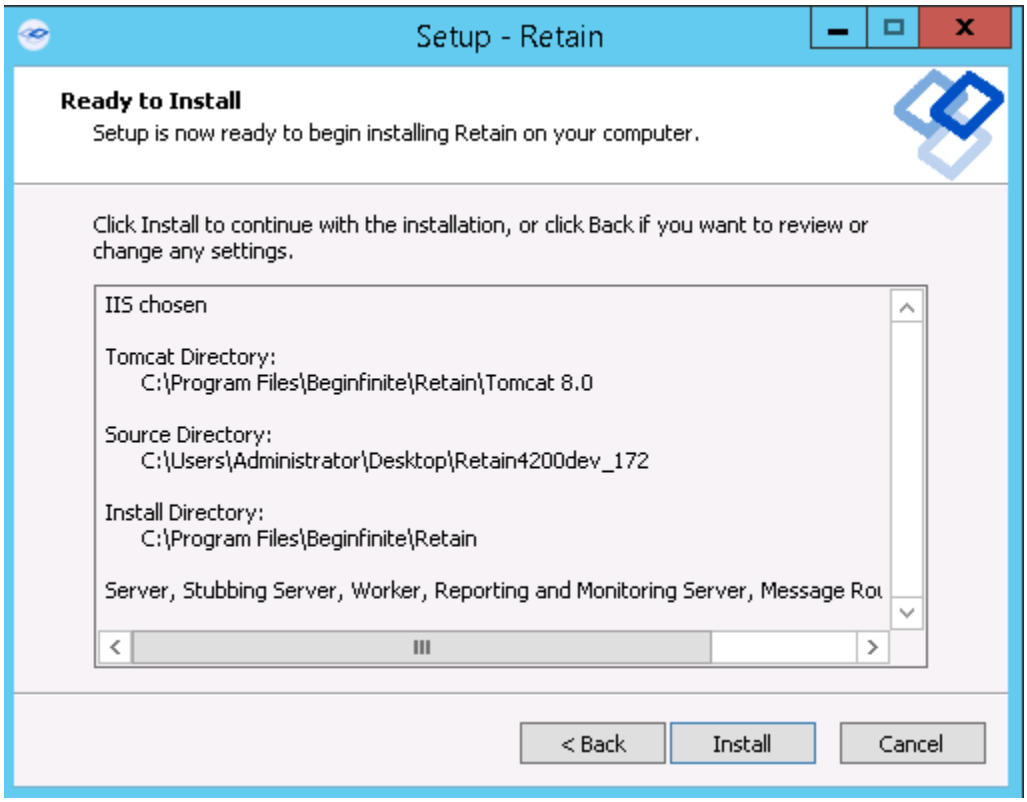
11. Select either Apache or IIS. (If apache, provide the installation directory.)



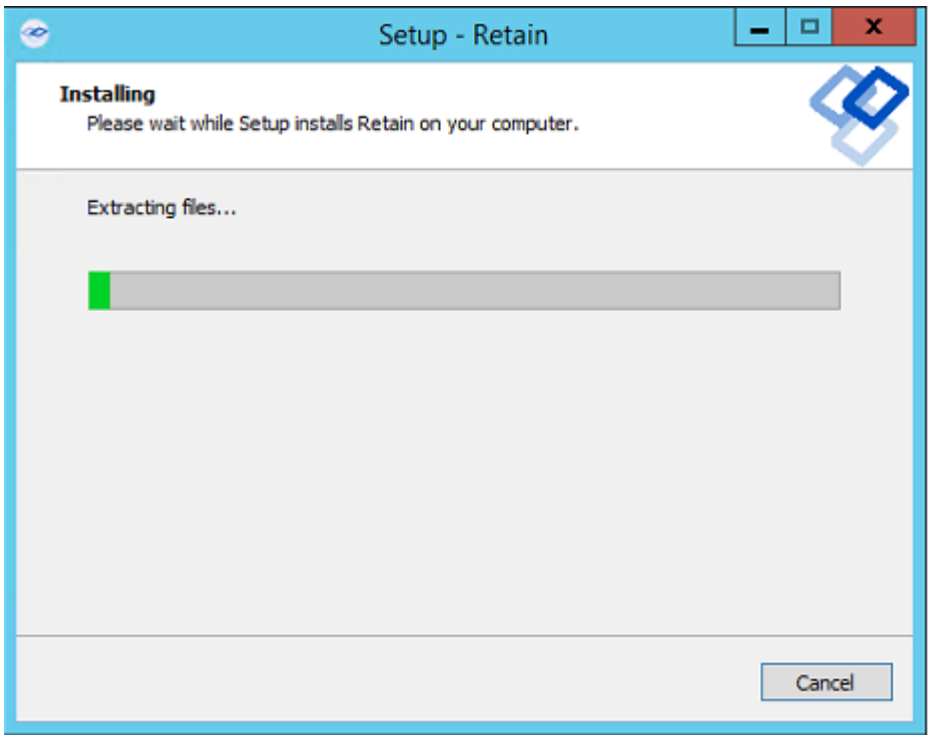
12. If a previous version of Retain is detected, upgrade options are displayed; Overwrite or Upgrade. Generally, you will want to choose Upgrade. Click Next.



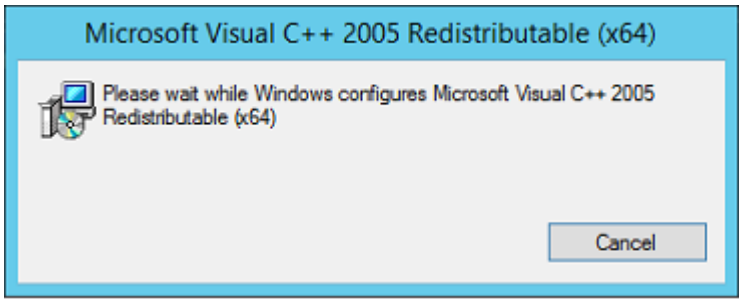
13. Confirm settings and select 'Install'.



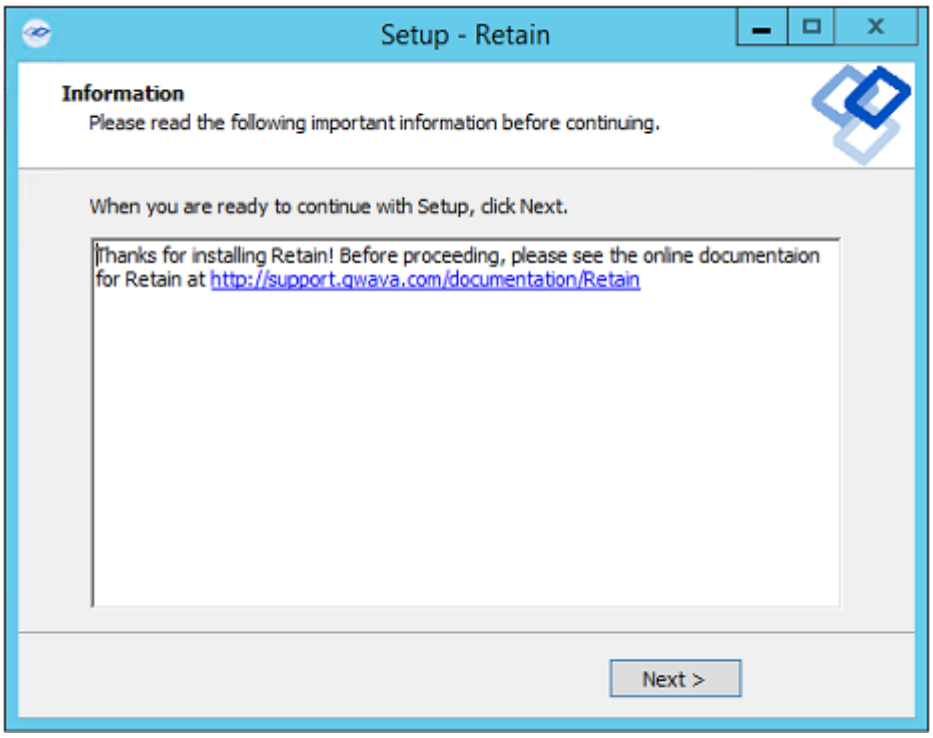
14. The product will install.



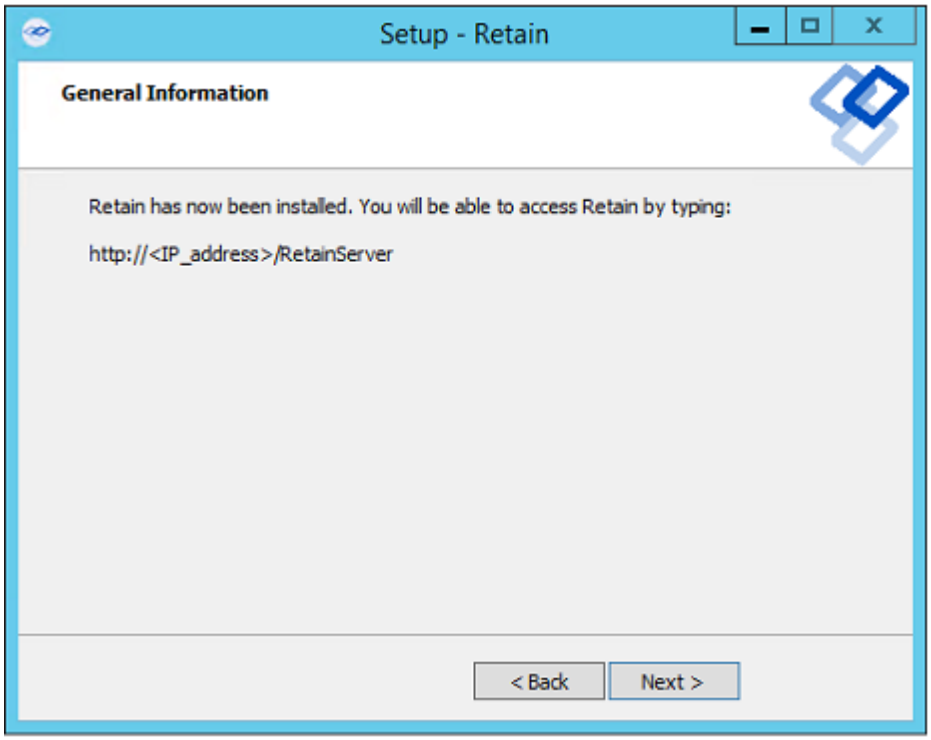
15. If Microsoft C++ 2005 Runtime distribution is not installed, the appropriate version will automatically be installed.



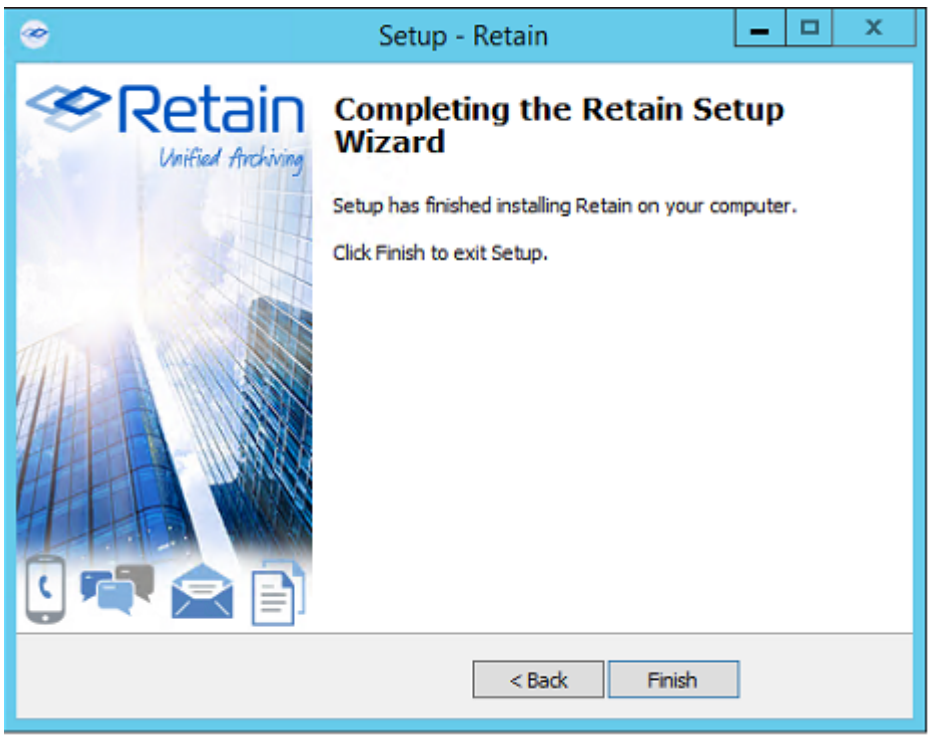
16. The Setup Wizard thanks you for installing Retain.



17. The Setup Wizard shows you the URL to the Retain admin console.



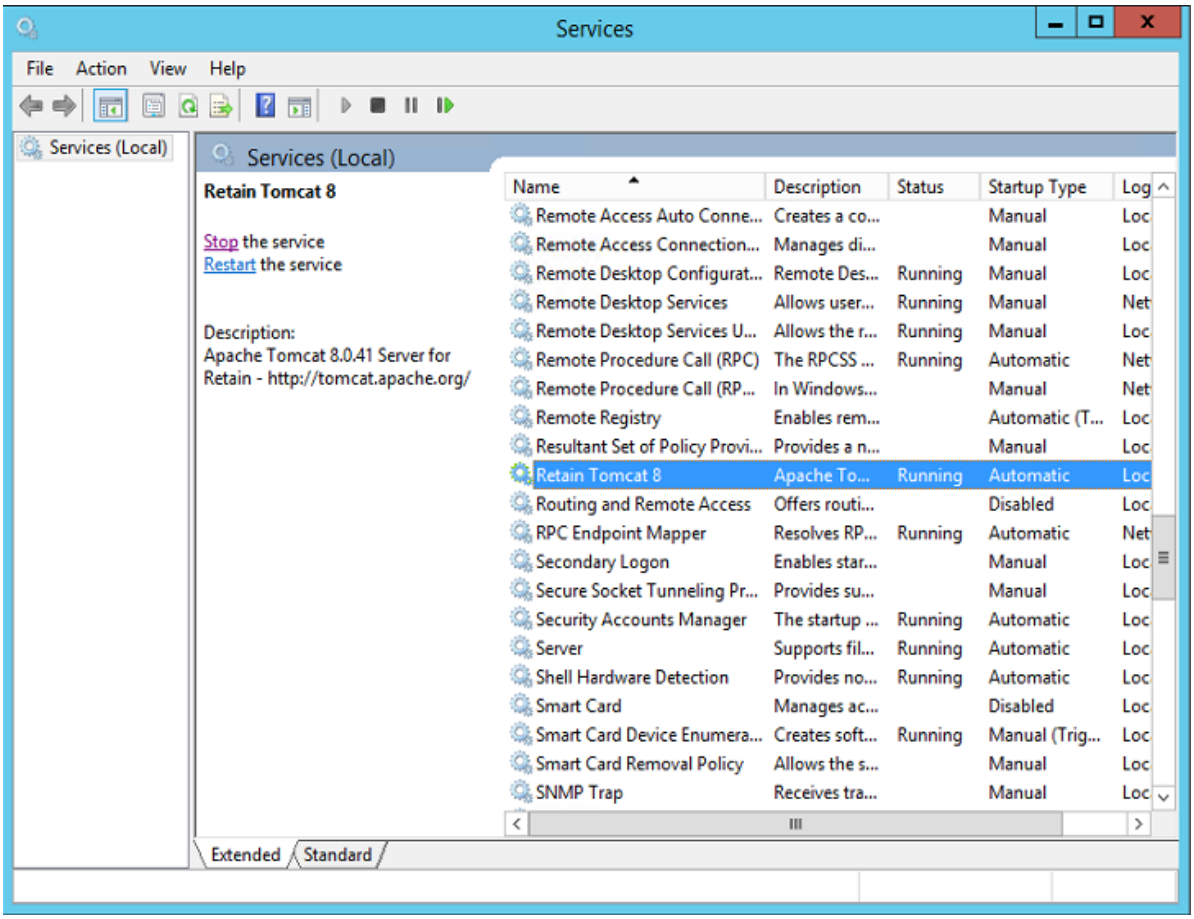
18. The Retain installation process is complete. Click Finish.



19. If using MySQL, stop the tomcat service and install the MySQL Connector/J (JDBC driver for MySQL), then start tomcat.
20. Browse to the Retain web console URL and complete the initial setup.

3.10.5 Controlling Tomcat

Tomcat may be started, stopped or restarted from the Services program.



3.10.6 Installing Just a Worker to a Server

The Worker may be installed on a server other than the Retain server itself. For example, the worker may be installed on the email server or on a dedicated worker server.

This is much simpler than installing the entire Retain server system. It only requires installing the webserver and the worker agent software. It uses the same installer as above.

Prerequisite:

1. Install the web server (IIS with ISAPI filter, as above).

Pre-install tasks:

1. Download the latest version of Retain.
2. Extract the archive. This will extract into a new folder.

   In the GUI right-click and select "Extract All..."
3. Run the installer, *RetainInstall.exe*, as above.
4. Select *only* "Worker" as the component to install.

5. Finish the install as normal.
6. Workers are accessed by browsing to *http://[IPAddress or HostName]/RetainWorker*

3.10.7 Installing Additional Workers

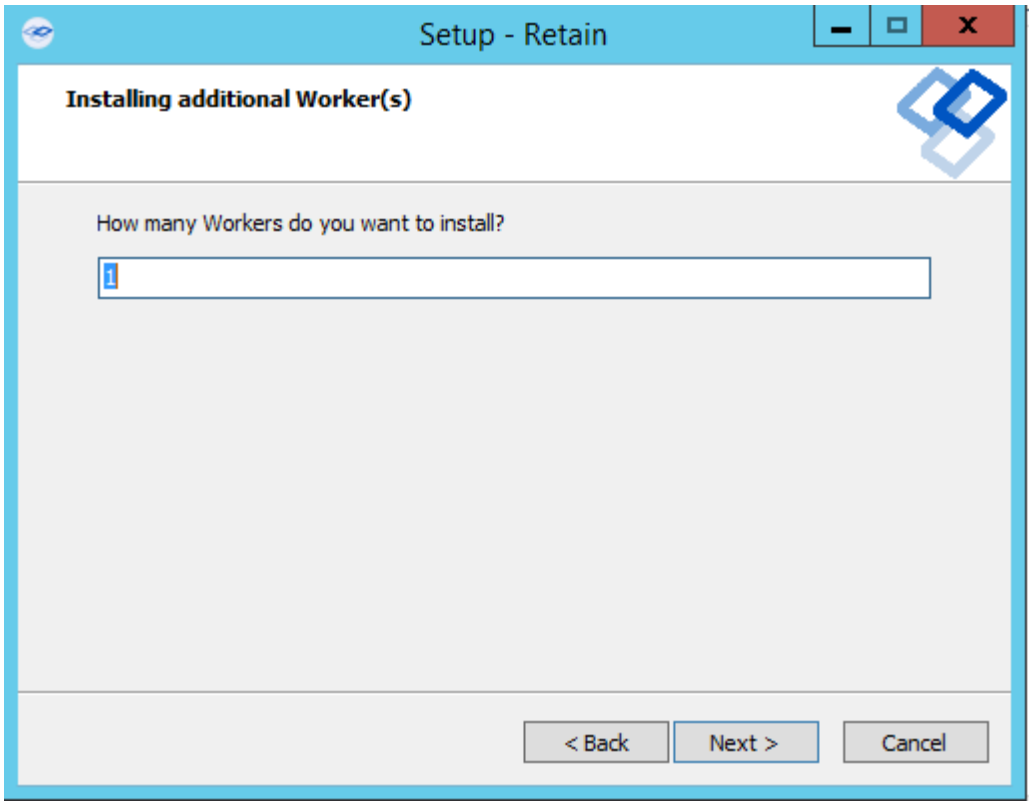
Overview

In larger email systems, it can take a long time to dredge all the users especially if the users are on multiple mail servers. You can install multiple workers that can run in parallel to speed up the archiving process. We generally recommend one worker per mail server. The Retain installer can install up to 10 workers on a single server. We recommend an additional 1GB for RAM per worker. Workers are accessed by browsing to http://[IPAddress or HostName]/RetainWorkerN

The Retain installer allows you to install up to nine workers at once.

Procedure

1. Run the RetainInstall.exe and install at least one Worker or upgrade the existing installation, if you haven't already done so.
2. From the command line (run as Administrator), type RetainInstall.exe /AddWorker to start the installer in Add Worker Mode.
3. Acknowledge requirements.
4. Accept the license agreement.
5. Confirm the install location of Retain.
6. Enter the number (1-9) of Worker agents to install. They will be installed in numerical order (e.g. RetainWorker, RetainWorker1, RetainWorker2..., RetainWorker9).



7. Confirm which webserver is installed.
8. Install the Worker agent software. After installing and acknowledging the location of the documentation the installer will show you how to access the Workers that were just installed.
9. Log into the Retain web console as admin
10. Create a new Worker and download the Bootstrap for each new Worker you installed. You may save it with a different name to organize it.
11. Browse to the new Worker at *http:///[ServerAddress]/RetainWorkerN* where N is the number of the worker you just installed.
12. Upload the bootstrap.
13. Click Continue and login to the worker.
14. Check the friendly name that it is the one you expect.

There is a separate log file for each worker, found in *C:\Program Files\Beginfinite\Retain\Tomcat 8.0\logs\RetainWorkerN.yyyy-mm-dd.log* unless otherwise specified in Configure Tomcat.

3.11 Post-Install Tasks

After running through the initial install guide for your OS, you are ready to configure Retain. From this point on, Retain works the same regardless which platform it is running on. To get Retain ready for regu-lar use, some final one-time procedures need to be run.

- Retain Server needs to be configured
- Configure the relevant module(s)
- At least one Retain Worker needs to be configured

3.11.1 If using MySQL

- Install Connector/J

3.11.2 Initial Setup Overview

- Choose your storage path
- Retain prompts you for access information to your SQL Database Server
  - You must create this database beforehand. (See system requirements for documen-tation for the supported SQL servers.)
  - Retain creates the tables, indexes, and prepares the database for use

You will be prompted to provide SQL connection information for both the configuration and message store databases. For convenience and simplicity most installations store this information in the same database. The database(s) will be used to store Retain configuration, cached address books, and all saved message data.

- Create at least one administrator account. The first account will be the default admin account. It is created during the initial setup
  - The ‘admin’ account is special:

    -It has ALL admin rights

    -It never expires

    -It always uses “offline authentication” (authenticates using a password stored in the SQL configuration database instead of authenticating against a message system)
  - Additional accounts can be created as desired with as many rights as you wish in the management interface. User accounts may possess a mixture of administrator level and user level rights. As a result, users may have some admin level rights, but would not be considered an administrator
- Retain prompts you for the following SMTP information for notification and forwarding. (If no SMTP is present, this step may be skipped by typing the word ‘skip’ in the SMTP Mail Server field.)
  - SMTP Server: Server to use for sending messages. (DNS or IP) (use DNS whenever possible)
  - From address: Enter the address that you want to appear in the “From” field on the messages Retain sends
  - To address: Enter the address to which Retain will send notification
  - Username: Enter the username to use for authentication with the SMTP server for relaying mail.
  - Password: Enter the password for this username
- Select the indexing engine. There are two different indexing engines available for use with Retain: The Standard Index engine and the High Availability Index (HAI). While essentially the same engine, the High Availability Index is scaled and designed for high availability and rapid searches on very large systems. The HAI requires at least 3 separate, dedicated, index servers in addition to the Retain Server. The Standard Index is embedded with the Retain Server and is more than sufficient for small and medium systems.
  - After initial setup, configuration of a module and creation of a profile, schedule, worker, and job are required before messages can be archived.

Continue with Retain Initial Setup

3.12 Initial Configuration

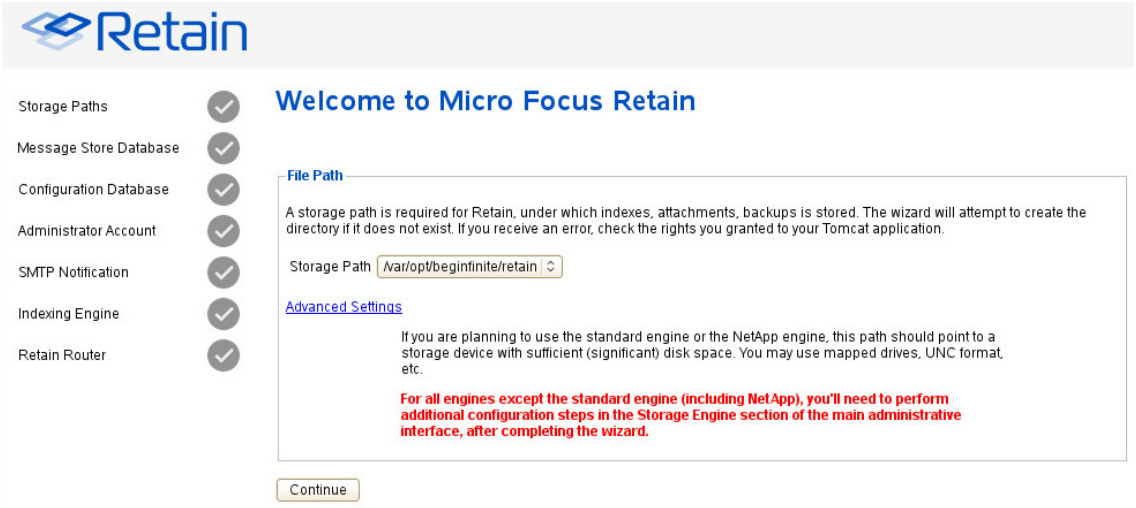
3.12.1 Moving Retain

If you are moving Retain to another server and you are connecting to an existing Retain database on an external database server. Do NOT go through the configuration wizard, or data will be deleted from the database. You MUST transfer the ASConfig.cfg from the previous server.

3.12.2 Run the Configuration Wizard

1. Open a web browser
2. Go to http://<your-RetainServer-ip>/RetainServer
3. Follow the wizard. You will see this screen initially:

   Please pay special attention to the Storage Path and the Advanced Settings.



3.12.3 About Storage Paths

Retain was designed to be very flexible in allowing you to choose where the data will be stored. Choosing the right hardware and software storage combination is crucial. Storage itself is a function of the operat-ing system and hardware you choose.

As storage needs approach enterprise levels, the need to specify your data storage locations becomes more critical. Some hardware and operating systems have a hard limit of 2 Terabytes, for instance. In sit-uations like these, a SAN might be a good choice for consideration. Either way, it's imperative that you take the time to investigate your storage needs BEFORE you begin. Retain will monitor the storage space available, and will issue email warnings at 20 GB, and will lock down Retain to a maintenance mode when the storage space drops to 10 GB. This is to avoid possible corruption to the storage and index which may occur when space runs out.

When choosing your storage system, bear in mind the following considerations:

- Storage size limits. Will storage needs ever exceed this?
- Ability to add additional storage capacity
- OS Version and kernel to support the storage system of your choice

- File system and partitioning
- Is a SAN right for you?
- Speed and reliability
- Backup/restore is your responsibility

3.12.4 Standard Storage Path

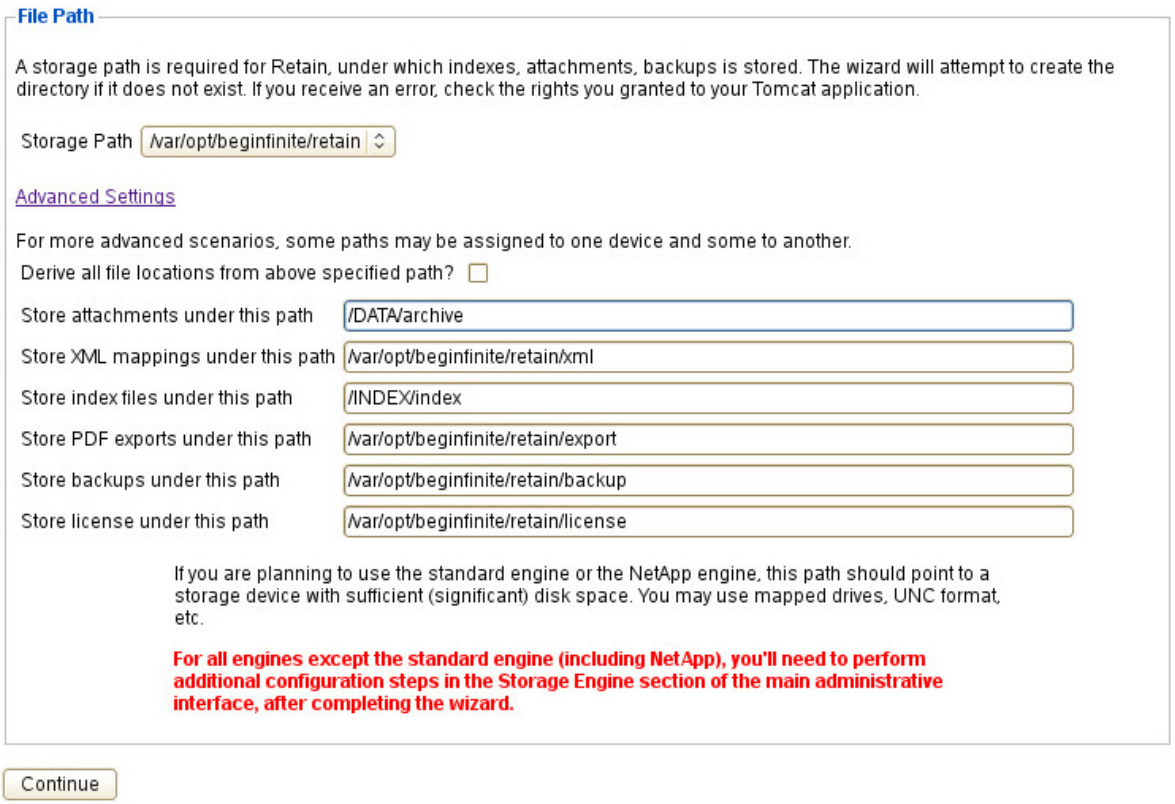
You are now ready to start your initial Retain configuration. Specify a storage path for Retain. Accept the default path unless you have a reason to choose a different one. See the next page for an important note about storage paths. If you are specifying a custom storage path, Retain must have rights to the location. For Linux instructions see Custom Storage Path Rights, at the end of the Linux install guide.

3.12.5 Advanced Storage Options

For this reason, Retain offers the advanced storage options. (Located under Server Configuration.) You can specify where you want Retain to store its information.

By default, all data is stored under one master path, specified in the Storage section. From the Retain Server main screen, you can also find this under Server Configuration under the Storage tab. Significant performance increases have been seen when data is stored on separate volumes.

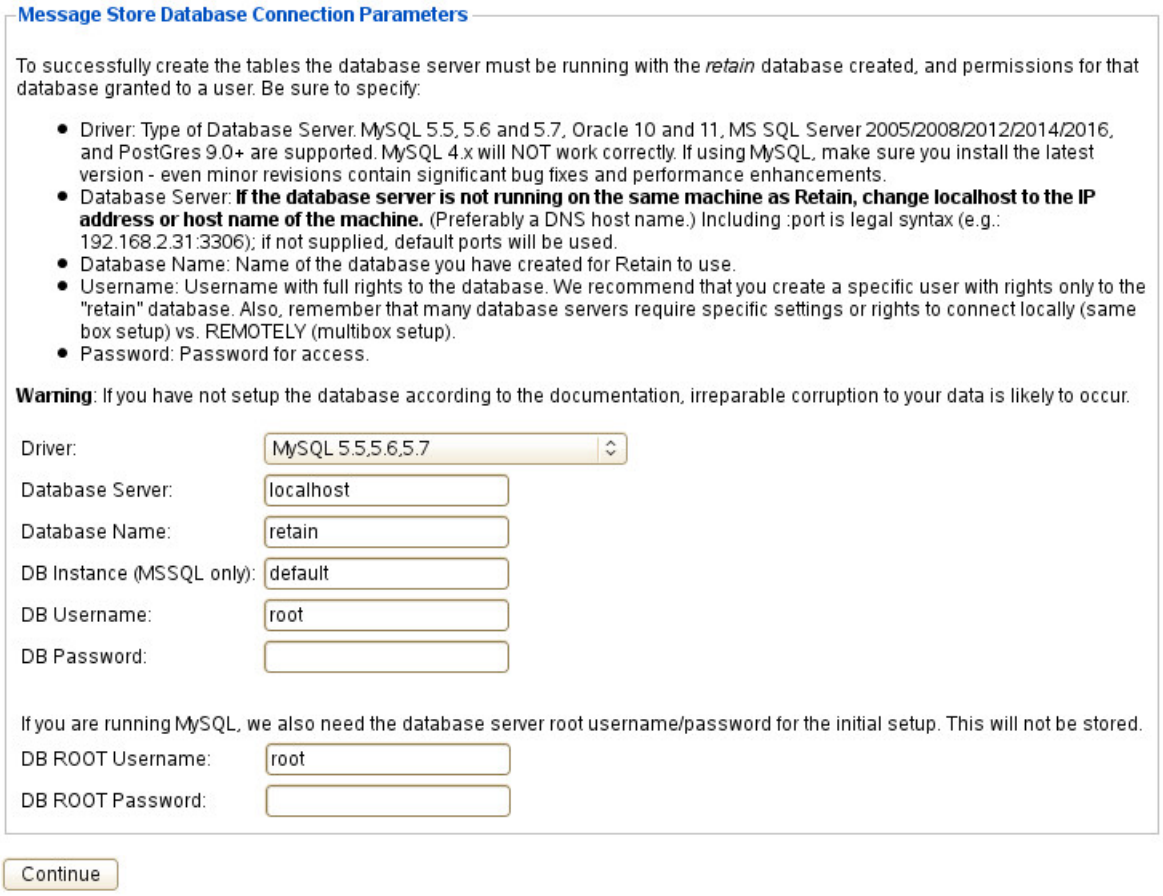
If you wish to specify your storage locations in more detail, click on the “Advanced Settings,” and disable "Derive all file locations from above specified path?".



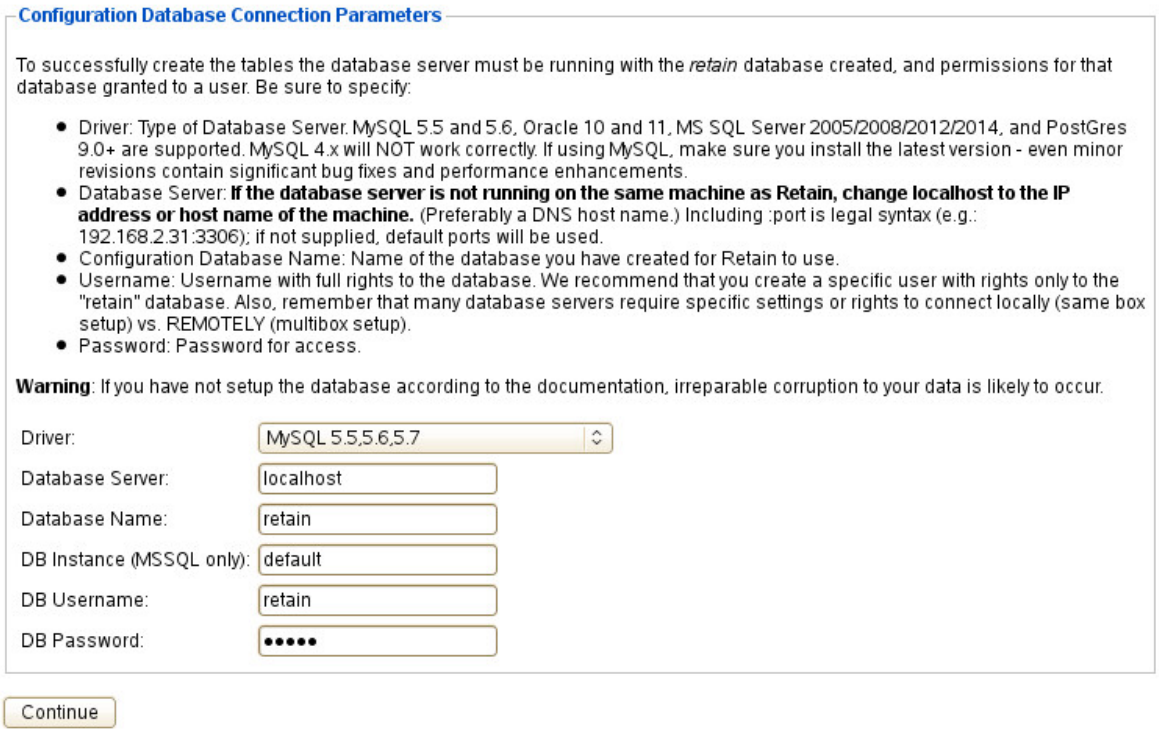
These fields are automatically populated if you only specify the main storage path. You may change them to point anywhere you wish. You may also change the storage locations later due to storage space needs. Please see the Partitions section for details. You also specify to encrypt the data as it's stored. We strongly recommend for performance and recoverability that you do NOT encrypt using Retain but con-sider Block level encryption options (such as BitLocker) that have escrow keys and are processed directly at the operating system level.

Click "Continue"

4. Next, the screen appears for entering the database connection parameters. This is the same information from the database created during the preparation phase. The connection informa-tion must be supplied for both the storage and Configuration databases, though they can be housed in the same database

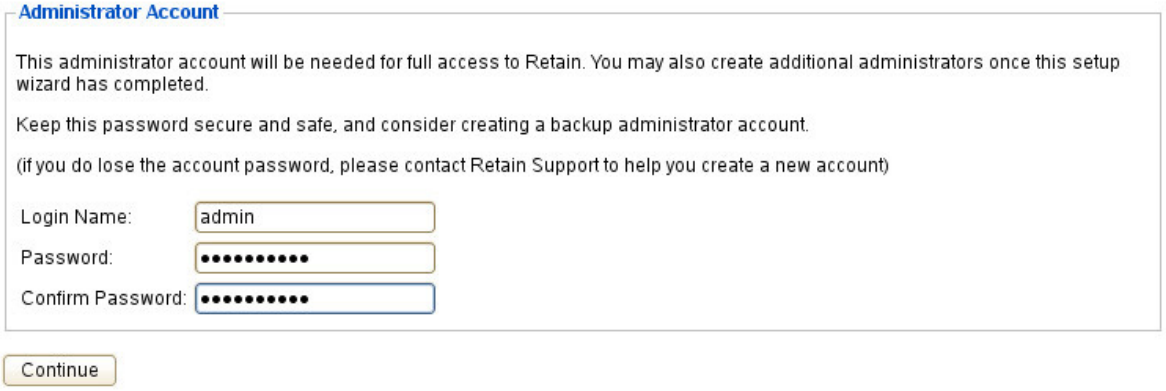


5. The configuration database must also be specified. Unless otherwise required, the settings here should be identical to the Retain storage database, as they are usually housed together.

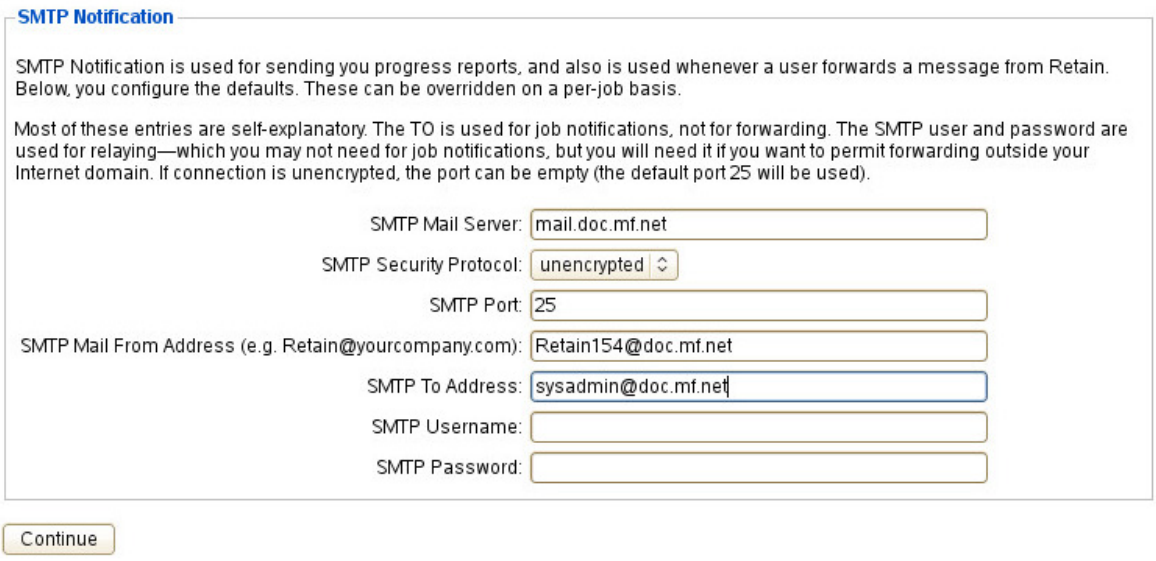


6. Create the administrator account. This account is special in that this account is the administrator with rights to everything, by default. While other accounts may be created, this administrator account should never be removed or changed. This account may be used by the index engine or other components. Remember the credentials created here; they will be used to log into the Retain system.

NOTE: Retain allows the creation of many user accounts that share administrative roles. The initial account created is not particularly special; it's merely a convenience to create one super user with all the Administrative rights granted. Do not delete this account. If the admin account is deleted, please con-tact support.

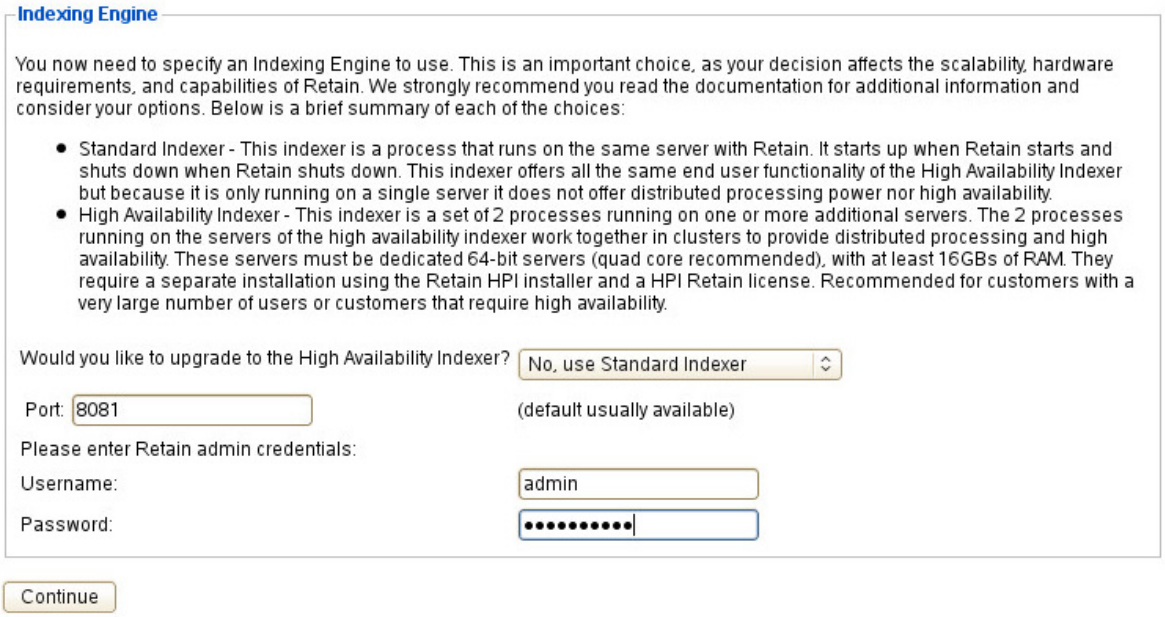


7. Enter the information for an SMTP relay host that Retain will use to send mail. Setup will attempt to send a test message. You cannot proceed with installation until the SMTP test message is suc-cessfully delivered. If you see an error, check firewall settings, (including local desktop/server firewalls such as McAfee, Symantec, and SUSE)

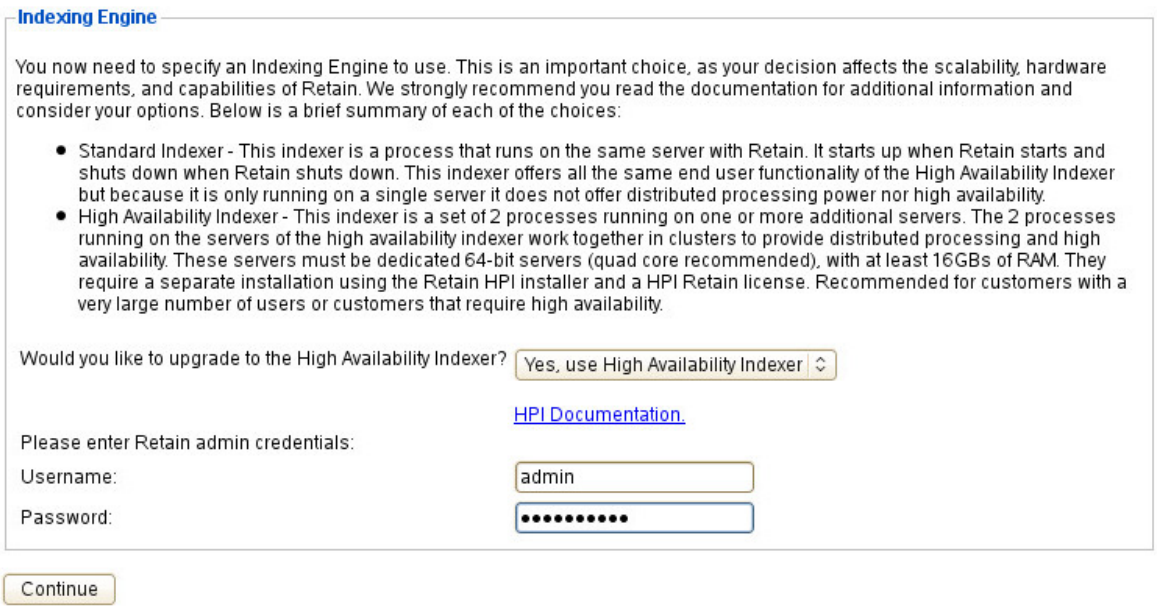


8. Index Engine Selection

8a. *Standard* is the High Performance index engine that comes standard with Retain. This engine will handle all but the largest system indexing needs, is housed in the Retain Server, and requires no further configuration.

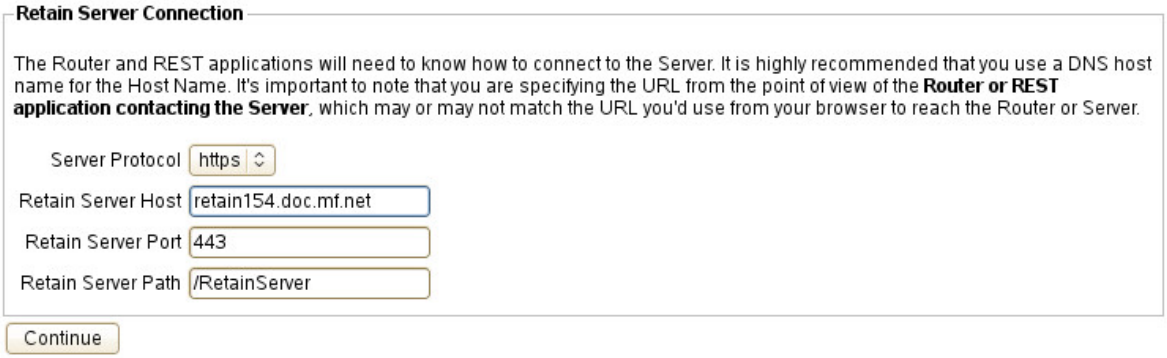


8b. *High Availability Index* engine. This index engine is designed for high load systems with very large data-bases. The HAI requires a separate license, several separate installations and a network plan created in coordination with Support. If you wish to use HAI, it should already be installed and running.

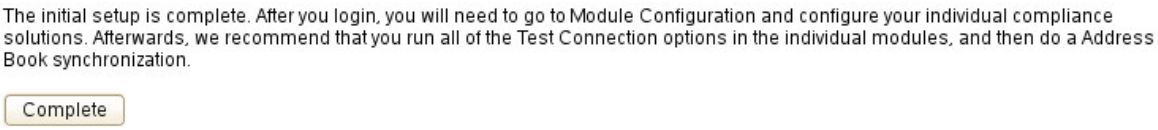


8c. Enter the Retain administrator account name and password for the index engine connection. Once the administrator account has been entered, the system will initialize and connect to the new index engine. This will take a few moments.

9. Setup the connection for the Router and REST application, such as Retain Publisher. Use a hostname that is set up in your DNS that external devices can resolve to.



10. Initial setup is now complete.



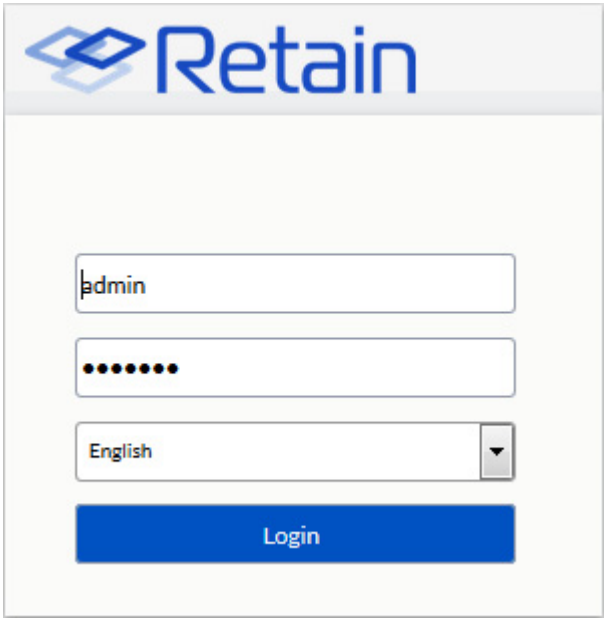
Select the 'Complete' button to be presented with the login prompt.

3.13 Front Page and First Login

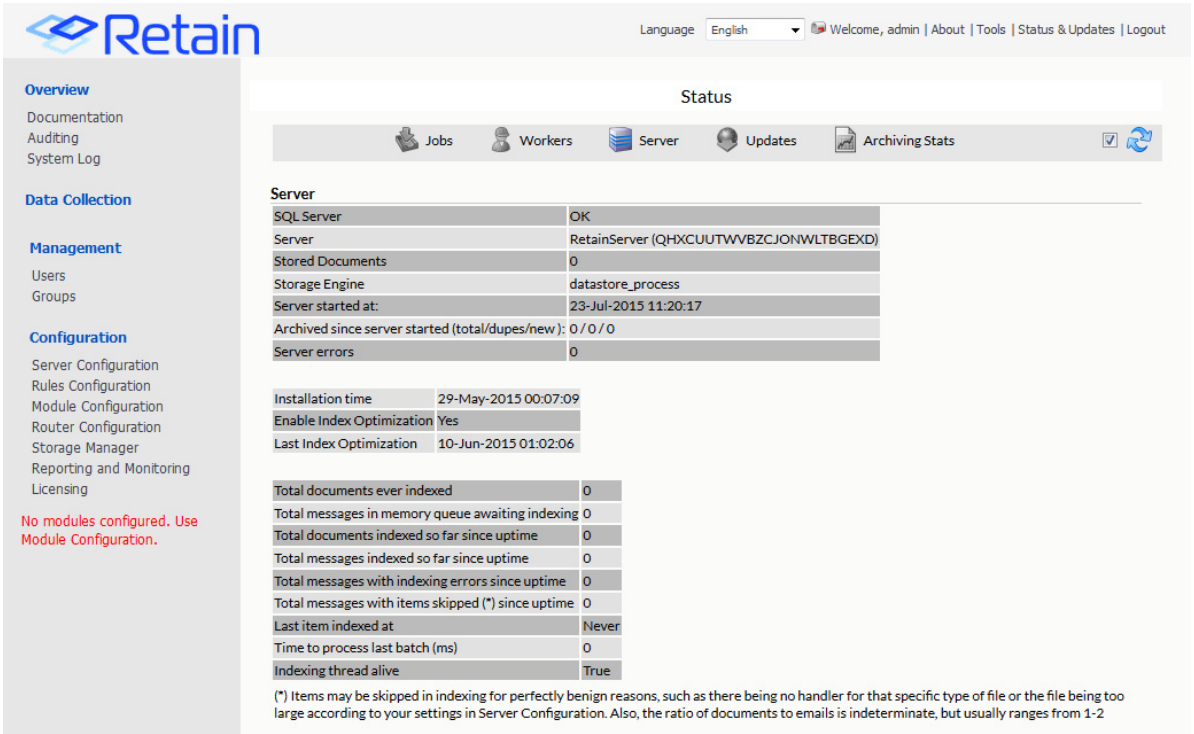
3.13.1 Logging In

From your web browser, type: http://<YourServerIP>/RetainServer

You'll be brought to the login screen:



After you log in, you will see the main administration page: (Non-administrators will only see the Search Interface.)



An informational screen is displayed, showing among other things the latest job activity, server health information, and software updates that are available.

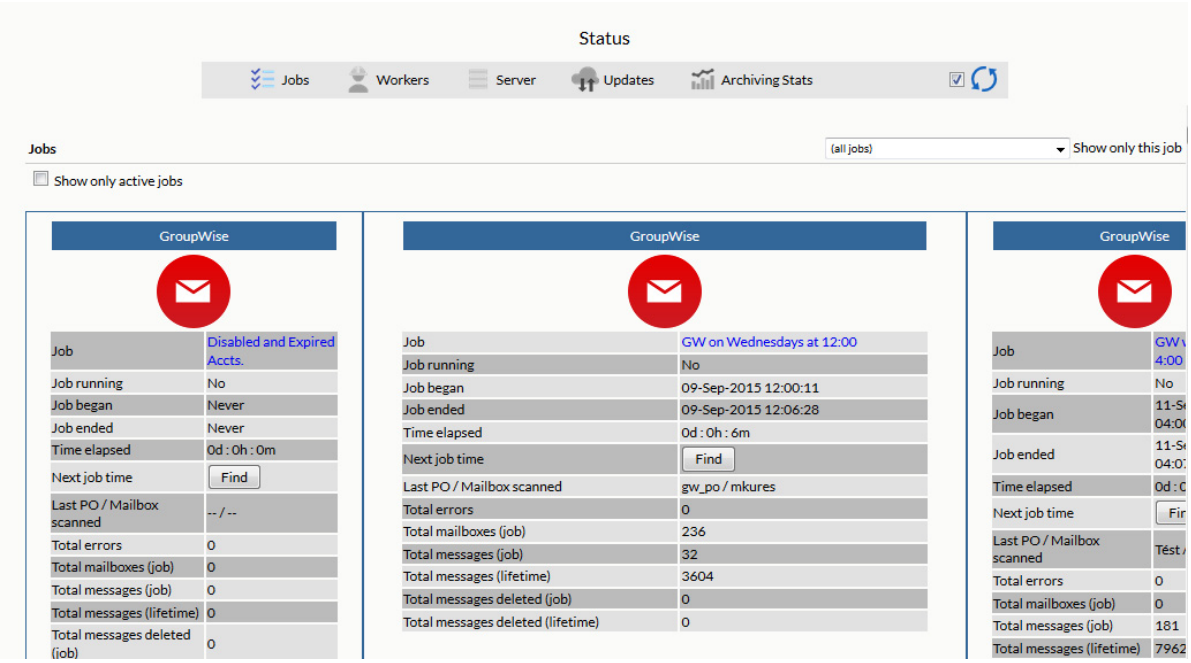
NOTE: Only users possessing at least one administrative role will see the administrator's screen login. Other users will be forwarded to the Search Interface.

3.13.2 Status and Updates

The status page has four different tabs which show the status of your system. The status page is set to default to the tab that is deemed to require the most attention, or is deemed to hold the most important information at a given time. For instance, if an error has occurred, that tab will be shown, likewise the update tab will be shown by default if there is an update available. The tabs are discussed below.

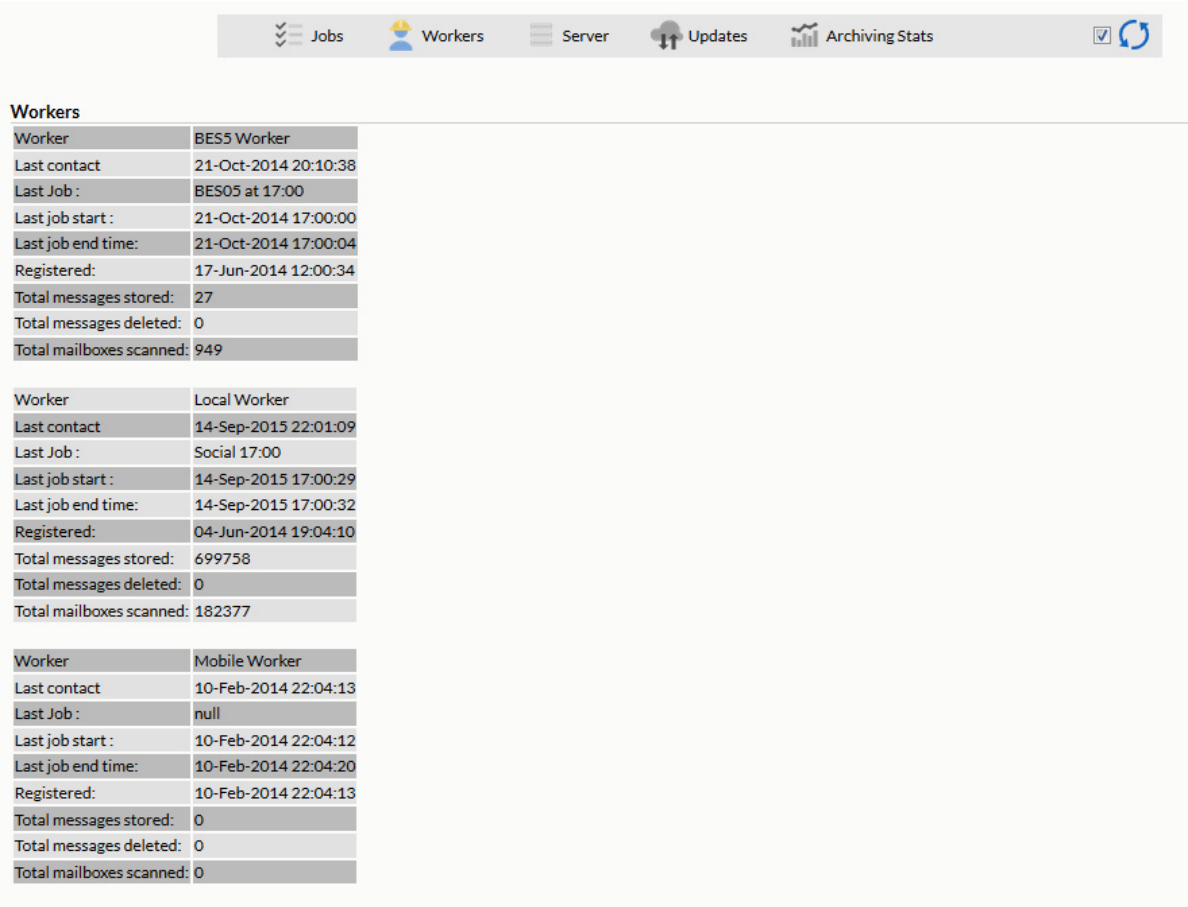
Jobs Tab

The Jobs tab, shown above, shows the jobs that are assigned to the server and their status. You also can select the "show only active jobs" check box, and only jobs that are currently running will be displayed.



Worker Tab

The workers tab displays the workers and their last active jobs, connections, and totals.

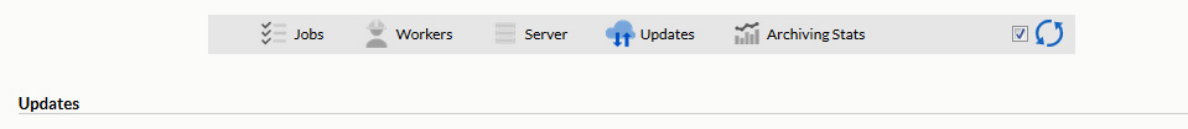


Server Tab

The Server tab shows the basic status of the Server with the messages archived and maintenance infor-mation.

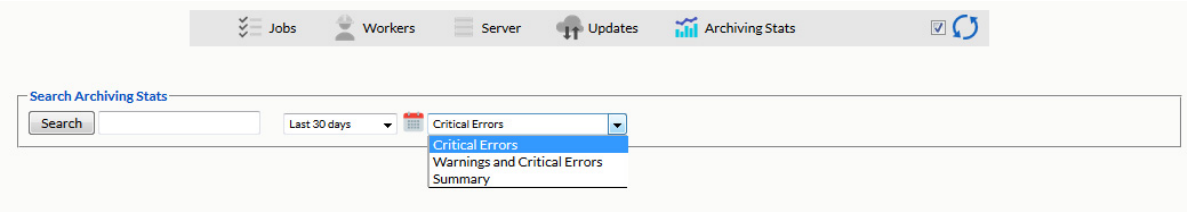
Updates Tab

The updates tab displays whether there is a later version of Retain available. This tab will be displayed, by default, only if there is a new update available.



Archive Stats Tab

The archive stats tab is a small statistics portal which allows rapid searching of the archive statistics in the server. The search allows for a time frame, and the categories shown below: Critical Errors, Warnings and Critical Errors, and Summary. Once the ‘Search’ button has been pressed, a new window or tab is opened to display the search results. This function had been deprecated as it has been replaced by Reporting and Monitoring.



3.13.3 Saving Changes

Retain uses a web interface for all administration, and is unaware of changes made on the webpage until the information is returned to the Server. To make sure you keep all changes you have made in the inter-face, select the “Save Changes” button to commit the change to the Server. When you make a change, the website is aware and you will see the ‘Save Changes’ button become colored at the top right of the screen:



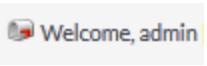
When the “Save Changes” lights up, click on the disk icon to save any changes you have made on this screen otherwise your changes will be lost.

The top right bar also contains a button to reload the current page and a help button which opens a win-dow displaying critical information.

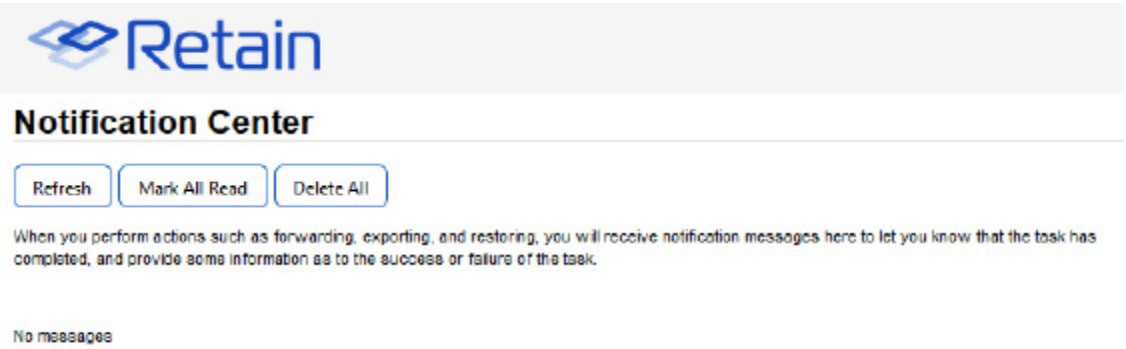
3.13.4 Notifications

Notifications are automatically sent to the administrator whenever data in Retain is moved. Moving data consists of system migration, exporting, forwarding, restoring, and publishing. By default, the address used is the address supplied during severs activation.

To spawn the notifications page, click on the ‘Welcome’ link at the top of the page. When notifications are pending, the red flag on the mailbox is lifted.



Selecting the mailbox or welcome link opens a new notifications window.



The highlighted notifications are new, unread notifications and events. Selecting the event notification provides all the pertinent information regarding the event.

If there are any errors, a truncated error log is provided which details the error and the reason for the error.

3.14 Licensing

The Licensing page displays, manages, and reports on the licenses active in the Retain System. The status displays active users, inactive users, merged mailboxes, and individual counts for active and inactive users on each licensed module.

3.14.1 Download Your License

You will get an email with your validation keys, there generally be more then one, including one for server access and another for the module used for archiving from your email system.

- Browse to the License Portal https://licenses.gwava.com/
- Select the proper Product
- Enter the Validation Key and click Next
- Fill out the contact information and click Next
- Download the .PEM file

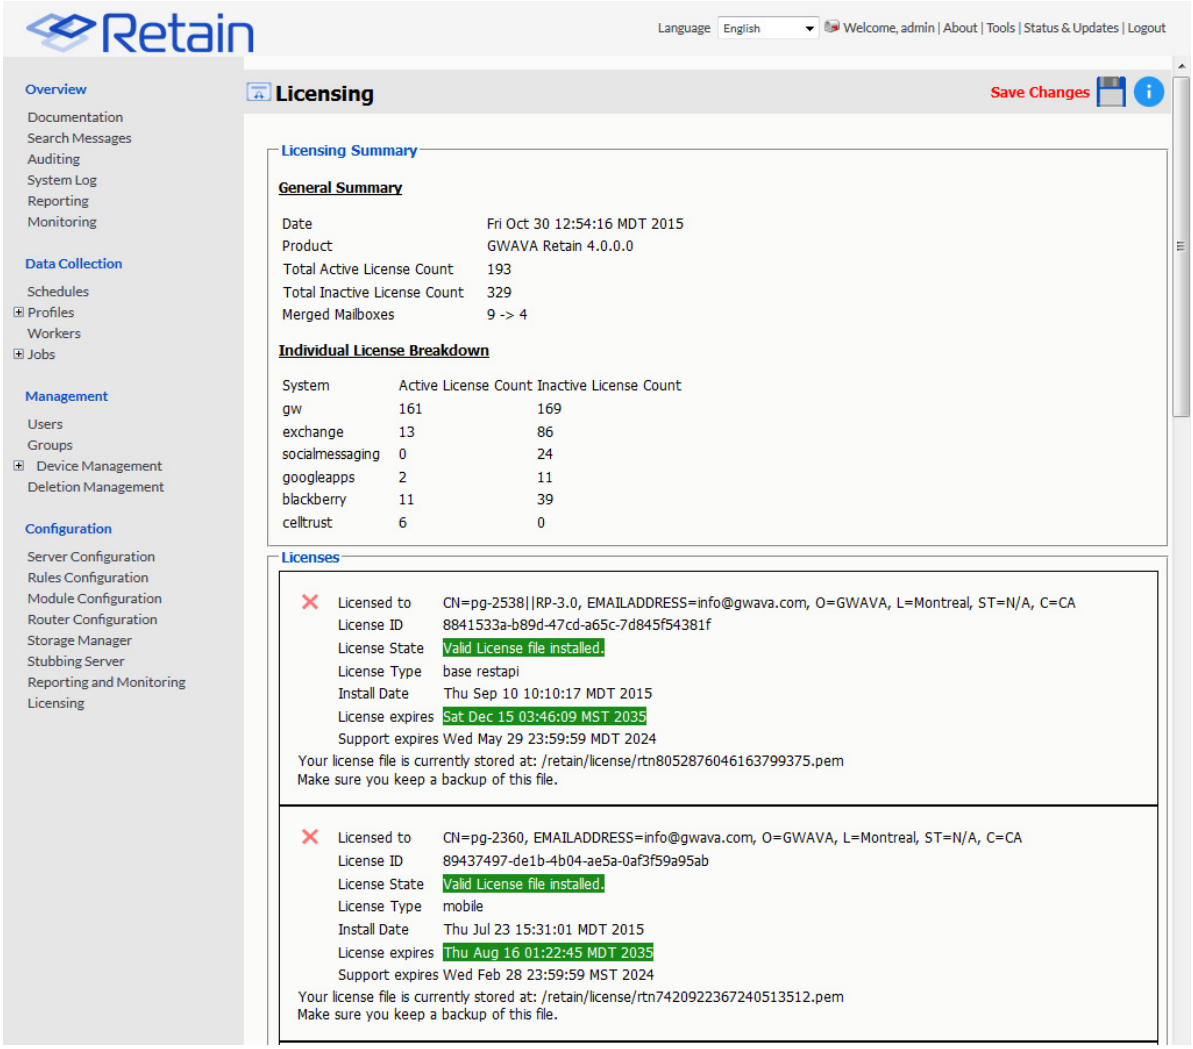
General Summary

- Date contains today's date.
- Product is what version of Retain is currently installed.
- License Count is the number of users archived by Retain. Only mailboxes that have been archived will be counted as active or inactive, not every entry in the mail system's address book. Mailboxes are marked as inactive if no new data from that mailbox had been archived in the last 90 days.
- Merged Mailboxes is the number of mailboxes merged with Mailbox Mapping Options and/or the Mailbox Merger/Splitter.

Generate Licensing Report will create a PDF file which displays the same information as seen in the ‘License Summary’ section and lists the licenses.

Licenses

Individual licenses which have been already added to the system are displayed under the individual license breakdown.



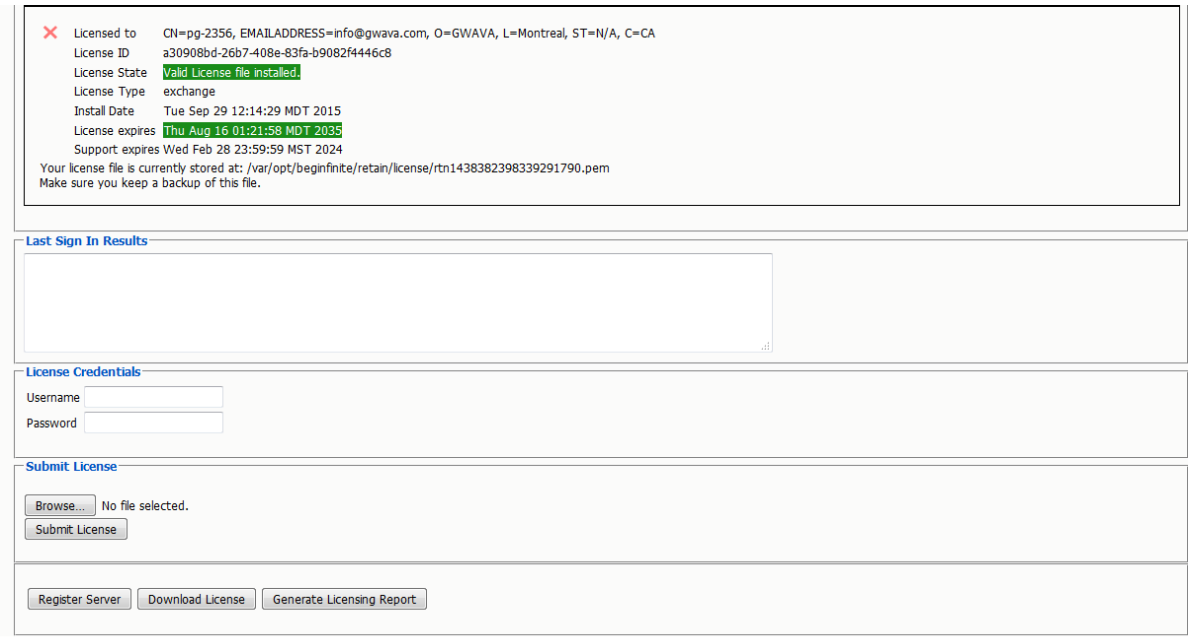
Without a license, Retain will function for 30 days for evaluation purposes.

After this period elapses, the following occurs:

- Users without the "Manage Servers" right are not allowed to log in at all.
- Users with the "Manage Servers" right are allowed access, but only to the Licensing Page.
- All jobs continue to be processed normally

Submitting a License

To upload a new license, browse to the bottom of the page and Select ‘Browse’ from the ‘Submit License’ section. Browse to and select the desired license file, and then select ‘Submit License’. The new license will be added to the system. Logout and re-login to the server to see the license in effect.



Licenses may also be downloaded from the Server through the ‘Download License’ button.

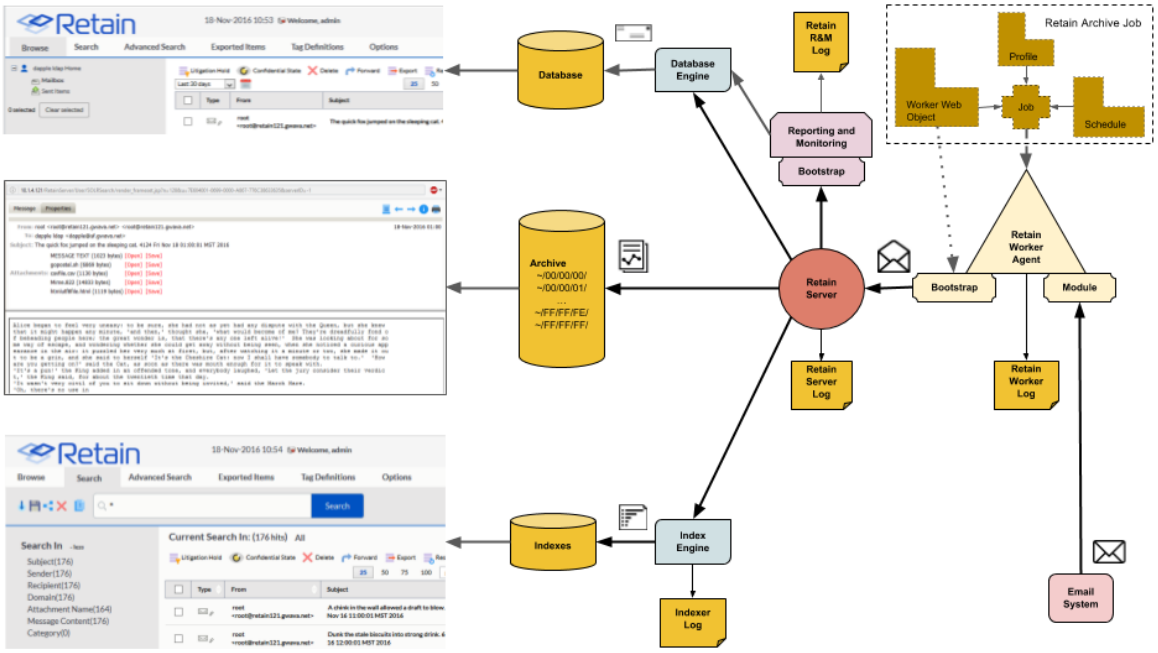
3.15 Backing Up Retain

When you first setup Retain and dredge the email system you have two identical data sets. Over time however, items are deleted from the email system. That means Retain is the only repository of the data. When it comes to backups: "Two is one, one is none." That means if one fails, a backup remains or not. Depending on the type of your organization the legal ramifications of lost data can be significant.

3.15.1 Where Data Is Stored In Retain

There are only a few major places where data is stored in Retain

1. Program directory
2. Archive directory
3. Index directory (this may be on an external cluster)
4. Database directory (this may be on an external server)



The Reporting and Monitoring Server data is stored in the database.

The Archive Job configuration is stored in the database.

Each major part of Retain keeps logs of what is happening the primary ones being the RetainServer, RetainWorker, Indexer and the Reporting and Monitoring log as RetainStatServer.

When you are interacting with Retain, depending on what you are doing you are viewing different parts of Retain the data.

- When you browse messages, you are viewing the metadata of the message that is stored in the database.
- When you search messages, you are viewing the indexes of the messages.
- When you open a message, you are viewing the message as saved on disk.

All message content and attachments to messages are stored on disk in the Retain storage area in a directory off of the "archive" directory (see the figure under "Lucene Indexes" to find the archive directory location). Every message and attachment gets assigned a "hash". Because the byte count of every message and file will be unique, its hash value will be unique. This is how Retain Server determines whether a message and/or attachment has already been processed and stored on disk when an archive job runs. That file's hash value is stored in the Retain database in the t_document and t_attachment tables.

The archive directory uses a load balancing strategy on disk. Thus, off the archive directory you'll find 256 two-digit subdirectories: 00 through FF. Each of those directories have their own set of 256 directories utilizing the same naming sequence (00 through FF). Additionally, those directories also have their own set of 256 subdirectories. Thus, if my filename were B4F05EECB7B21D9014A86C32291C913D190C33394365AC79ED3E1F6849532, I would find it under .../archive/B4/F0/5E.

When a user clicks on a message link in the Retain mailbox - whether from the Browse tab or the Search tab's search result list - Retain finds the file on disk and places the contents in the message window. If the original message was known to have text and the message window comes up blank, the file is missing from the location that Retain thinks it is in. This is extremely rare and usually only

happens as a result of moving the archive directory to a new location. In such cases, we find that either the files did not all copy over properly from the old location or the administrator forgot to tell Retain where the new location is at.

3.15.2 Backing Up Retain

The archive directory consists of up to approximately 16.7+ million directories and the archives are stored evenly across them. This makes it hard (if not impossible) for the traditional file-based backup systems to back it up; thus, you either need a disk image (block level) backup or you need to use Micro Focus's Reload for Retain backup solution, which was recently developed due to high demand by Retain customers for a more elegant and simple solution.

The three most critical pieces that must be backed up are the archive data (specifically, the "archive" directory), the Retain database and the configuration files. Losing even one of those pieces would mean that your archive data is completely lost. They are interdependent. Indexes can be recreated, which will only take time.

Virtual Machines

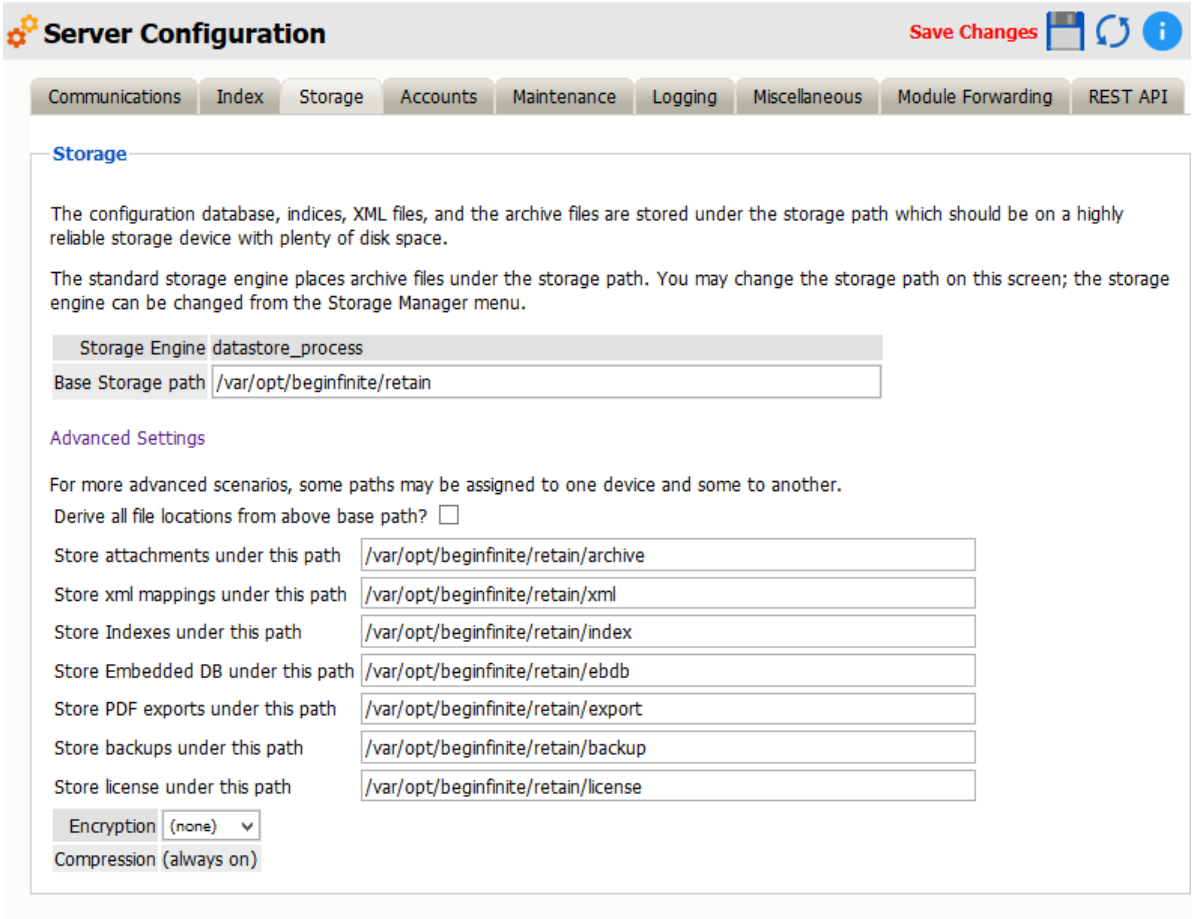
If you are running Retain on a VM and if you are using a purchased version of VMWare at any level, the purchased version provides a disk backup utility. This can be used to back up Retain if the VM's local disks are part of the VM guest itself. However, those backups can also take a lot of time as the data grows, so even in those circumstances, you still may want to consider using Reload for Retain which knows what items are new and backs up only them new items.

If the disks are external to the VM guest, then those disks need to be backed up. This article assumes that the reader understands how disks work with VMs. If you do not fully understand virtual machine concepts, we recommend that you consult with the person that set up and maintains your VM environment.

Storage Path

The storage path(s) is where you'll find your archive files, the indexes, license file, and Retain's backup of the indexes and the ASConfig.cfg.

To find your Retain system's storage path, login to the Retain Server administration web interface (http://[ipaddress/DNS hostname]/RetainServer). Starting at the navigation pane on the left side under "", click on Server Configuration | Storage. Most customers leave it at the default setting under "Advanced Settings" of "Derive all file locations from the above base path", but you will have to click on that checkbox to have it expand to show all the paths.



The information shown in the Storage tab in the Server Configuration screen is also stored in the ASConfig.cfg.

What to Back Up

Other than backing up the archive directory structure, there are a few areas of Retain that are important to backup that would not require a disk image (block level) backup:

1. Configurations files
   - ASConfig.cfg
   - Indexer configuration files
2. Database (critical)
3. License
4. Index files (these can be rebuilt but that process can take days, weeks or months during which time your searches will not produce full results)

File Locations

1. Configuration files.
   - ASConfig.cfg is stored in a directory off of your Retain installation:

     Linux default: /opt/beginfinite/retain/RetainServer/WEB-INF/cfg

Windows default: [drive]:\Program Files\Beginfinite\Retain\RetainServer\WEB-INF\cfg

- Indexer configuration files (the entire directory's contents):

  Linux default: /opt/beginfinite/retain/RetainServer/WEB-INF/solrweb/WEB-INF/cfg

  Windows default: [drive]:\Program Files\Beginfinite\Retain\RetainServer\WEB-INF\solrweb\WEB-INF\cfg

2. Database.
   - The location is too varied to mention here. Each customer should know where their Retain database resides.
3. License.
   - The license directory is located under your storage path.
4. Indexes.
   - Because the Index directory can be in a constant state of change, it is recommended that you back up the index subdirectory located beneath the backup directory (also found under the storage path).
   - When the Retain maintenance routine runs, it makes a backup copy of the index directory and places it here. The frequency of this backup is configured in the RetainServer interface under Server Configuration | Maintenance.

If performing an upgrade of the Retain software you are strongly advised to manually back up all the files mentioned in this article before performing the upgrade.

NOTE: on backing up during the nightly maintenance cycle:

In Retain 4.0.3 and higher the indexes are optimized during maintenance every night. This may take a hours after an upgrade. During optimization the index directory may grow by 2-3 times as temporary files are created and removed. Backing up during this time is not recommended.

3.16 To uninstall Retain

To uninstall Retain, use the following general procedure. Actual locations may vary according to your spe-cific setup, storage location, and host operating system.

Linux

1. Shut Down Tomcat: “service retain-tomcat8 stop”
2. Shut Down Apache: “service apache2 stop”
3. (Optional) Remove the data storage files

   Can be: /var/opt/beginfinite/retain. See the Retain Administration Guide section on storage to see where your data actually is placed.
4. Remove the program files

   Can be: /opt/beginfinite/retain
5. Remove the RetainWorker.xml, RetainServer.xml, RetainRouter.xml, and RetainStatsServer.xml from Tomcat

   Can be: /opt/beginfinite/retain/tomcat8/conf/Catalina/localhost
6. Remove the directories under tomcat/work (but NOT work itself)

   “rm –r /var/cache/retain-tomcat8/work/*”

7. Remove references to Retain in httpd.conf

   Can be: /etc/apache2/conf.d/retain-apache.conf
8. Remove tomcat service file: “rm /etc/init.d/retain-tomcat8”
9. If Apache and Tomcat were not removed, restart them both to verify that they load without problems.

Windows

1. Shut down Tomcat service
2. Remove Retain Tomcat 8 service

   Open a terminal as admin and type: “sc delete tomcat8”
3. Remove the program files

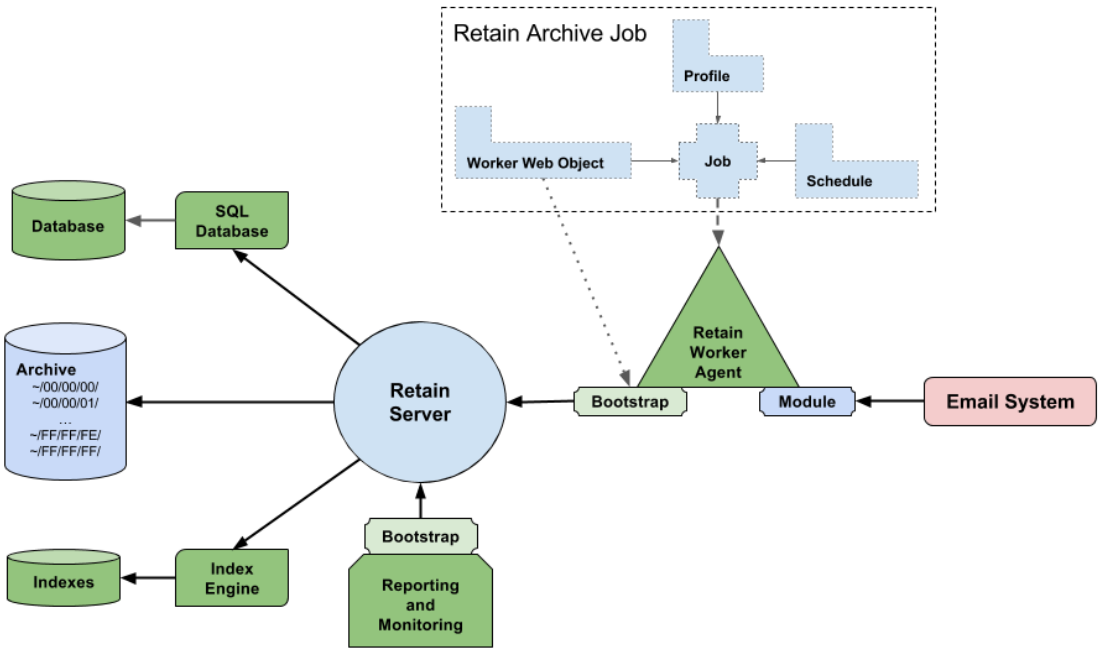
   Can be: C:\Program Files\Beginfinite\Retain
4. Remove the environment variables: RETAIN_CATALINA_HOME, RETAIN_HOME, RETAIN_JAVA_HOME, RETAIN_JAVA_VERSION, and RETAIN_JRE_HOME
5. (Optional) Remove the data storage files

   Can be: C:\Retain
6. (Optional) Remove SQL Database, SQL Server, Apache or IIS

4 Modules and Jobs

4.1 Modules and Job Overview

Modules are how Retain connects to a messaging system and jobs specify what Retain archives.



4.1.1 Modules

Modules are where you provide the data needed for Retain to connect to the messaging system. This includes the address of the messaging system server and credentials to enter the system, such that Retain can access all the mailboxes.

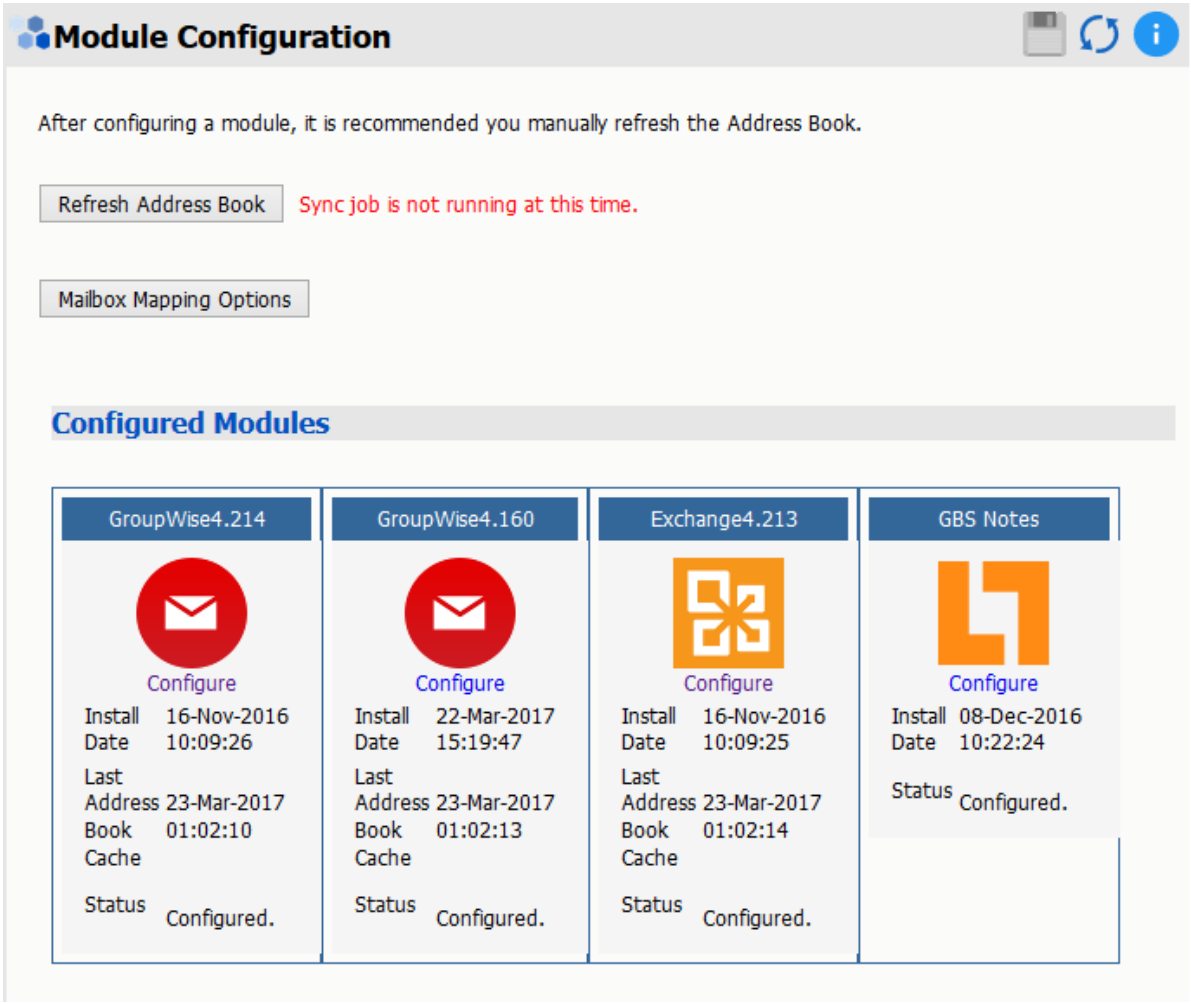
This is used by the Retain Worker to connect to the messaging system. It is the Retain Worker that does all the work of bringing the data into Retain.

Since the Worker Agent software can be installed on a server separate from the Retain server itself we need to provide it with a bootstrap so it knows how to connect tot he Retain server.

Multiple Modules

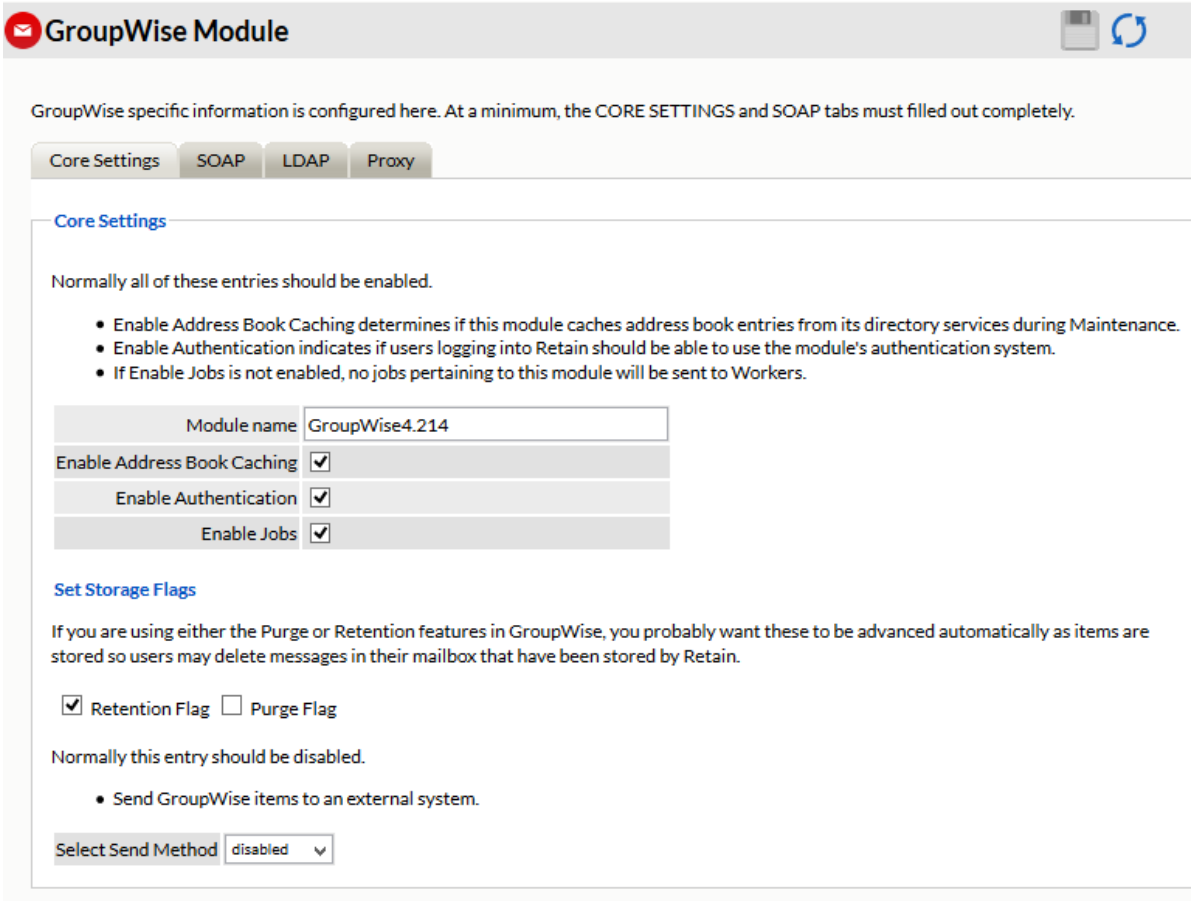
In Retain 4.2 and above, multiple modules of the same type can be enabled. This allows Retain to connect to multiple messaging systems of the same type at the same time that do not share common access. The modules can be given different names to make it easy to distinguish.

Multiple modules is enabled for GroupWise. Exchange, Google Apps (G Suite) and Mobile modules. It does not apply to the other modules.



You enter the Module name in the Module.

For example, the GroupWise4.214 module has a different domain and trusted application key compared to the GroupWise4.160 module.



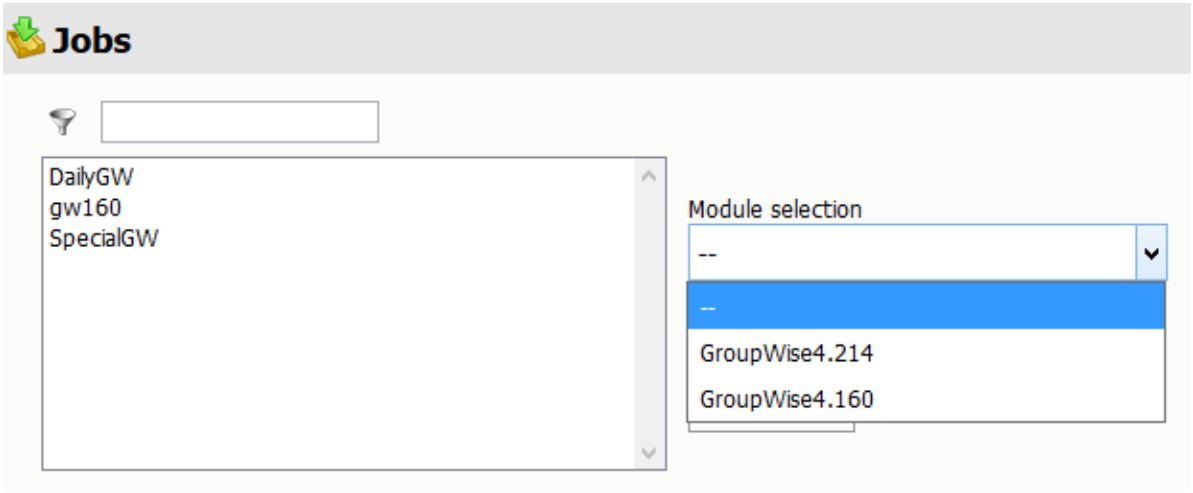
4.1.2 Jobs

Jobs specify what is to be dredged or downloaded from the messaging system by specifying four things:

- *Schedule*: When the dredge is to take place.
- *Profile*: What will be dredged from the messaging system
- *Worker*: Who will do the dredging.
- *Jobs*: Which mailboxes will be dredged.

Jobs with Multiple Modules

In Retain 4.2 and above, when you create a job with multiple modules you will have to select the module by name before creating the job.



You can select any existing jobs without having to preselect the specific module. If you do select a module, then only the associated jobs will be shown.

4.2 GroupWise Module

The GroupWise module allows Retain to archive messages from GroupWise. The GroupWise module must be fully configured before Retain can archive or communicate with the GroupWise mail system.

With GroupWise, the Stubbing server can be installed which can allow users to search the Retain archive from their GroupWise client.

4.2.1 Prerequisites

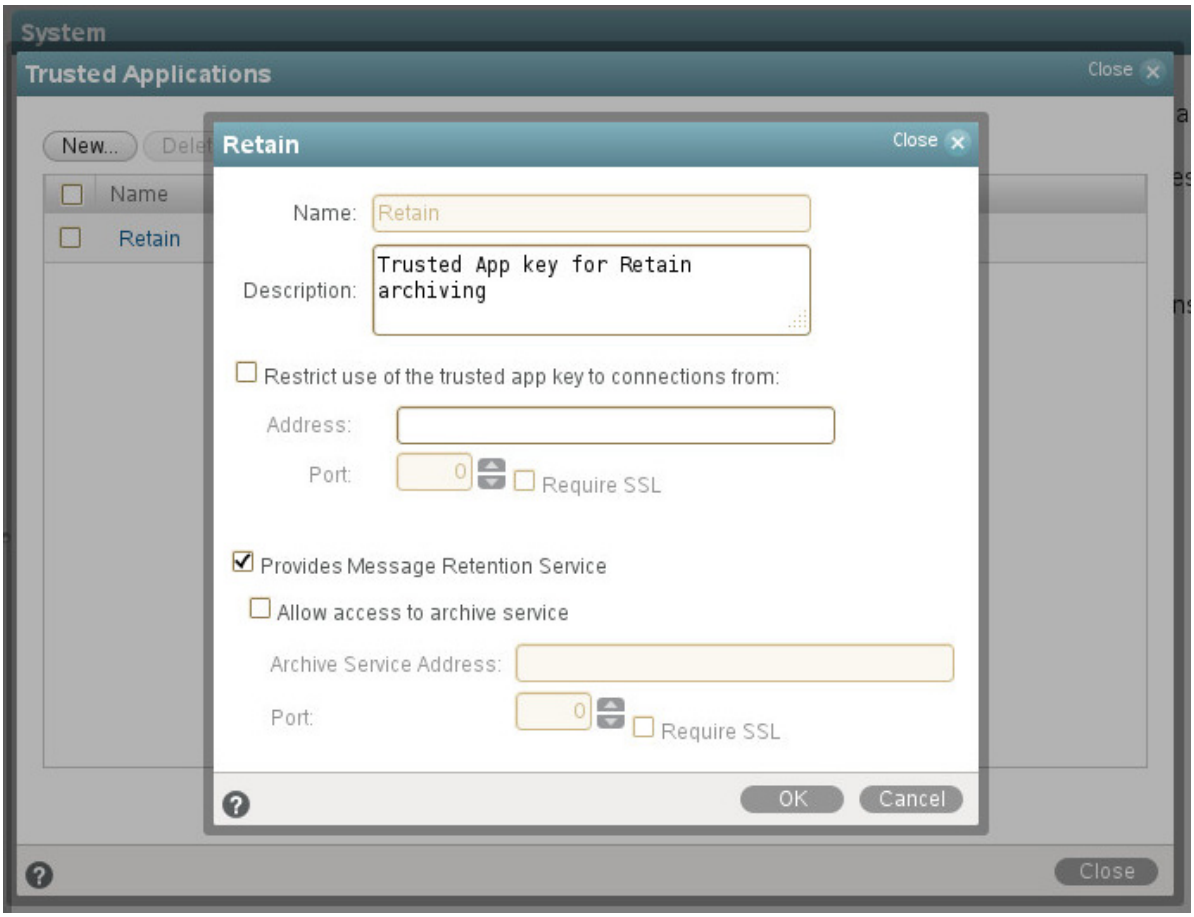
- IP Address and SOAP port of a POA
- A Trusted Key Name and Trusted Application Key

Setting up the Module

The GroupWise module page opens first with the Core Settings.

GroupWise

Browse to the GroupWise Administration page. Under *System* | *Trusted Applications*, create a key called Retain and export the key. You will need this for the SOAP section below.



We highly recommend enabling Message Retention Services in GroupWise, so users cannot delete messages before they are stored by Retain.

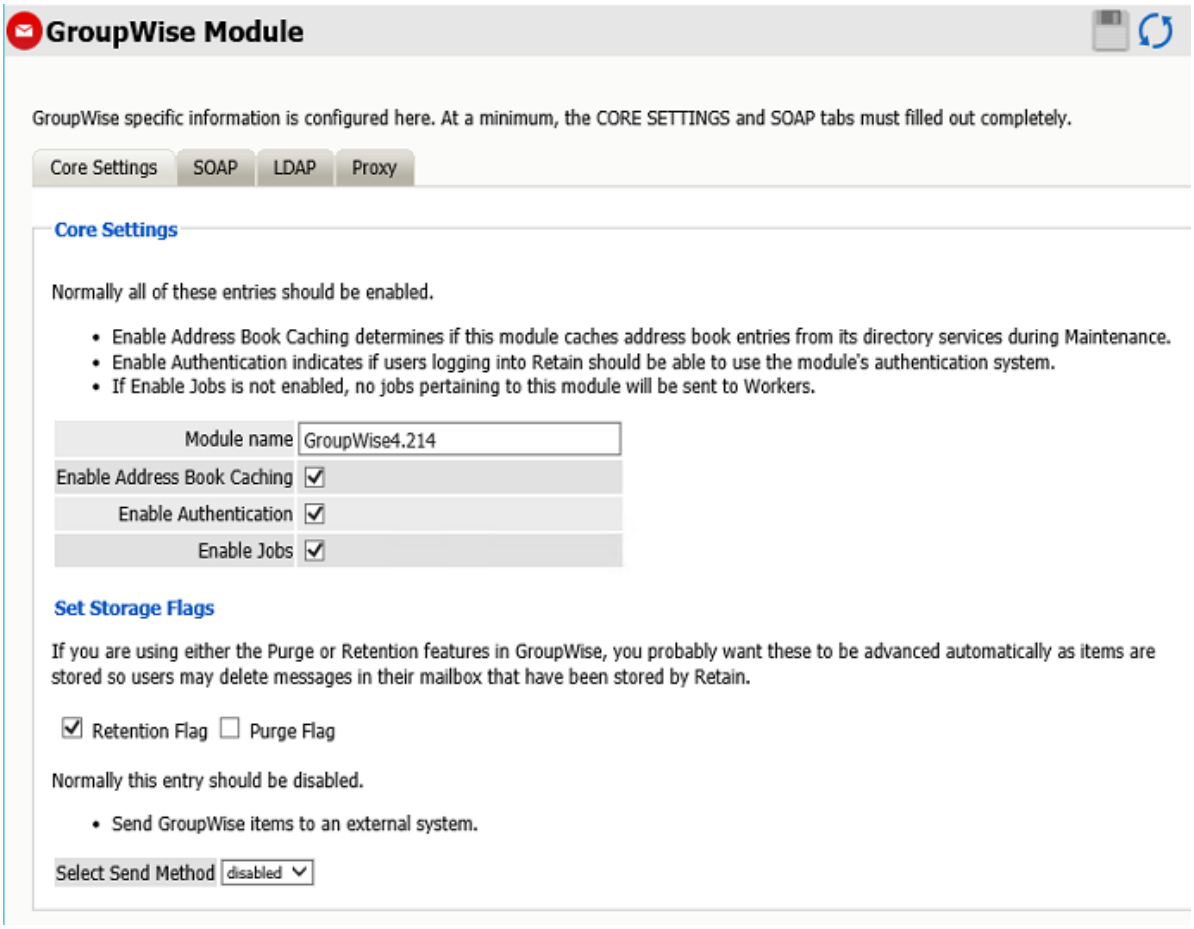
4.2.2 Core Settings

Normally all the checkboxes on this tab are always left selected. It is rare that you would ever deselect any of them. Two cases where you might, would be: troubleshooting (as instructed by Technical Support), and retiring an old email system.

The module needs to be enabled on this page to make it active in the Retain system.

The module can be given a name.

The Send Method option enables either the SMTP Forwarding or FTP features. For either feature to appear and function, the Module Forwarding tab must be configured on the Server Configuration page. See that section for more information.



The Enable Address Book Caching function allows Retain to regularly cache the online email systems address book and synchronize it with Retain. This is critical for administration, authentication, and archiving purposes. It is recommended to cache the Address Book once every 24 hours to keep the Retain storage system up to date. By default, maintenance is set to cache the Address Book once every 24 hours.

The Enable Authentication checkbox determines if end-user authentication is performed when the user logs into Retain. If it is deselected, the Retain system will NOT authenticate the user against the email sys-tem and the user will not be able to log in unless another authentication method is enabled.

The Enable Jobs checkbox determines if configured data retrieval jobs are ever passed to the Worker. Even if the individual job is fully configured and enabled, if this checkbox is switched off, no jobs config-ured for this module will be processed.

Set Storage Flags

If you have Retention Services enabled, which we highly recommend, you will need the Retention or Purge flags updated depending on which one you are using. Generally, you would use Retention flag for Retain and the Purge flag for your backup software, such as Reload.

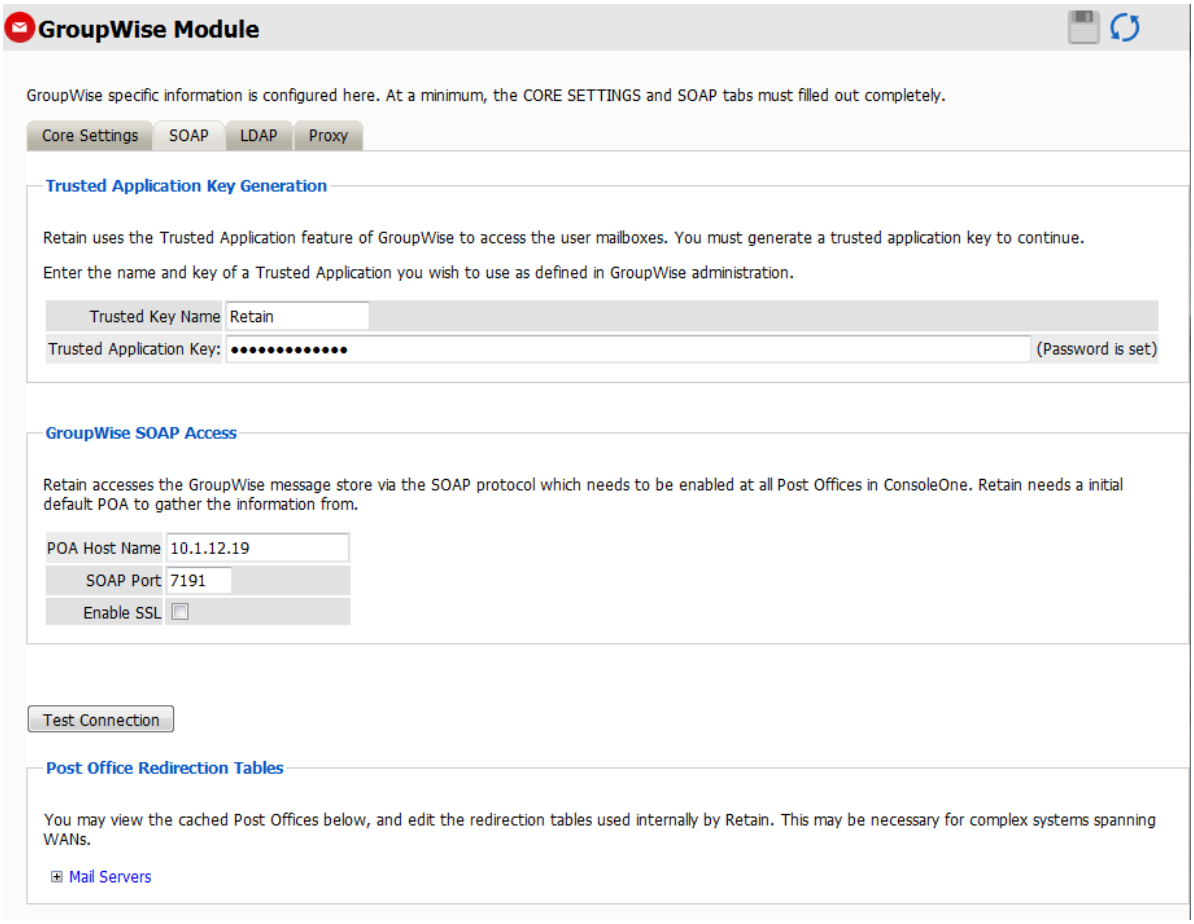
You would enable Retention Services in GroupWise Administration under Domain or Post Office. Then in Client Options, Integrations, Retention, and *Enable Message Retention Service.*

This would be used in conjunction with a Trusted App Key which the GroupWise module needs for the SOAP tab.

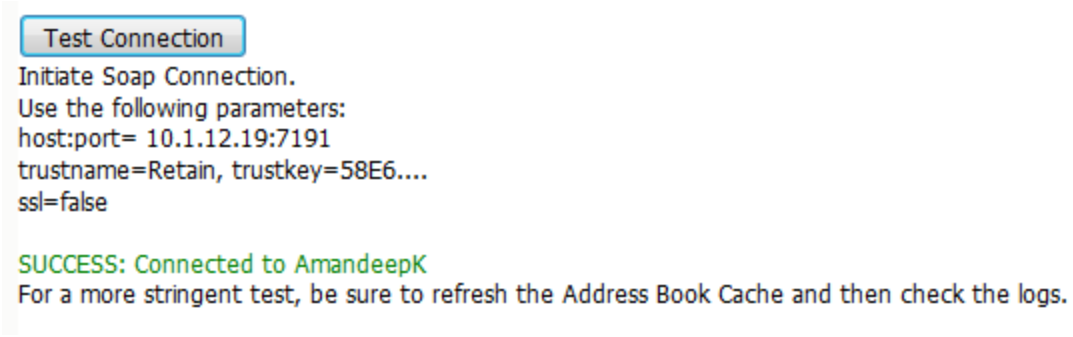
The Send Method option enables either the SMTP Forwarding or FTP features. For either feature to appear and function, the Module Forwarding tab must be configured on the Server Configuration page. See that section for more information.

4.2.3 SOAP

Provide the POA Host Name and SOAP port. IP addresses are acceptable, but host names are preferred as IP addresses may change. SSL is supported, but comes at the price of speed, and will slow down the archiving process. Create a Trusted Application key for Retain, either manually from GroupWise Administration, System, Trusted Applications in GroupWise 2014 or above or Console One for GroupWise 2012 or lower.



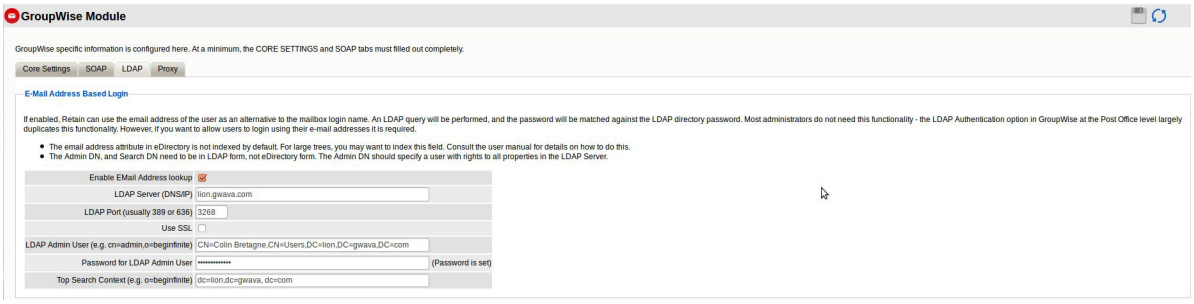
The SOAP access information must be provided, and the connection tested and verified before the sys-tem can connect. After providing the required information, click the 'Test Connection' button. The results are displayed. A successful result must be reached before Retain can archive messages from GroupWise.



If mail server Redirection is required for mail servers which are not contained on the local WAN and must have the connection addresses manually specified, the appropriate information may be modified in the redirection table. Most installations will not require any modification.

4.2.4 LDAP

LDAP can be used to authenticate users against other directory services such as eDirectory or Active Directory.



LDAP may be used for individual users wishing to access their respective archives. If LDAP is setup and desired to be used for Retain user authentication, it must be fully configured in the GroupWise module.

Utilizing LDAP allows users to log into their respective archives using the user's full email address. This authentication requires that the email attribute be marked indexed in GroupWise.

Active Directory Authentication

Prerequisites

Active Directory LDAP Services setup

LDAP Authentication setup in GroupWise

- Go into GroupWise Administration/System and open LDAP Servers

  Create a New Directory with the IP address of your AD server

  Set the LDAP user name (for example, Administrator) and password

  Set the Base DN (for example,DC=sf,DC=gwava,DC=net)

  Choose the Sync Domain (the GW domain)

Enable Synchronization.

- Create a New LDAP Server

  Select the Directory

  Enter the IP) address of the AD server.

  Under the Post Offices tab, select a post office

- Go to Post Offices and choose the post office that will be using LDAP.

  Under the Security tab choose LDAP Authentication and select Preferred LDAP Servers

- Go to System and open Directory Associations

  Enable Search Sub Tree

  Associate Mailboxes with their LDAP Names

Retain

LDAP allows the users to authenticate against Active Directory, using their email address, by following these steps:

- Fill out the LDAP tab

  Enable EMail Address lookup

  Provide the LDAP Server hostname or IP address

  Provide the LDAP Port (usually 389 or 636)

  Enable Use SSL, if necessary

  Provide the LDAP Admin User's distinguishedName in the form CN=Administrator,CN=Users,DC=company,DC=com

  Provide the Password for LDAP Admin User

  Provide the Top Search Context, usually the distinguishedName of the domain in the form DC=company,DC=com

- Edit the file ~\Beginfinite\Retain\RetainServer\WEB-INF\classes\config\misc.properties

  Change the following lines:

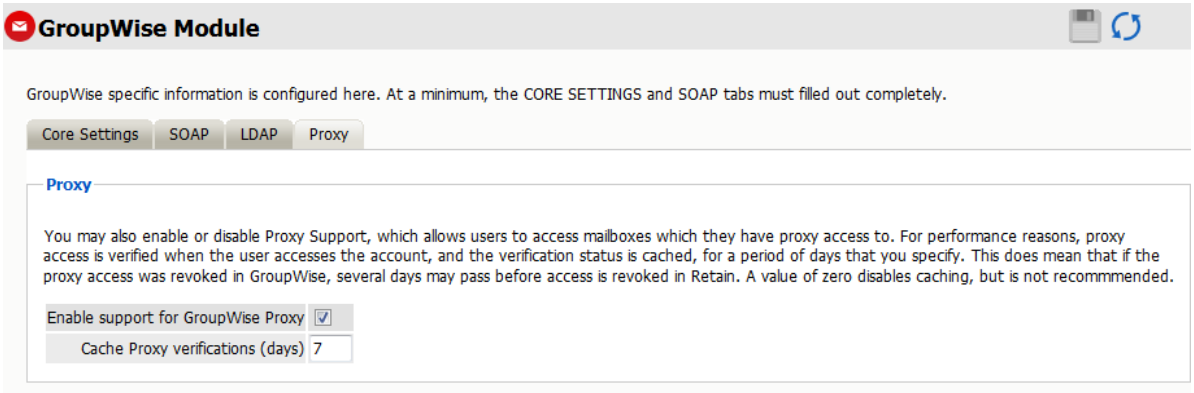
  custom.ldap.enable=1

  custom.ldap.class=*com.gwava.authenticate.gw.AlternativeGWLDAPAuthentication*

- Restart tomcat

4.2.5 Proxy

Users who have been given proxy to another account in GroupWise may be granted access to proxy accounts in the Retain archive as well.



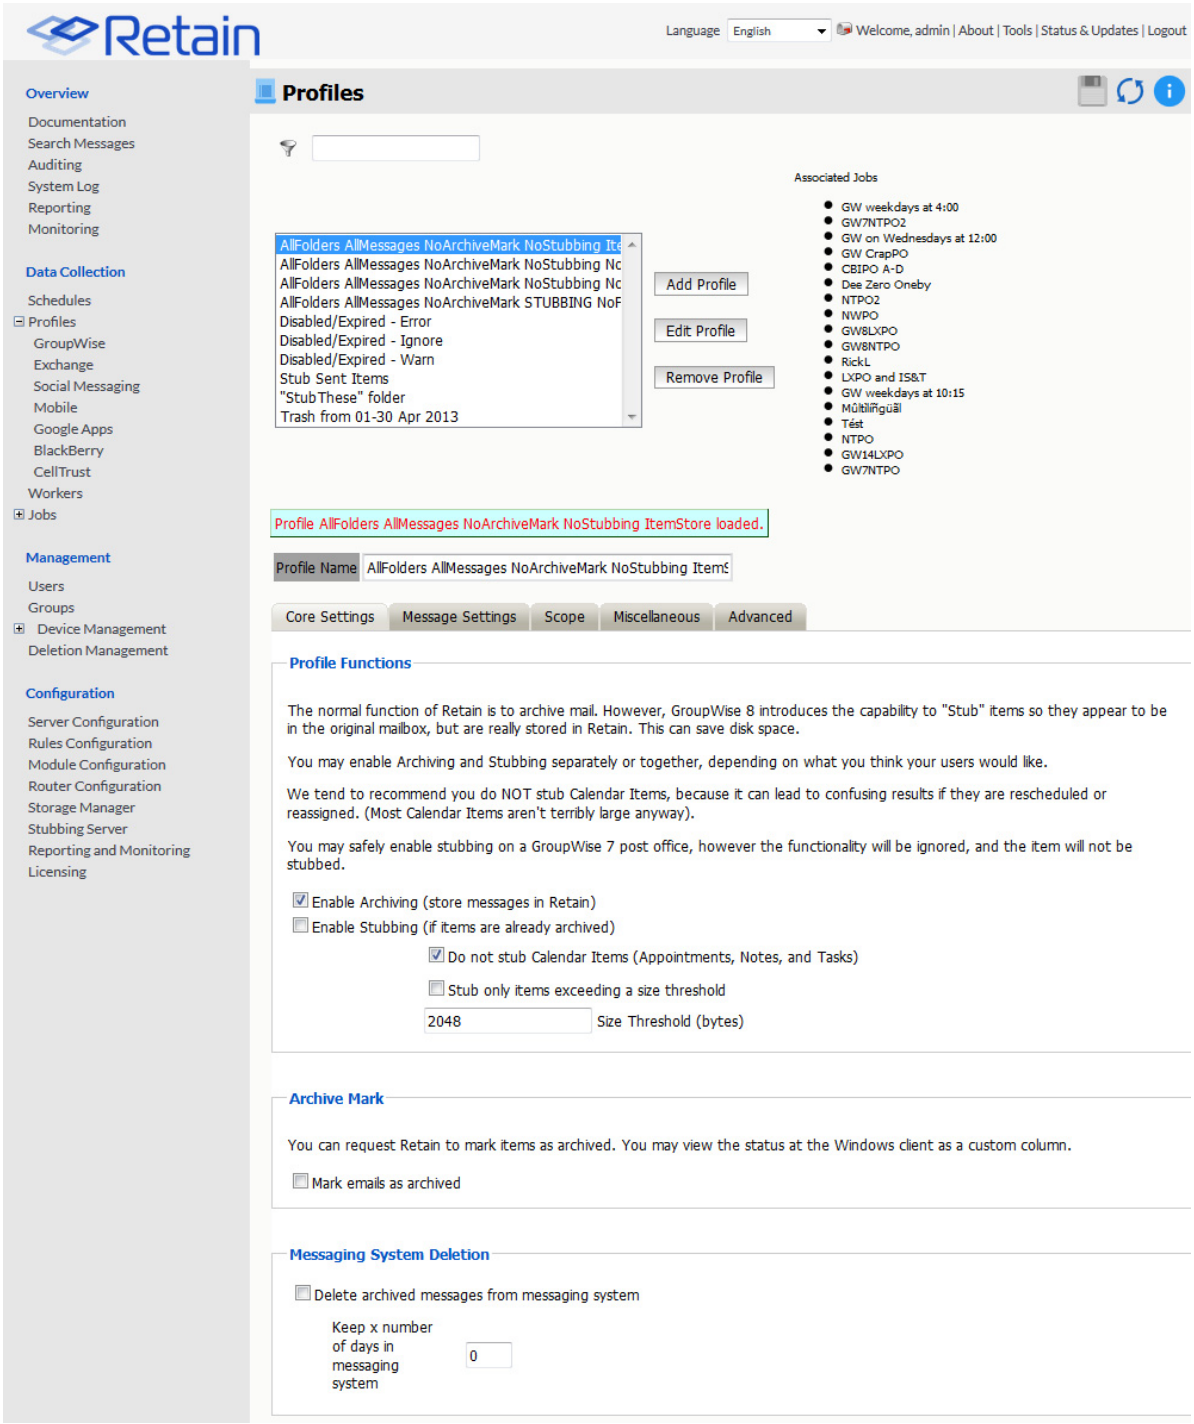
For performance reasons, Retain caches proxy verifications for a period of days (default is 7 days). Revocation of proxy access might not be reflected immediately in Retain. The caching period may be reduced or even disabled, (a value of ‘0’ disables caching), but this is not recommended.

4.2.6 GroupWise Profile

The Profile allows you to set what is to be dredged from your GroupWise system.

Click on “Add Profile” and provide a profile name, or select an already existing profile to access the con-figuration tabs. All changes made on this page must be saved by selecting the “save changes”, disk icon, at the top right of the page. Tabs may be changed and navigated through without affecting new settings, but any move to another page will require saving, or abandoning the changes made.

Core Settings



4.2.7 Profile Functions

The Profile Functions tell the Retain Server what to do with the mail it archives from the GroupWise sys-tem. If Archiving is not enabled, mail will not be archived by Retain.

4.2.8 Messaging System Deletion

For systems where the administrator wishes to have archived messages removed from the system auto-matically, the Messaging System Deletion option may be used. Messaging System Deletion will remove messages from a mailbox after they are archived, according to the time frame specified in the settings. The amount of time to keep messages is specified in days. The recommended setting depends on the archiving scheme in the system. For instance, if messages are to persist in the system for 30 days, then the system deletion setting should be set to ‘30’ and enabled. A setting of ‘0’ will remove messages from the system as soon as they are archived. Be sure to configure the system before enabling the setting in the profile.

4.2.9 Stubbing

Before enabling stubbing, make sure you understand the way stubbing works and how that will work in your system, because, once enabled, the stubbing process cannot be reversed. Stubbing, which requires GroupWise 8.02HP1 or above, allows Retain to replace archived messages and items in the GroupWise Post office with ‘stubs’ that link back to the full message or item in the Retain Archive. End users should not notice a dif-ference as they have transparent access to the stubbed item through the GroupWise Post office. The stub behaves much like the native item is still in the GroupWise system.

Stubbing cannot be reversed.

Plan your system and understand the positive and negative sides to Stubbing before you continue. To use stubbing, it is strongly recommended that the GroupWise system be at least 8.02HP1. Previous versions of GroupWise do not contain critical stability fixes.

Why stubbing is good:

- Save Disk space. Mail can be removed from the system without being ‘deleted’.
- Users have easy access to items stored in Retain
- Transparent access, no need to login to Retain for common tasks.

Why stubbing may not work for your system:

- Mobile, Web, IMAP, and POP3 clients will only see posted empty items, not actual contents.
- Stubbing cannot be reversed. (Once a message is stubbed, it cannot be returned to the same position in the system)
- User training for stubbed items may be necessary due to decreased access speed
- No method for protecting stubs from deletion in the GroupWise client. (Recreation is difficult and complicated.)
- Non-mail items should not be stubbed. (Calendar and appointment items behave erratically when rescheduled).
- Stub items are not compatible with third party API’s. These applications view the stubs as posted items only, with no text or attachments.

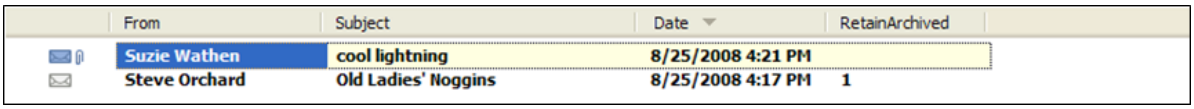
Determine the correct option for your system. If you wish to enable Stubbing, you need to install and configure the Stubbing Server as well. Refer to the install guide for your specific OS and the Stubbing Server sections. Stubbing can only be active for items which have been archived by Retain.

Calendar Items should not be stubbed unless necessary, as this can cause erratic behavior if archived items are rescheduled. If you wish to only stub large items in your Post Office, you may set a size thresh-old. Any item larger than the specified size will be stubbed. If this option is not set, all messages will be stubbed, regardless of size.

4.2.10 Archive Mark

Some users may opt to use the Archive Mark in GroupWise for messages that have been archived by Retain. The archive mark is a custom flag and may be modified, therefore is not secure and should not be used for compliance. Archive Mark slightly degrades job performance. Check the check box to enable Archive Mark for the selected profile.

When the Archive Mark is active, Retain creates a custom column for mail, called “RetainArchived” which users and administrators may add to their GroupWise clients to view mail which has been archived. The “RetainArchived” column indicates an archived mail item by displaying a ‘1’ in the message row, while remaining blank when the message is not archived.



To add the column, in the client: highlight desired mailbox or folder, right-click on columns header, select More Columns, select RetainArchived, and select Add.

Messaging System Deletion

Retain can delete items from GroupWise after they have been archived by enabling "Delete archived messages from messaging system".

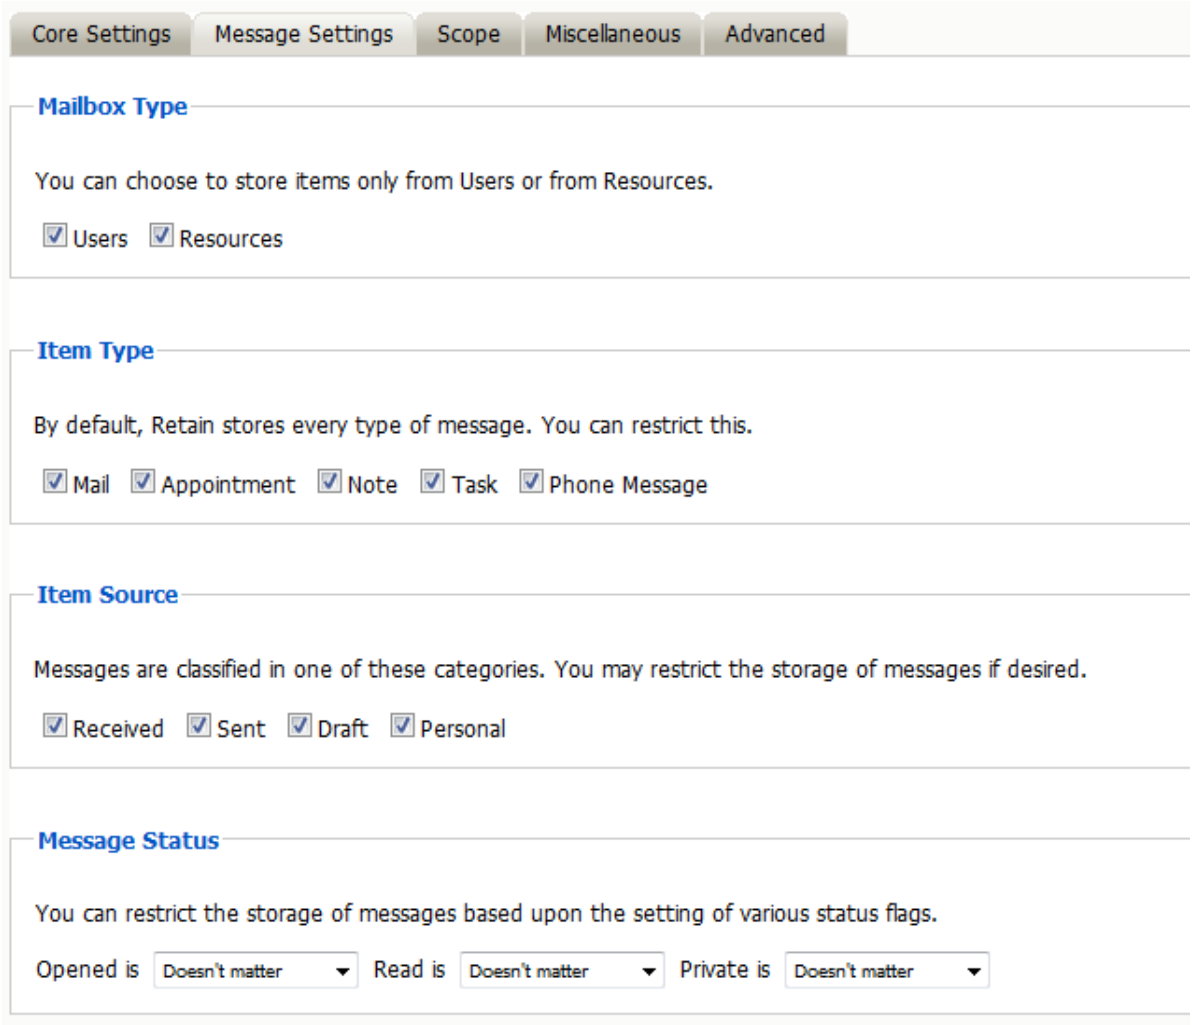
Set how far back the messages will be kept by entering a reasonable number. like 90, of "Keep x number of days in messaging system".

Message Settings

The message settings tabs allows you to restrict the types of messages to be archived.

- Mailbox type: You can restrict the type of mailbox users, resources, or both
- Message Item: Select what type of items within these mailboxes to dredge
- Message Source: Select whether incoming, outgoing, draft, or posted items get dredged
- Message Status: You can decide whether to dredge an item based on its GroupWise flags

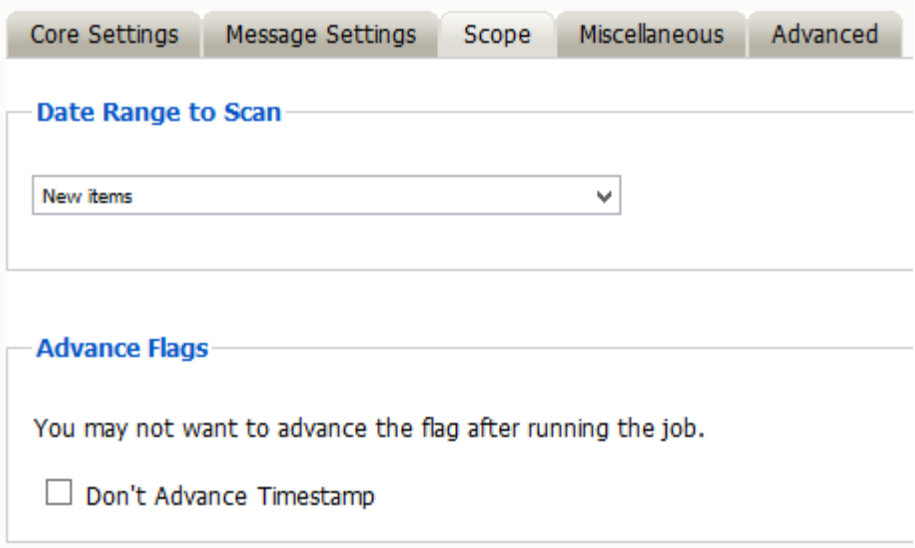
The default settings dredge all items, regardless of flag status.



By default, all items are selected for collection.

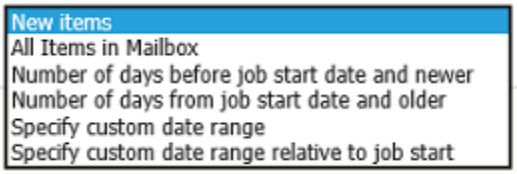
4.2.11 Scope

This is the most critical tab to fill out as it sets the limits on how much to archive.



4.2.12 Date Range to Scan

The Date Range determines which message items are collected, depending on the date of the message.



New Items: All items that have not been archived by Retain since the last time the job ran.

All Items in Mailbox: All items in the mailbox starting from 1/1/1970, duplicates will be processed but not stored if they already exist in the Retain archive.

Number of day before job start date and newer: Only items from the relative number of days from the time the job began will be archived. E.g. messages that came into the email system 7 days ago or less.

Number of days from job start date and older: Only items previous to the relative number of days from the time the job began will be archived. E.g. messages that came into the email system 7 days ago or more.

Specify custom date range: Only items between two absolute dates will be dredged.

Specify custom date range relative to job start: Only items between two relative dates will be dredged. E.g. messages that came into the email system between 7 and 5 days ago.

Advance Flags

Enabling "Don't Advance Timestamp" will not update the timestamp flag. Items that are dredged will still be considered new by Retain the next time the job runs.

This is useful when troubleshooting, but is generally not used for normal jobs.

4.2.13 About GroupWise’s Purge and Retention Flags

The Purge, and Retention flags perform similar functions. The Purge and Retention flags are stored internally in the GroupWise post office on a per-mailbox basis. The Retain Server also keeps its own timestamp flag internally in the Retain configuration database.

Each can be thought of as simply a “time stamp” indicating the last time data was retrieved by a third-party application. There are a few subtle differences between the three flags, which are discussed below.

The Purge flag was introduced by Novell in GroupWise version 6.0.

With Purge enabled in GroupWise:

- The flag is a time stamp, stored in the user's mailbox.
- Items older than the time stamp can be deleted freely.
- Items newer than the time stamp cannot be deleted.
- A 3rd party application, such as backup software can retrieve and/or set this time stamp value. GroupWise never changes the time stamp value.
- The initial purpose of this was for backup software. Backup software can either manually or automatically, (via the SMS TSAGW component of GroupWise), advance this time stamp once all items are backed up. The backup software can also use this time stamp to “know” that all items older than this time stamp have already been backed up, and can be skipped.
- Retain can use this flag to skip archived items, and automatically advance it as well. However, in practice, you'll rarely want to do this, because backup software, including Micro Focus Reload, already uses this flag. There's no way to “share” the flag.

Because makers of retention and compliance software wanted a flag of their own, one that wasn't shared with the backup software Purge flag, the Retention flag was introduced by Novell in GroupWise 6.5.1.

In almost all respects, the Retention flag is identical in purpose and function to the Purge flag. Like the Purge flag, it is stored internally, and is primarily used to prevent premature deletion, and to skip previously stored messages. In fact, these are so identical in purpose, in theory virtually all backup software could switch to using the Retention flag, and all Retention software to the Purge flag, and you'd notice almost no difference. The primary difference is simply it's a different value, used by different vendors.

There are some *small* distinctions:

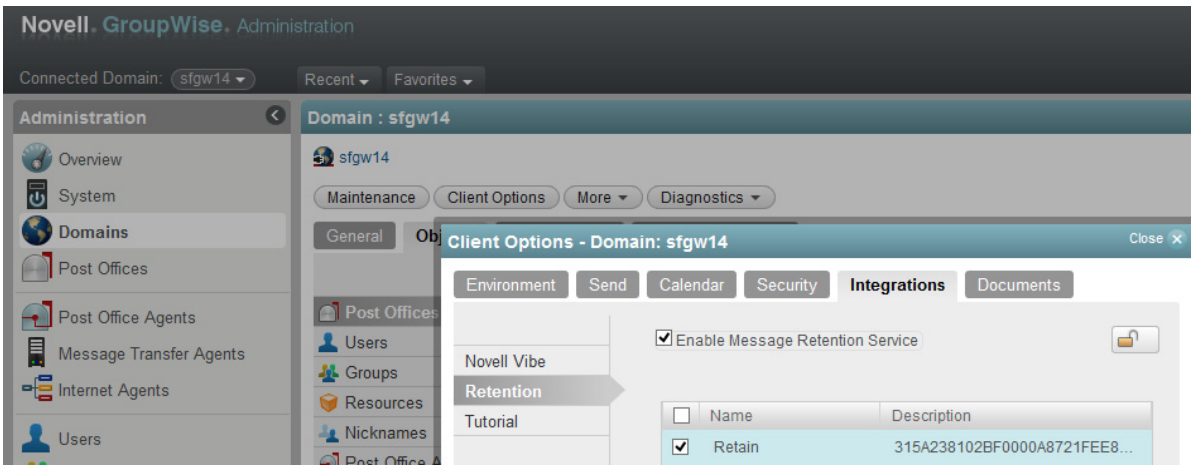
- It's enabled and configured in a different location in ConsoleOne.
- It cannot be enabled or configured until a Trusted Application (like Retain) is installed and specifically requests it be made available.
- The Purge flag is turned on/off globally. The Retention flag can be configured at the Domain, PO, or user level.

Most customers using Retain will want to use the Retention flag. It's logical and consistent to do so, and avoids conflict with backup software. It is recommended that you use this flag unless there is a compel-ling reason not to. A few possibilities are discussed below.

Retain can use the Purge and Retention flags.

4.2.14 Setting up Purge and Retention flags in GroupWise

It is highly recommended that Message Retention Service is enabled, so users cannot delete items before the item is archived.

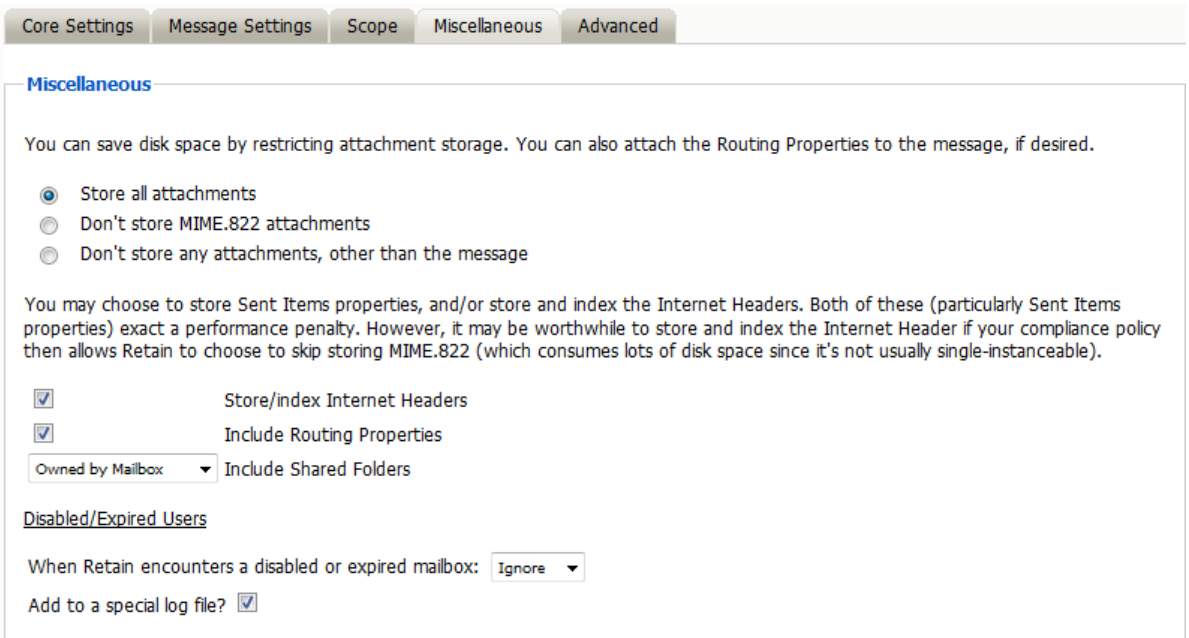


In GroupWise Administration, go to the Domain or Post Office you wish to enable Retention on.

- Open Client Options. Select Integrations and then Retention.
- Check the box to enable the Retention Flag service.
- Select the Trusted Application that will use the Retention Service.
- Finally, Lock it on and Ok to exit.

4.2.15 Miscellaneous

Miscellaneous: What to do with attachments? Most of the time, you will want to store all attachments. That is the default setting.



But if you don't want to store attachments, you can control this here. Note that the attachment names and sizes will still be listed when viewing the message, but the attachments will be clearly marked as unavailable.

Retain can stores the Index or Internet Headers and the Routing properties of messages. This information is stored in a XML attached to the message document and is searchable.

Store/index Internet Headers

This option tells Retain to store the headers in the SQL database. In order to use the header search option in the search interface you must have this option selected. To save disk space and still satisfy storage compliance, select this option and select "Don't store the mime.822 attachments". All other parts of the MIME file are already being archived, so if you store and index the Internet Header, you will already be archiving all information required.

Include Routing Properties

Routing Properties are kept in the MIME.822 file, if you are not storing that file, and are required to store the routing properties of the message then enable this option.

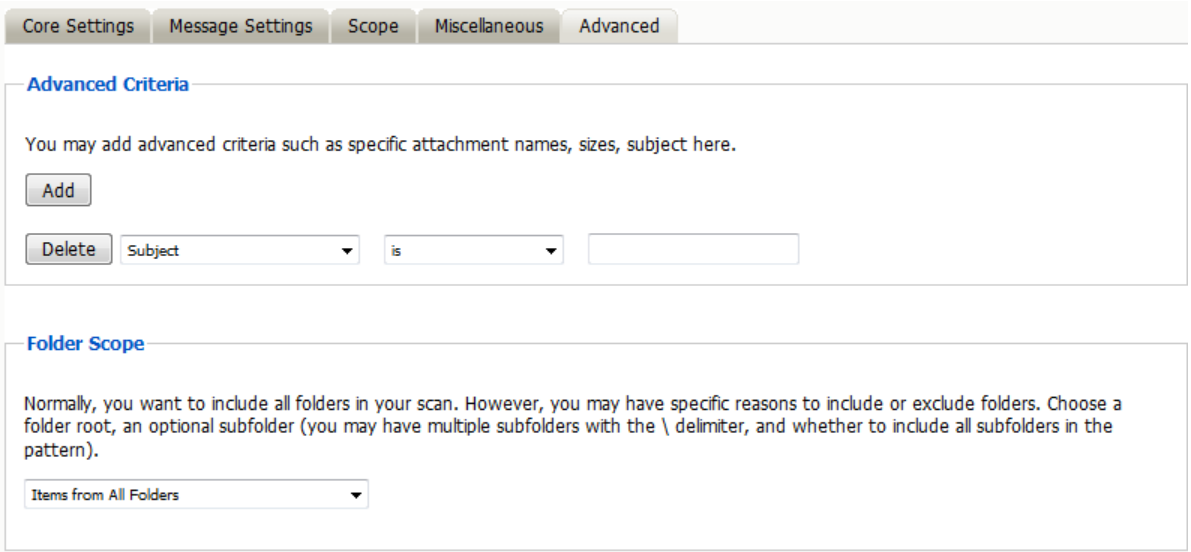
Disabled/Expired Users

This option tells Retain how to react to a disabled or expired user in GroupWise. Be default, Retain will Ignore a disabled mailbox and continue to the next mailbox. The Warn option will show a warning in the logs and the Reporting and Monitoring Server reports. The Error option will show an error in the log, will send an error email, if the job is configured to send error messages, and be shown in the Reporting and Monitoring Server reports.

A special log file can be created in the logs location.

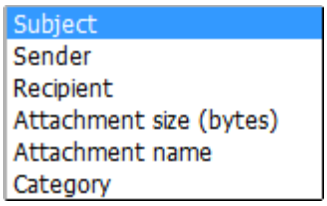
4.2.16 Advanced

The Advanced tab allows you to limit what is stored by Retain. This must be used with caution as this opens holes for data to be lost through. It is recommended to store everything since storage space is inexpensive.



Advanced Criteria

If you want to be more specific as to what to dredge or not to dredge, add the criteria here. Each line will be logically AND-ed together. Think “Dredge all items where the following is true:” Criteria A AND Criteria B AND Criteria C AND etc.



You may select based on:

- Subject
- Sender
- Recipient
- Attachment Size (in bytes)
- Attachment Name
- Category

And whether they are equal to, not equal to, contain or do not contain the item you specify.

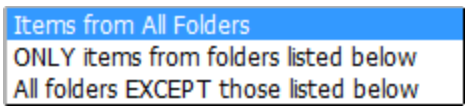
This gives you great flexibility and granularity. It allows you to customize dredges and retention for many different groups, or even individuals.

4.2.17 Folder Scope

By default, we dredge items from all folders. You can specify one or more inclusions or exclusions.

Your choices are:

- Dredge everything
- Dredge only these listed folders
- Dredge everything except these listed folders



How to specify the list of folders to dredge/exclude:

1. Specify a System Folder (mandatory). Example: Calendar.
2. You specify a subfolder of that folder (optional).

   Example: entering “old” would mean the folder “old” under “Calendar”.
3. You can have multiple hierarchies under that with the / delimiter.

   Example: “old/mail” would mean the subfolder “mail” under “old” under “Calendar”.
4. You specify if the option includes subfolder.

   Example: If you select “old” and “includes subfolder” is unchecked, “Calendar/mail” is selected. If “includes subfolder” is CHECKED, “Calendar/old/mail” would also be selected.

You may now configure Schedules, Workers and Jobs.

4.3 Exchange Module Office365

Retain supports Office 365.

4.3.1 Prerequisites

There are several prerequisites that need to be completed for setting up a job.

4.3.2 *Test Autodiscover*

If Autodiscover fails for a user, serious errors will occur when Retain attempts to archive them. This option allows users to be skipped or have the job immediately aborted.

The first thing to do is to test that autodiscover is enabled for the domain.

1. Go to the Microsoft Remote Connectivity Analyzer (https://testconnectivity.microsoft.com/)
2. On the Office365 tab, select "Outlook Autodiscover" from the Microsoft Office Outlook Connectivity Tests
3. Enter your credentials and run the test. If the test fails, you will need to call Microsoft to have them turn autodiscover on.

4.3.3 Get the SMTP URL

Retain needs to know what SMTP server to use for sending notifications.

Query the DNS for the MX record or

1. Go to the Microsoft Remote Connectivity Analyzer (https://testconnectivity.microsoft.com/)
2. On the Office365 tab, select "Inbound SMTP Email" from the Microsoft Office Outlook Connectivity Tests
3. Enter your credentials and run the test.
4. Record the SMTP server URL for Retain.

4.3.4 Set up DNS SRV Record

Microsoft has an article describing how to set up a DNS SRV record titled "A new feature is available that enables Outlook 2007 to use DNS Service Location (SRV) records to locate the Exchange Autodiscover service".

In short, the process is:

1. Go to the DNS Manager
2. Expand Forward Lookup Zones
3. Locate and right-click on the external DNS zone and choose Other New Records
4. Click Service Location (SRV) and enter:

   Service: _autodiscover

   Protocol: _tcp

Port Number: 443

Host: [your SMTP host URL]

5. Click OK
6. Check that this works by opening a command prompt and try to ping autodiscover.yourdomain. This would be one of the successful ones from the Test Autodiscover step above. This should return a Microsoft IP address. If ping cannot connect, there is something mis-configured in the DNS or a firewall blocking the connection.

For more information on Autodiscover, please see: http://support.gwava.com/kb/?View=entry&EntryID=2402

4.3.5 Set Rolling In-Place Hold

To prevent data loss, it is highly recommended that a rolling in-place hold be set so users are unable to remove items from disk before Retain has a chance to archive them.

Not all O365 licenses allow the setting of a hold, in such cases there is no way to prevent data loss.

In Retain, set Profile/Miscellaneous to Include user's recoverable items.

In Exchange Admin Console, set up a distribution list, for example All_Mailboxes, that contains all mailboxes. It will be best to create a policy to add new users to this distribution list by default.

Place the distribution list under a 90-day hold.

In the Exchange Management Shell:

An In-Place Hold can be set up for all mailboxes for 90 days:

New-MailboxSearch "Retain90DayHold" -ItemHoldPeriod 90 -InPlaceHoldEnabled $true -SourceMailboxes All_Mailboxes

It will take time for the hold to take effect. You can determine how many mailboxes were placed under hold with the script:

((Get-Mailbox).InPlaceHolds).Count

4.3.6 Setup Users

NOTE: Retain uses PowerShell to connect to O365. PowerShell does not allow the following special characters in names or passwords: # $ ( ) * + . [ ] ? \ / ^ { } |

4.3.7 Retain Impersonation User

In O365 Exchange Admin Center create the Retain Impersonation user with a mailbox, making sure to give it a license.

Give the Impersonation user the proper rights. Under Permissions, create a new Admin Role (e.g. Retain Impersonation Management), add the ApplicationImpersonation right and the Retain Impersonation user as a member.

4.3.8 Retain Administrator User

Retain needs a user with Administrator rights to download the address book from Office 365 every day with the Office365 Address Book Synchronization Script. This can be an existing administrator account or you may create a separate one. It needs to have sufficient rights to see all the users in the address book.

Shared Mailboxes, Rooms and Equipment

Impersonation rights allow the Retain user to enter other mailboxes but those rights do not extend to shared mailboxes. To access a shared mailbox the Retain user would need delegate rights to each shared mailbox that is to be archived. These rights can be granted through the Exchange Management Shell.

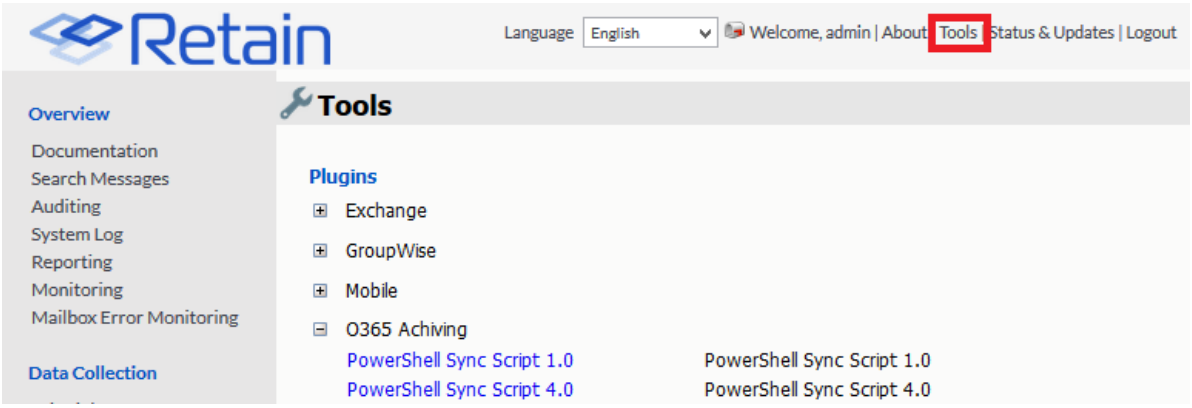
If the shared mailbox was John Doe's and your Retain impersonation account were "Retain", you would issue the following command in an Exchange Management Shell (EMC):

Add-MailboxPermission -Identity "John Doe" -User Retain -AccessRights FullAccess -InheritanceType All -AutoMapping $false

Address Book Synchronization

For Retain to authenticate users and access mailboxes for archiving, it needs to know what mailboxes are in Office365. To download the domain address book, Microsoft requires third-parties to use Powershell.

There are two scripts to do that both found under the Tools menu at the upper right of the Retain admin console.



The two scripts do the same thing: download the address book from Office365.

PowerShell Sync Script 1.0 saves the username and password in plaintext in the script.

PowerShell Sync Script 4.2 encrypts the password to a separate file.

Retain will not be able to archive members of the distribution lists if the field: HiddenFromAddressListsEnabled is True.

If using multiple modules, you will need to create separate folders for the script and the resulting CSV files. The folder location that Retain should pull the CSV file data from is set in the module as detailed below. You will also need to create a scheduled task for each script.

4.3.9 PowerShell Sync Script 1.0

If the hosted system in use is an Office 365 system, further configuration on the Office 365 system is required. Because Office 365 Exchange Servers do not directly expose LDAP information, a PowerShell script must run periodically to cache user and group information.

The script requires PowerShell 2.0. Without this script, you cannot run jobs on Office365, or authenticate users. This script connects to O365 and downloads the address book and distribution group lists into two address book CSV files, exchangeuser.csv and exchangegroup.csv.

4.3.10 Setup Process Overview

1. Install PowerShell 2.0 or higher (not needed for Windows 7 or Windows Server 2008 R2)
2. Install the Office 365 PowerShell cmdlets
3. Allow PowerShell script execution policy
4. Adapt and execute the “sync365.ps1” script provided by Micro Focus
5. Automate the script to run regularly.
6. Place the two resulting csv files in Retain

4.3.11 Setup

1. Install PowerShell 2.0 or higher. (Windows 7 and Server 2008 R2 already come with PowerShell 2.0)
2. Enable Microsoft .NET Framework 3.5.1.
3. Install the Office 365 PowerShell cmdlets. Two packages are needed which can be downloaded from the Microsoft Azure Active Directory PowerShell Module Version Release History (which is currently found at: http://social.technet.microsoft.com/wiki/contents/articles/28552.microsoft-azure-active-directory-powershell-module-version-release-history.aspx#Where_can_I_find_the_latest_version_of_AAD_PowerShell) page:

   Microsoft Online Services Sign-In Assistant for IT Professional RTW (this is the prerequisite to the Azure AD Module)

   Azure Active Directory Module for Windows PowerShell (64-bit version)
4. You may need to change the execution policy to allow these scripts to function:

   Allow PowerShell script execution

   The Default Execution Policy is set to restricted, it can be viewed by entering this command in PowerShell:

   Get-ExecutionPolicy

   The script provided by Micro Focus must be run in PowerShell while running under Administrator rights:

   Set-ExecutionPolicy RemoteSigned
5. Download the O365 Archiving Powershell Sync Script, found on the tools page. To access the tools page, select the 'tools' link from the top right of the Retain Server administration page.
6. If using the script on the Retain server itself: Move the “sync365.ps1” script to the "~\Program Files\Beginifinite\Retain\RetainServer\WEB-INF\cfg" folder. Otherwise, make sure that the exchangeuser.csv and excahngegroup.csv end up in that directory on the Retain server.
7. Edit the “sync365.ps1” script with Microsoft Integrated Scripting Environment (ISE) editor.

   a. At the top you will find 3 settings: $User, $PlainPassword, $ExportBasePath:

Set $User to the UPN of an administrator account in Office 365.

Set $PlainPassword to the plain text password of the administrator account.

Set $ExportBasePath to a directory where the two resulting csv files will be saved. If the path does not yet exist, it must be created manually as the directory will not be created automatically.

Keep in mind that your backslashes must be quoted, and a final backslash is not needed.

Example:

$ExportBasePath="C:\\Program Files\\Beginifinite\\Retain\\RetainServer\\WEB-INF\\cfg"

NOTE: If it is unacceptable to set a password in clear text, please view the following solution or use the secure script below:

---

http://bsonposh.com/archives/338 -presents an alternative, which loads and saves encrypted credentials from file.

b. Execute the script by clicking the play button. This process can take a while if there are many users. Once the script has completed, a message indicating that the script is done should be displayed along the bottom status bar. The resulting files, exchangeuser.csv and exchangegroup.csv, must be placed in the "~\Program Files\Beginifinite\Retain\RetainServer\WEB-INF\\cfg" folder for Retain to access the address book.

8. Set Task Scheduler to run the script automatically once per day. If you create it at the Task Scheduler (Local) level you will be able to find it in the Task Scheduler Library folder after it is created, in the center pane.

   a. Create a New Task

   b. On the General tab, give it a name; and, for the description, you may want it read something like "Retain address book synchronization script for Office365".

   c. Under "Security options" section, choose: “Run whether user is logged in or not”.

   d. Under the Triggers tab, click on "New..."

   e. Under "Settings", set it to "Daily" and set the "Start" time: [set it to run an hour before the archive job is set to begin, but it doesn’t take that long, Exchange does maintenance between 1-5 AM]

   f. Do not expire

   g. Check Enabled

   h. Under the Actions tab: Create a New action.

   i. Set the Action to “Start a program”

   ii. Program/script: powershell

   iii. Add arguments: -NoProfile -ExecutionPolicy Bypass -file "[drive]:\Program Files\Beginfinite\Retain\RetainServer\WEB_INF\cfg\sync365.ps1" -Verb RunAs

   iv. Start In: (leave blank)

4.3.12 PowerShell Sync Script 4.2

The standard sync365.ps1 script requires entering a plaintext administrator password. If that is not acceptable by your security policy, there is the O365 Archiving PowerShell Sync Secure Storage Script: Save-CredentialsEncrypted.ps1 which allows you to securely save the password used by the admin user.

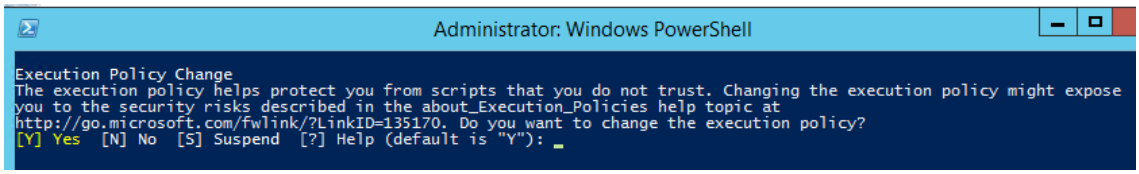
Requirements

- The script can be found under O365 archiving under the Tools menu at the top-right of the Retain admin console.
- This must be run on the same machine as it was created on as the encryption is dependent on the machine it was run on.
- This script requires PowerShell 4.0 or higher or the script will abort.
- If using Windows Server 2008R2 or earlier, the script will throw errors because the Task Scheduler cmdlets are not supported and the Scheduled Task must be created manually.

Running the Script

1. After downloading and extracting the script, open PowerShell and change to the directory the script is in and type *.\Save-CredentialsEncrypted.ps1*

   1a. If you haven't run a PowerShell script before, you may have to change the Execution Policy to get a script to run. Set the execution policy to *Set-ExecutionPolicy -ExecutionPolicy Bypass -Scope Process*. If PowerShell 4.0 or higher is not installed the script will abort.



2. A dialog box will appear requesting the credentials for an administrator user and the destination folder of the address book files.

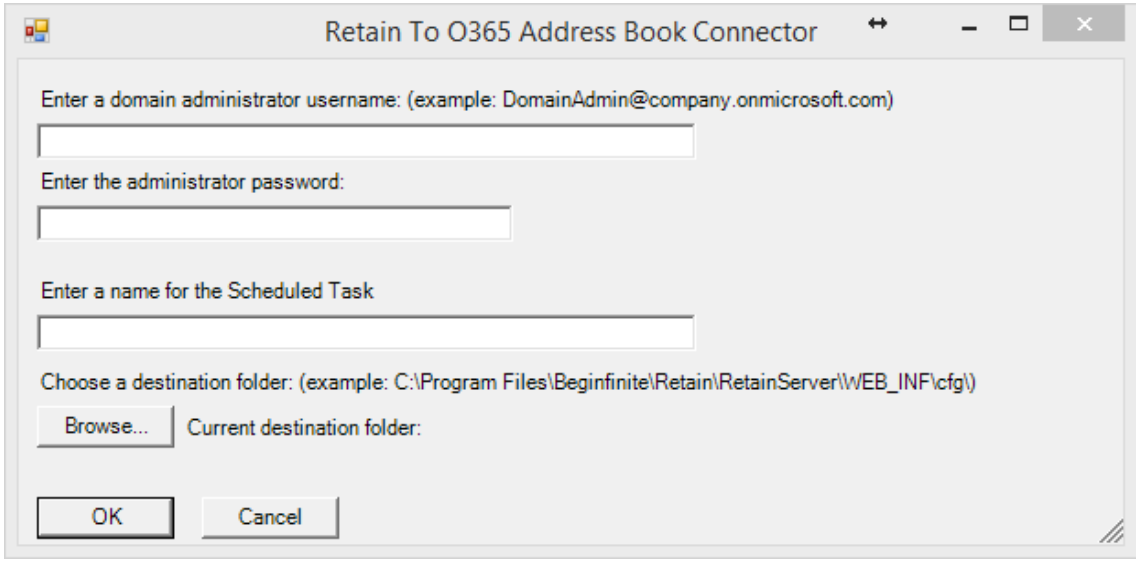
   The administrator user must have sufficient rights to download the address book from O365. This would be the logon name and password of an administrator user (e.g. DomainAdmin@company.onmicrosoft.com).

   Provide it with a name for the Scheduled task.

   You can set the destination folder to anywhere you wish. However, the address book files exchangeuser.csv and exchangegroup.csv must end up in the ~\Program Files\Beginfinite\Retain\RetainServer\WEB_INF\cfg\ folder of the Retain Server. Retain will look for the address book files there when refreshing the address book.

   If you are using a Linux-based Retain Server, you will need to set up a process to move the address book files to the Linux-based Retain Server's opt/beginfinite/retain/RetainServer/WEB-INF/cfg folder.

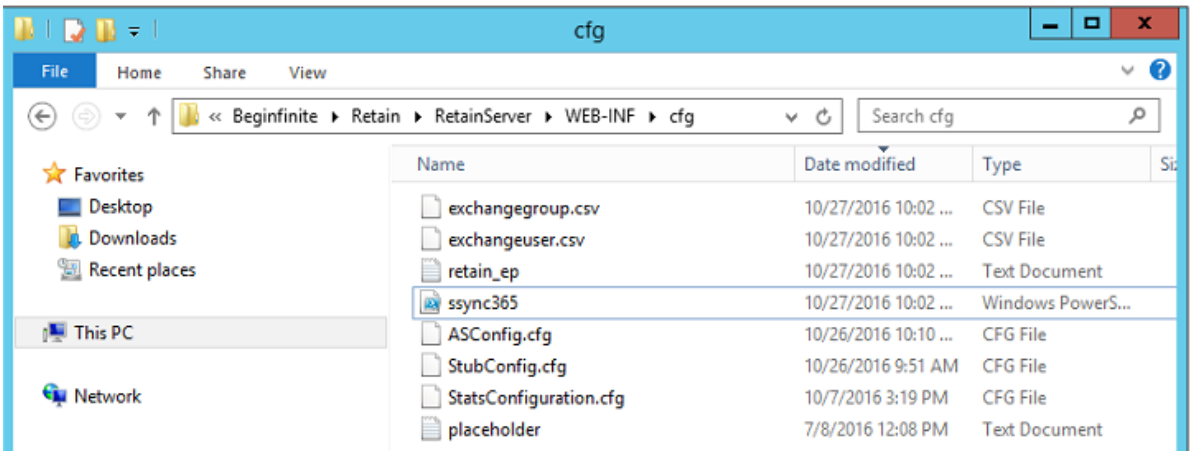
   If you are running Retain will multiple modules each module will need a separate task and destination folder.



3. The script will then setup a recurring task (SyncO365) in Task Scheduler to download the address book every day at 12:30am.

   The script will request your logon credentials as it sets up the task. The script will start the task before exiting. If you are using Windows Server 2008, multiple errors will appear because the Task Scheduler cmdlets do not exist, this is expected and the task can be created manually in Task Scheduler.

4. It will take a moment to download the address book files: *exchangeuser.csv* and *exchangegroup.csv*. For a small system (<100 users) it may take a few minutes, for a large system (>10k users) it may take more than half an hour. After the script completes, make sure that the address book files are filled with items. If this does not generate the files it may be an execution policy issue, see notes section below.



Notes

We have found that Task Scheduler will sometimes stop working, so we recommend monitoring the .CSV files that they are updated every day.

Blank .CSV files:

Office365 requires regular password changes. If the wrong credentials are entered or the password has expired two blank .CSV files will be created. You will need to run this script again entering the Administrator logon name, the new password and the destination folder. The task will be started by the script to update the address book files.

Common Errors:

If you see red text and the window closes immediately that means, there was an error of some kind. Open a PowerShell window, change to the script folder, and run the script manually:

.\Save-CredentialsEncrypted.ps1

It may be that the proper version of PowerShell is not installed. This script requires at least version 4.0. You can determine the installed PowerShell version by running the cmdlet:

$PSVersionTable.PSVersion

To change the execution policy manually, the following command in a PowerShell window will allow execution for the current process:

Set-ExecutionPolicy -ExecutionPolicy Bypass -Scope Process

Then run the script from the same process window.

4.3.13 Configure SharePoint Online/OneDrive for Archiving

In Office365 you can attach documents stored in SharePoint Online/OneDrive. However it doesn't send the actual document but only a link to that document.

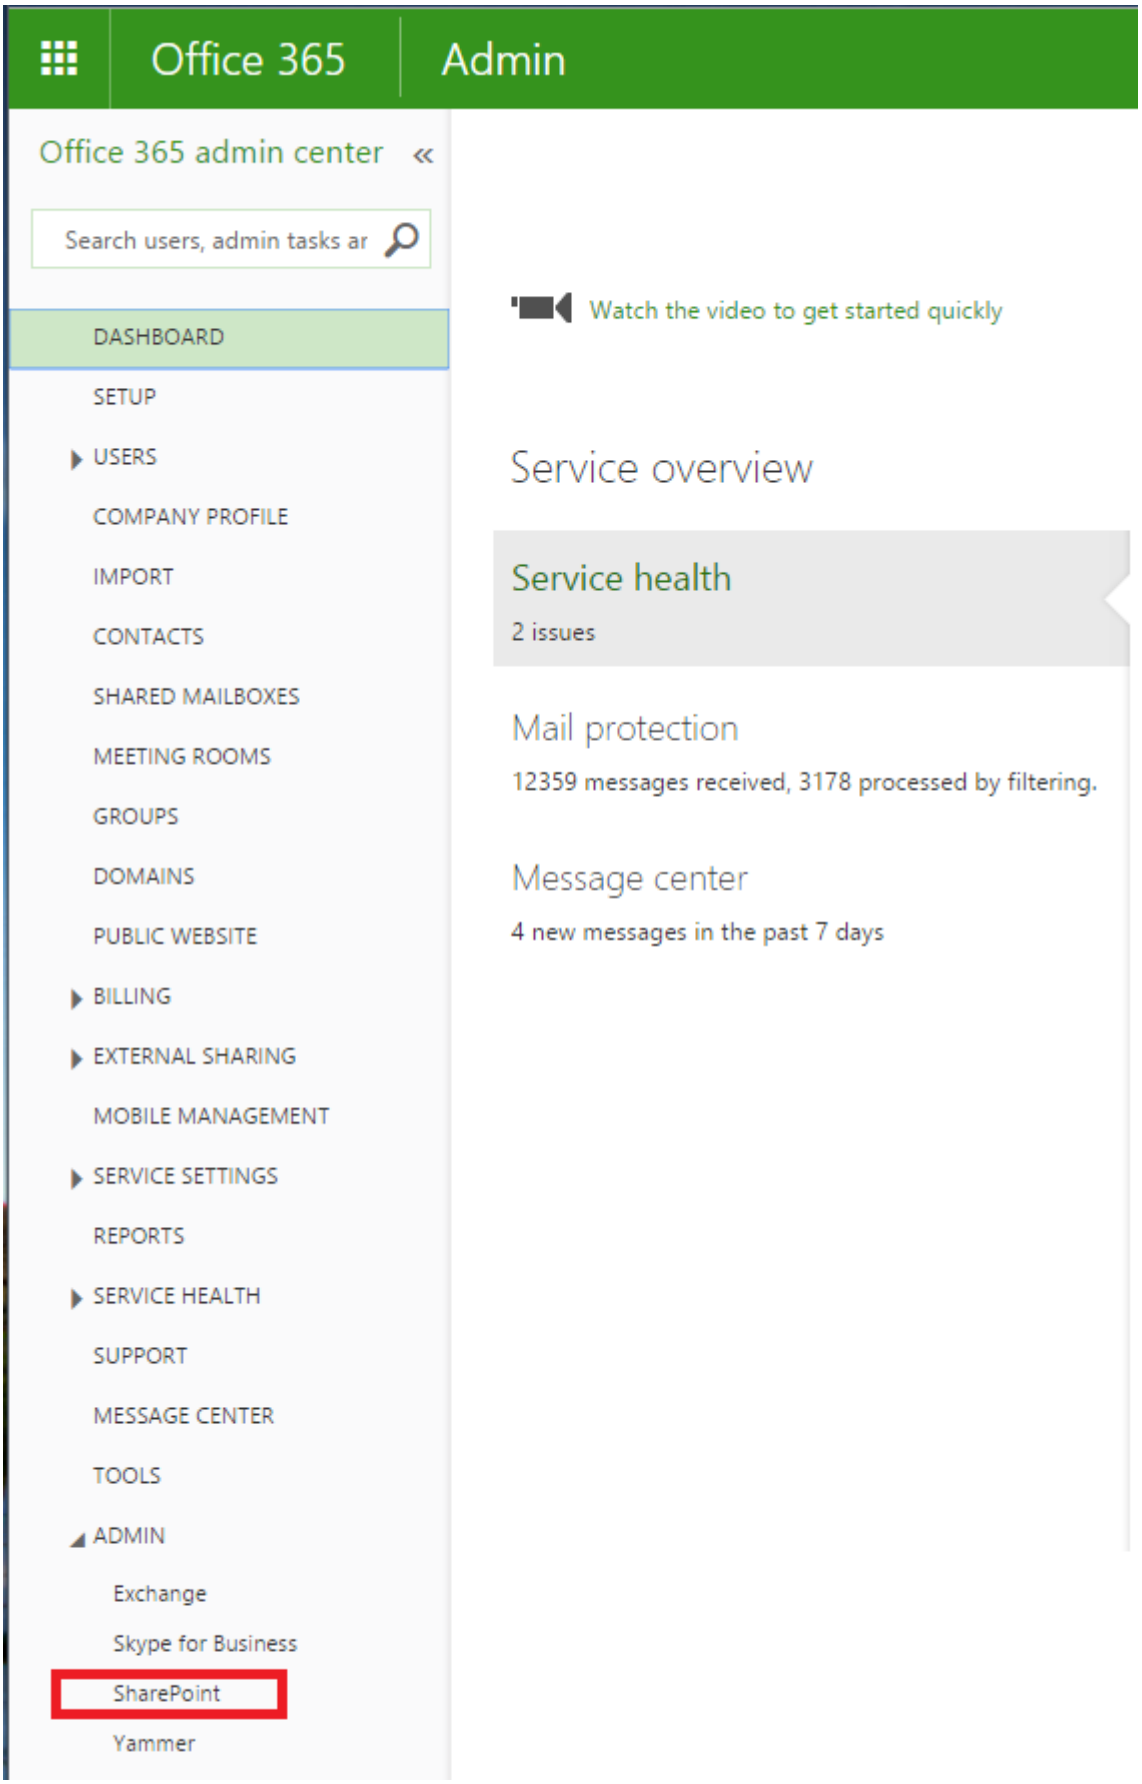
NOTE: As a cloud service these screens may change at anytime.

When Retain attempts to archive an attached document it often runs into trouble because the retain user, while it does have ApplicationImpersonation right, doesn't have rights to anything outside of Exchange. So Retain will error because it knows there is supposed to be something there but can't reach it.

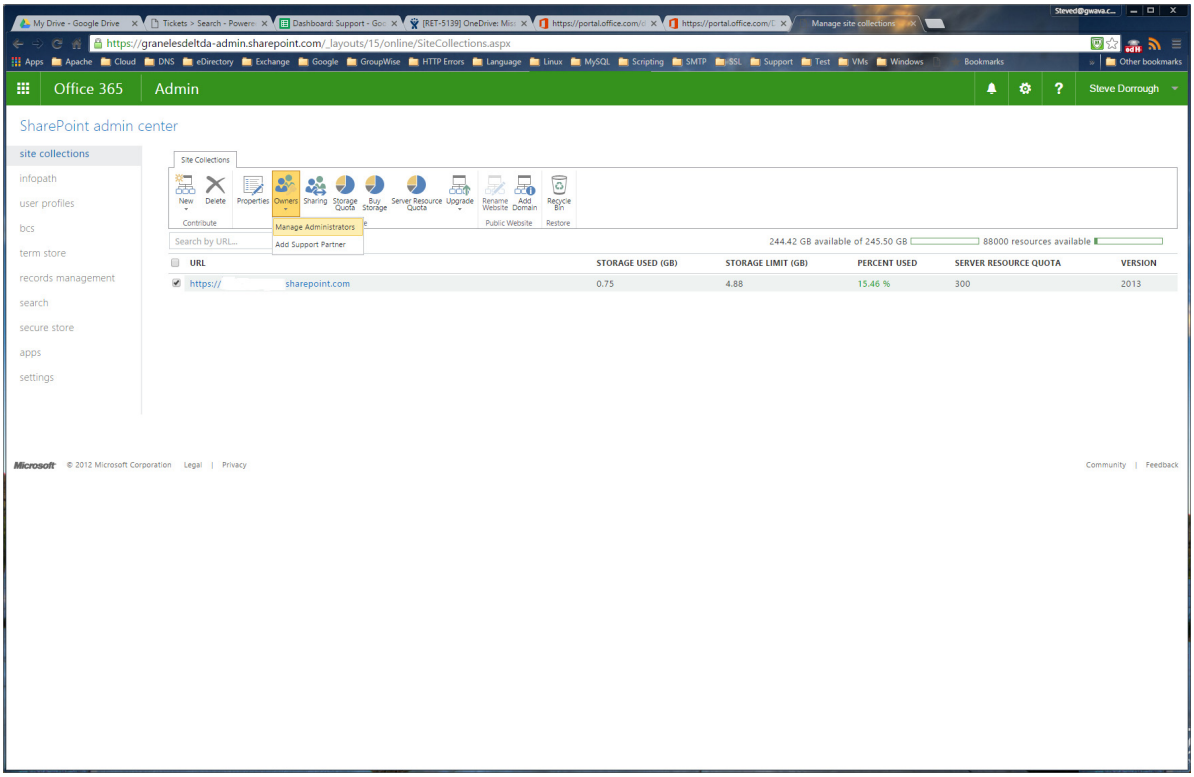
11:04:16, 704[Thread-4920] [ERROR] ExchangeAttachment: error while creating attachment.

java.io.IOException: SharePointError - Impersonation has no access to: https://gwava-my.sharepoint.com/personal/user08_gwava_onmicrosoft_com1/Documents/Email attachments/office 365 users(1) (1).txt

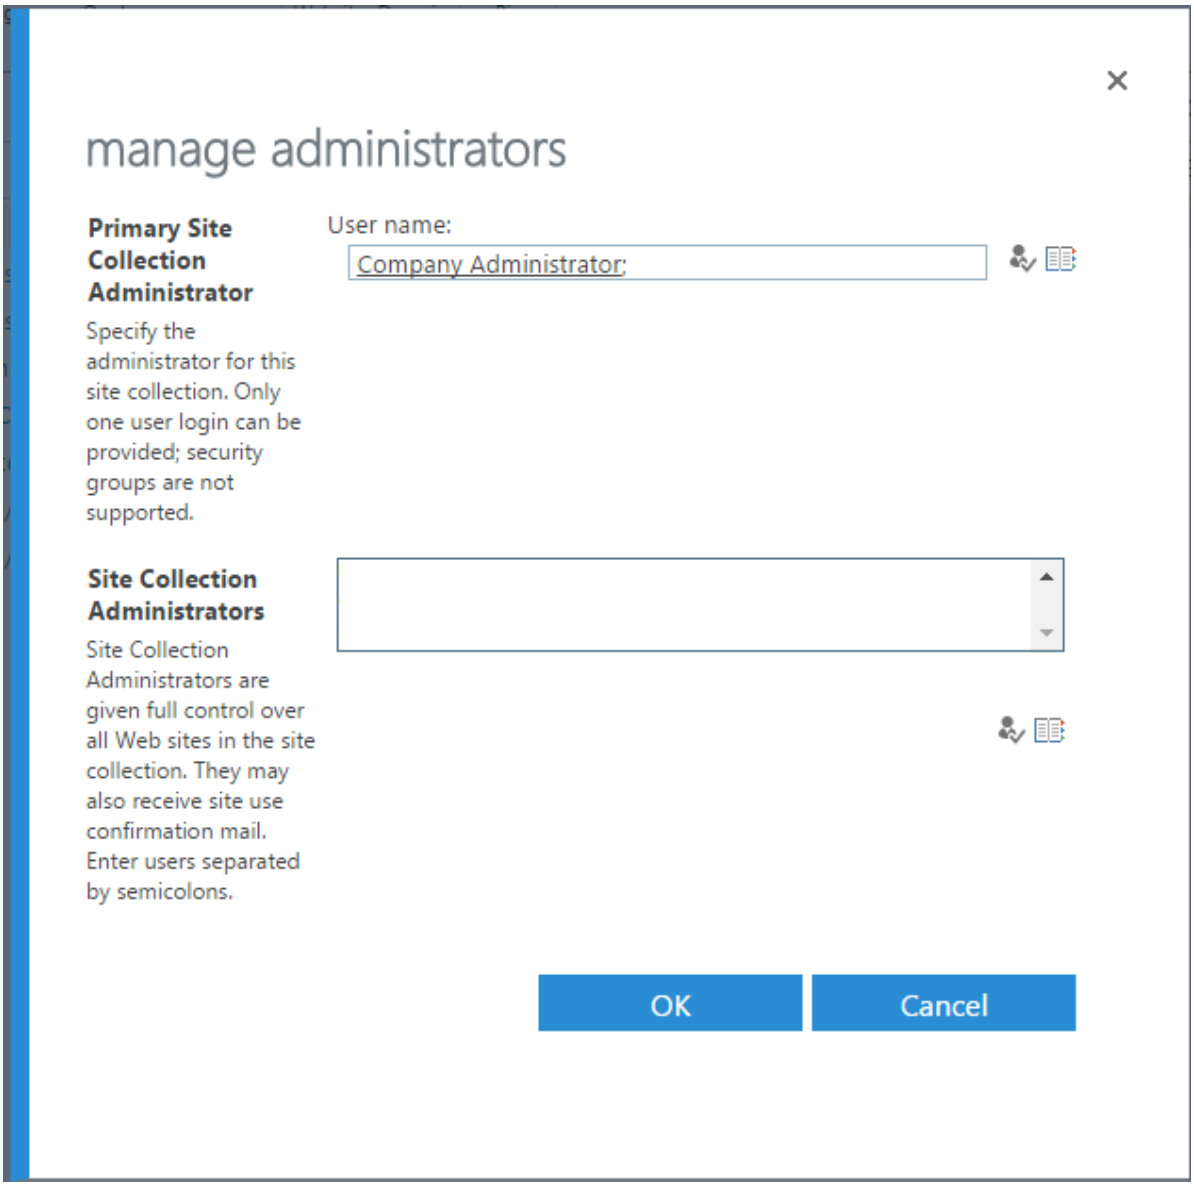
What needs to be done is the retain user needs be added as a "site collection administrator" in SharePoint/OneDrive. Start by going to the Office365 admin page and selecting SharePoint.


Office 365
Admin
Office 365 admin center
Search users, admin tasks ar
DASHBOARD
SETUP
USERS
COMPANY PROFILE
IMPORT
CONTACTS
SHARED MAILBOXES
MEETING ROOMS
GROUPS
DOMAINS
PUBLIC WEBSITE
BILLING
EXTERNAL SHARING
MOBILE MANAGEMENT
SERVICE SETTINGS
REPORTS
SERVICE HEALTH
SUPPORT
MESSAGE CENTER
TOOLS
ADMIN
Exchange
Skype for Business
SharePoint
Yammer
Watch the video to get started quickly
Service overview
Service health
2 issues
Mail protection
12359 messages received, 3178 processed by filtering.
Message center
4 new messages in the past 7 days

Go to the Site Collections, select the collection to add the retain user too, click the “Owners” tab.



Site collection administrator: You can now add the retain user as a “”. This way the retain user will have rights to access the data.



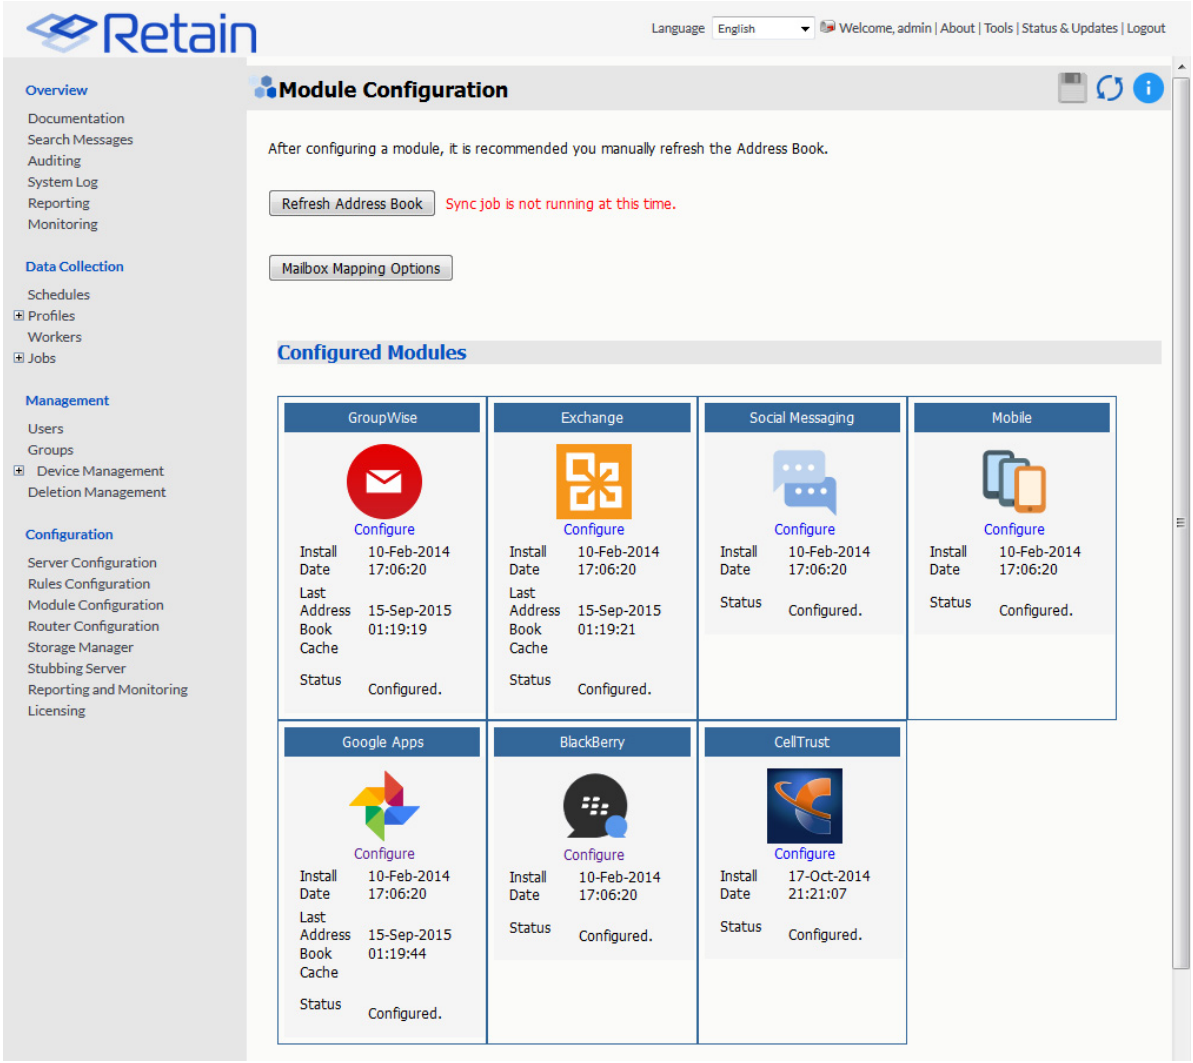
This works for individual site collections, but it would be difficult to do that individually if you have many site collections. There is no way in SharePoint to change more than one collection at a time.

Changing Across The Domain

However, there is a SharePoint Online Management Shell script that you could modify for your use to do all site collections at once.

4.3.14 Configure Exchange Module for Office365

The Exchange module must be configured in the Retain Server before any communication between Retain and an existing Exchange message system can occur. Open the Retain management page on the Retain Server, and select Module Configuration.



Select the ‘Configure’ option in the Exchange module. A new window or tab will open with the module configuration.

NOTE: Ensure that your Retain Server DNS is set to the same DNS server that your Exchange server uses. The Exchange module uses these DNS setting to auto discover critical information about Exchange stored in Active Directory and will not function correctly unless both systems are pointed to the same DNS server.

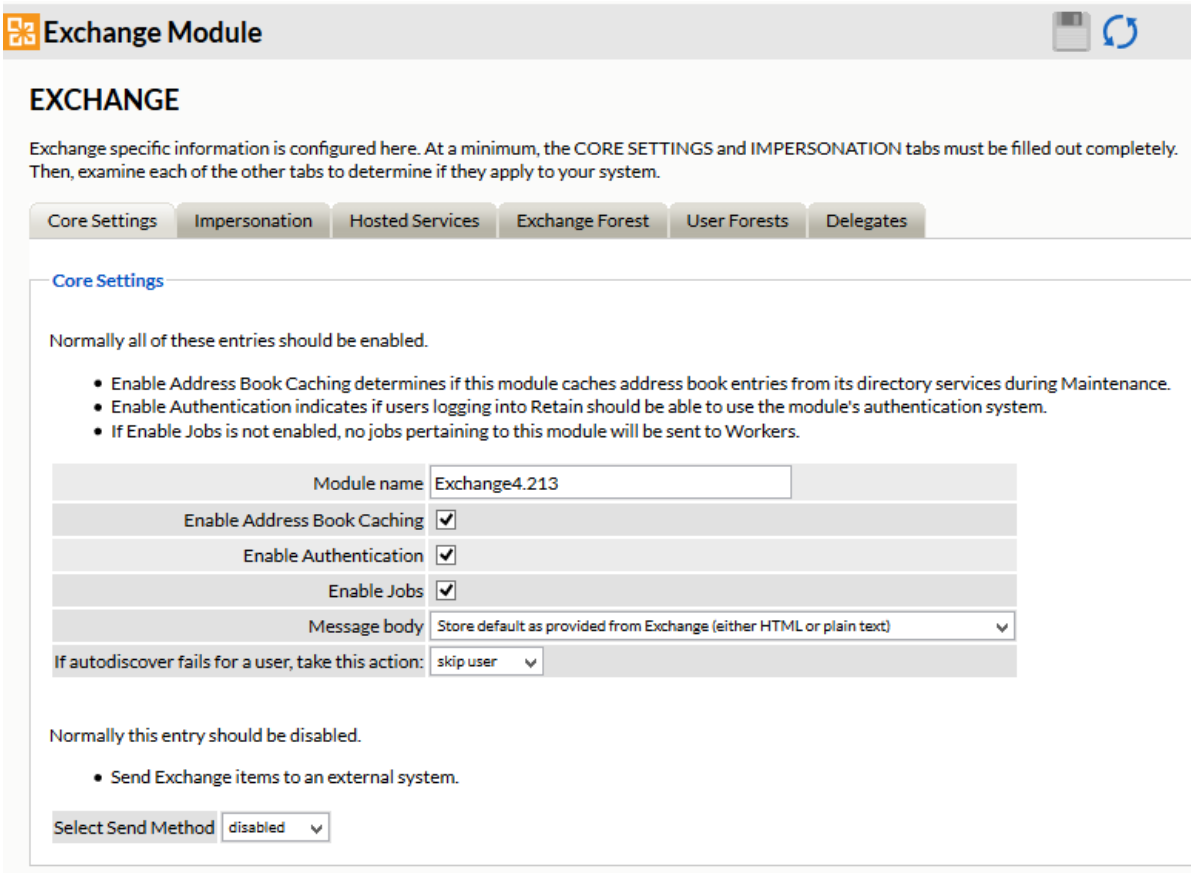
4.3.15 Core Settings Tab

The Core Setting Tab allows you to disable all jobs and disable users logging into Retain.

The module needs to be enabled on this page to make it active in the Retain system.

The module can be given a name.

The Send Method option enables either the SMTP Forwarding or FTP features. For either feature to appear and function, the Module Forwarding tab must be configured on the Server Configuration page. See that section for more information.



Normally all the checkboxes on this tab are always left selected. It is rare that you would ever deselect any of them. Two cases where you might, would be: troubleshooting, (as instructed by Technical Support), or retrieving an old email system.

The Enable Address Book Caching function allows Retain to regularly cache the online email systems address book and synchronize it with Retain. This is critical for administration, authentication, and archiving purposes. It is recommended to cache the Address Book once every 24 hours to keep the Retain storage system up to date. By default, maintenance is set to cache the Address Book once every 24 hours.

The Enable Authentication checkbox determines if end-user authentication is performed when the user logs into Retain. If it is deselected, the Retain system will NOT authenticate the user against the email sys-tem and the user will not be able to log in unless another authentication method is enabled.

The Enable Jobs checkbox determines if configured data retrieval jobs are ever passed to the Worker. Even if the individual job is fully configured and enabled, if this checkbox is switched off, no jobs config-ured for this module will be processed.

The Message body allows the administrator to decide whether to store either the HTML or plain text message body, or both.

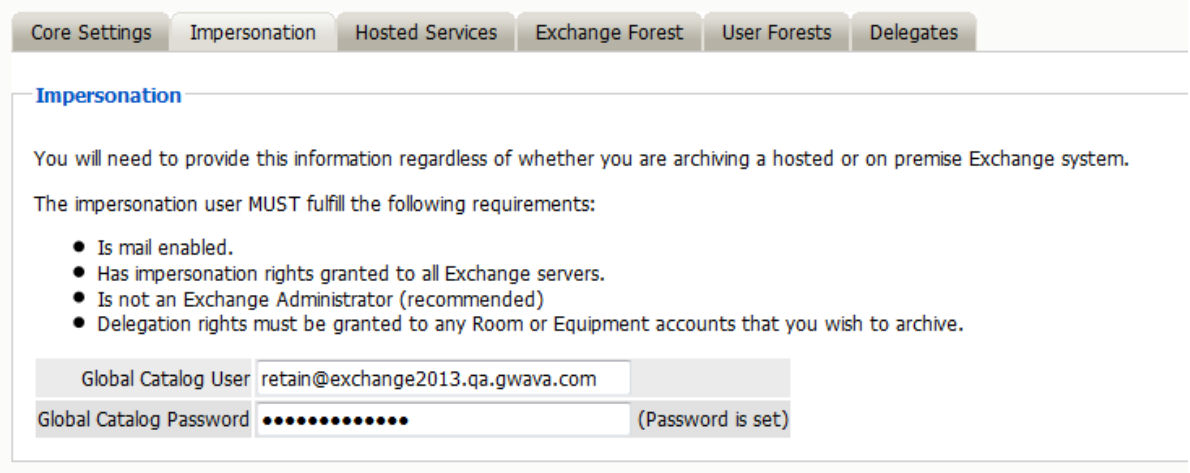
Send Method

The Send Method option enables either the SMTP Forwarding or FTP features. For either feature to appear and function, the Module Forwarding tab must be configured on the Server Configuration page. See that section for more information.

4.3.16 Impersonation Tab

If the Impersonation and Core Settings tabs are not completely configured with the correct information, the hosted system will not be archived correctly.

Enter the Impersonation user credentials.

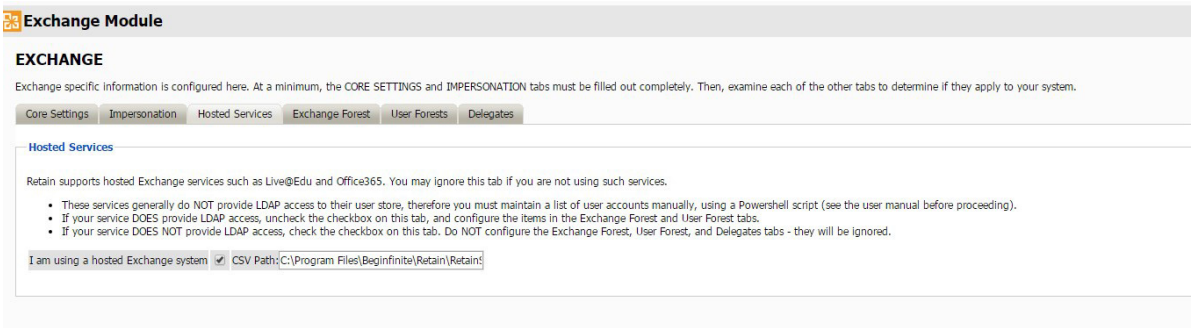


4.3.17 Hosted Services Tab

Mark the checkbox to enable Retain support for hosted Exchange services.

Provide the path that Retain will use for the CSV files that the PowerShell Sync script will save the user list. By default: C:\Program Files\Beginfinite\Retain\RetainServer\WEB-INF\cfg

All other tabs will be ignored. Save all changes before closing the Exchange Module page.



Multiple Modules

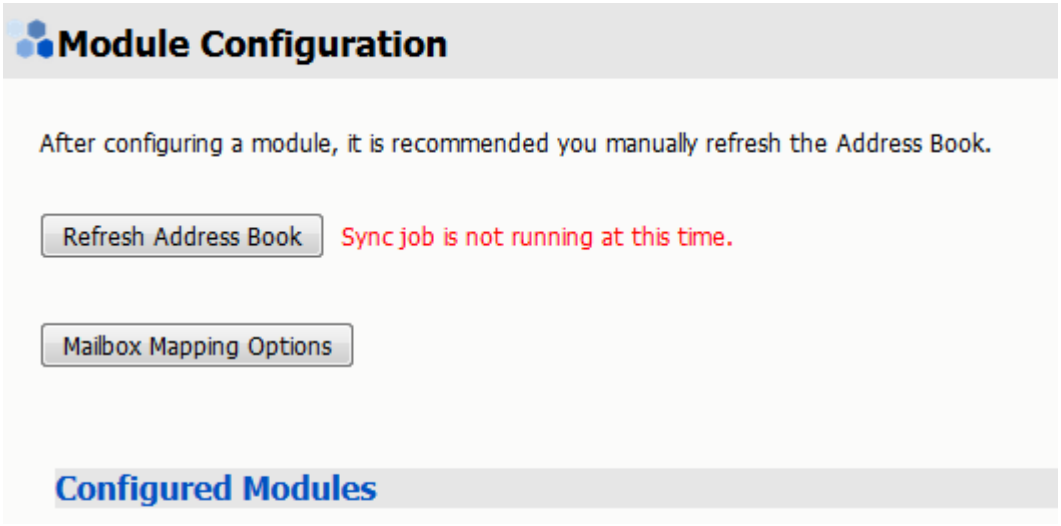
must: If using multiple modules, each module have a different folder for the script and CSV files.

For example, C:\Program Files\Beginfinite\Retain\RetainServer\WEB-INF\cfg\domain1 C:\Program Files\Beginfinite\Retain\RetainServer\WEB-INF\cfg\domain2 and *C:\Program Files\Beginfinite\Retain\RetainServer\WEB-INF\cfg\domain2*

Each module will also need a separate Scheduled Task.

4.3.18 Refresh Address Book

After saving changes, return to the Retain Server's Module Configuration page, and trigger a refresh of the Address Book.



Depending on the size of the address book, it may take several minutes to return with information, but a successful configuration will return a correct address book cache date and no errors. The date should reflect the date of when the address book refresh was triggered.

The Status may show “Address Book Cache Never Run” or may list commonly mis-configured or missed items if the Refresh job fails.

Once the status is configured and the Address Book has been cached, Retain can connect to and archive messages from the Exchange server. The system is ready to have workers, schedules, profiles, and jobs configured, and those options will now appear on the main administrative interface.

4.4 Exchange Profile

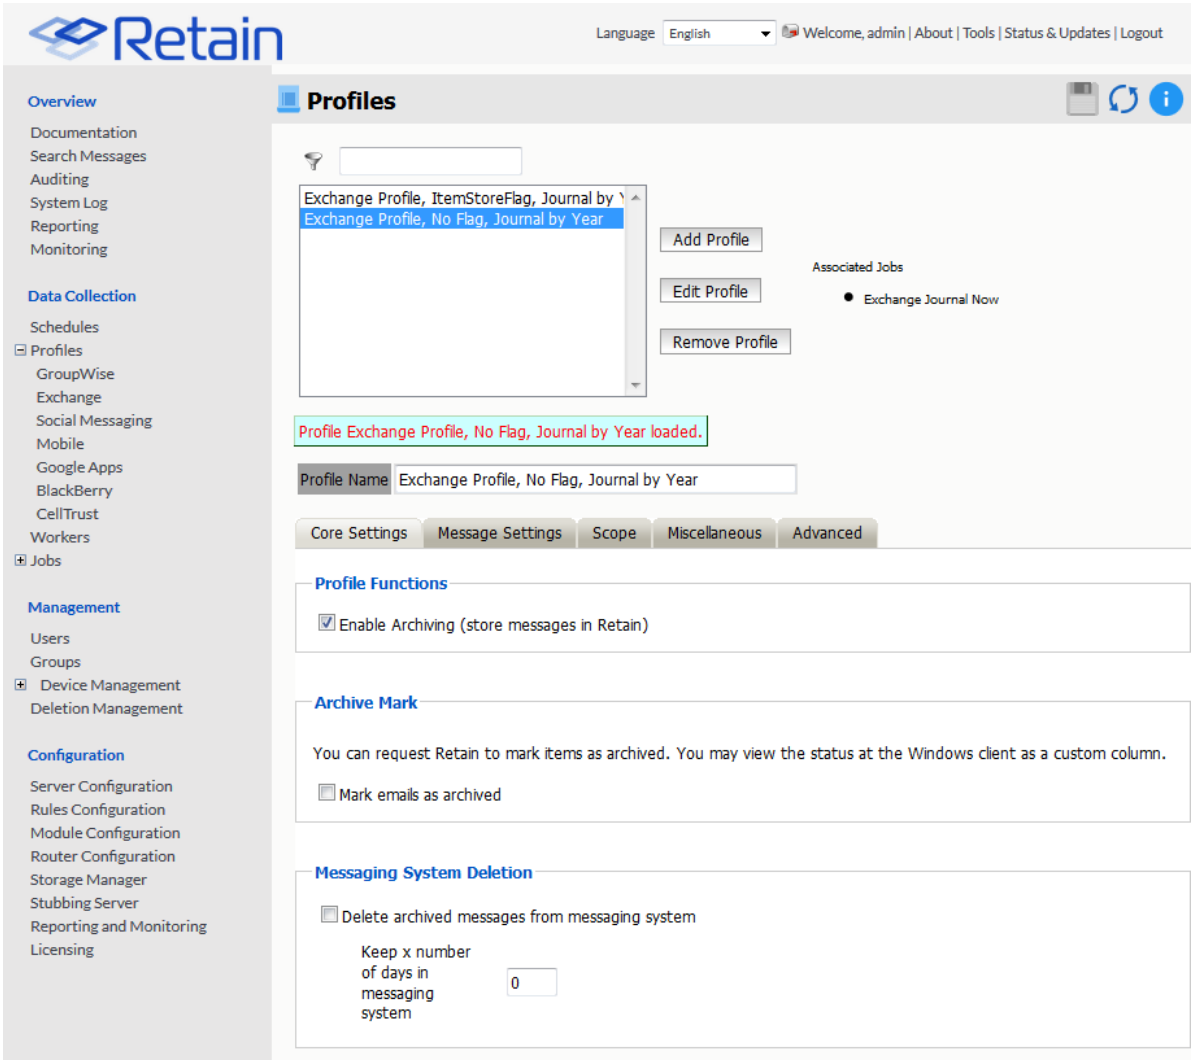
The job will need an Exchange profile setup to connect to the email system properly.

After the Exchange Module has been configured, the Exchange Profile will be available for configuration. If an Exchange Profile is not configured, jobs cannot be run against the Exchange system.

Click on “Add Profile” and provide a profile name, or select an already existing profile to access the con-figuration tabs. All changes made on this page must be saved by selecting the “save changes”, disk icon, at the top right of the page. Tabs may be changed and navigated through without affecting new settings, but any move to another page will require saving, or abandoning the changes made.

4.4.1 Core Settings

The core settings consist of an enabled/disabled option which must be enabled for any jobs based on this profile to archive anything.

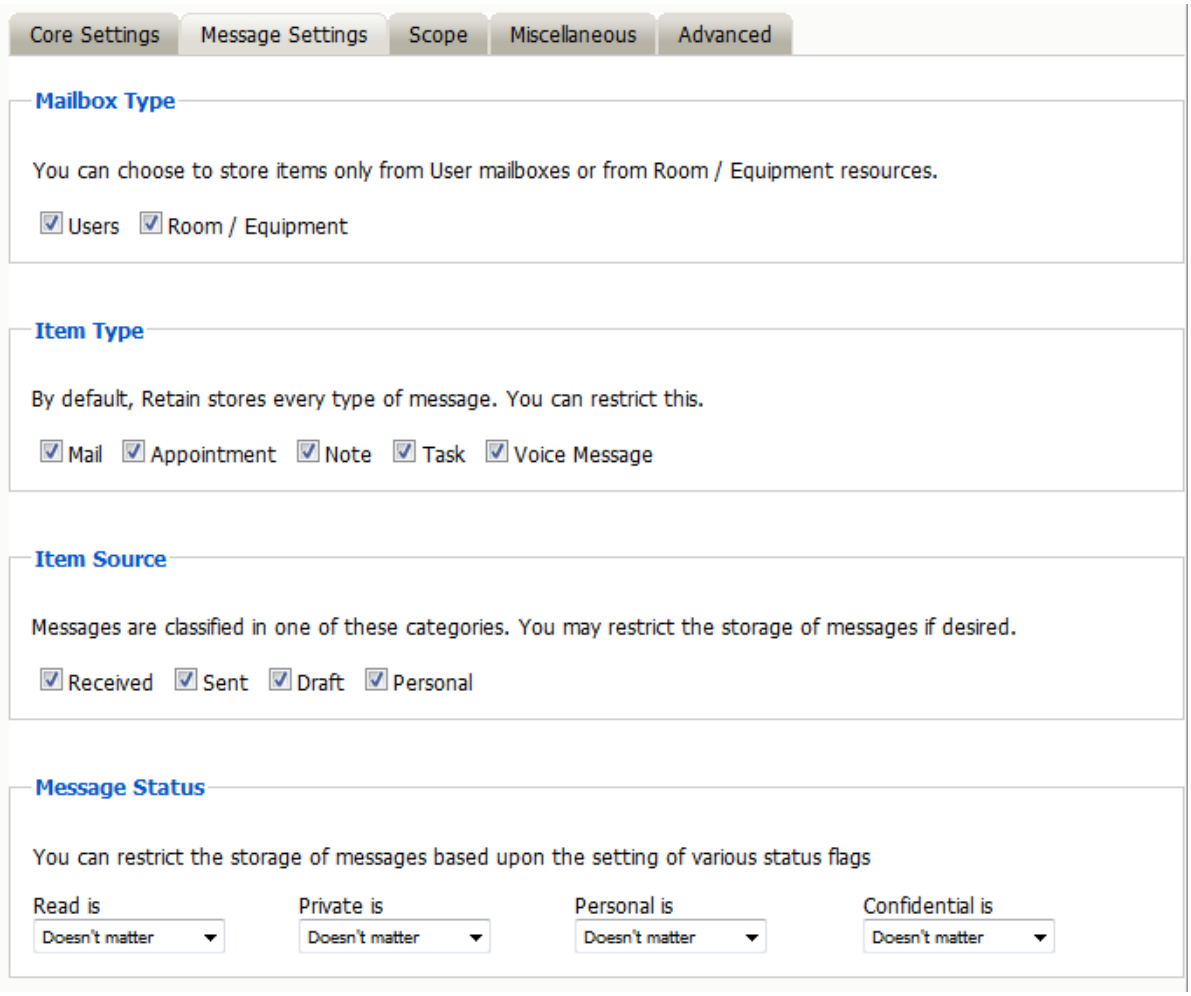


4.4.2 Messaging System Deletion

For systems where the administrator wishes to have archived messages removed from the system auto-matically, the Messaging System Deletion option may be used. Messaging System Deletion will remove messages from a mailbox after they are archived, according to the time frame specified in the settings. The amount of time to keep messages is specified in days. The recommended setting depends on the archiving scheme in the system. For example, if messages are to persist in the system for 30 days, then the system deletion setting should be set to 30 and enabled. A setting of 0 will remove messages from the system as soon as they are archived. Be sure to configure the system before enabling the setting in the profile.

4.4.3 Message Settings

Retain can archive and select specific types of mail and Exchange system items to be archived. The Man-age Settings tab provides access to manage those settings.

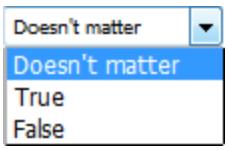


The Mailbox type specifies whether to include or exclude the available types of mailboxes. Because there can be multiple profiles and jobs, it may be advantageous to archive the Users and Room / Equipment mailboxes separately as needed and appropriate for the system.

The Item Type option specifies the different types of messages found in Exchange that can be archived, and allows the exclusion of or inclusion of the different individual types.

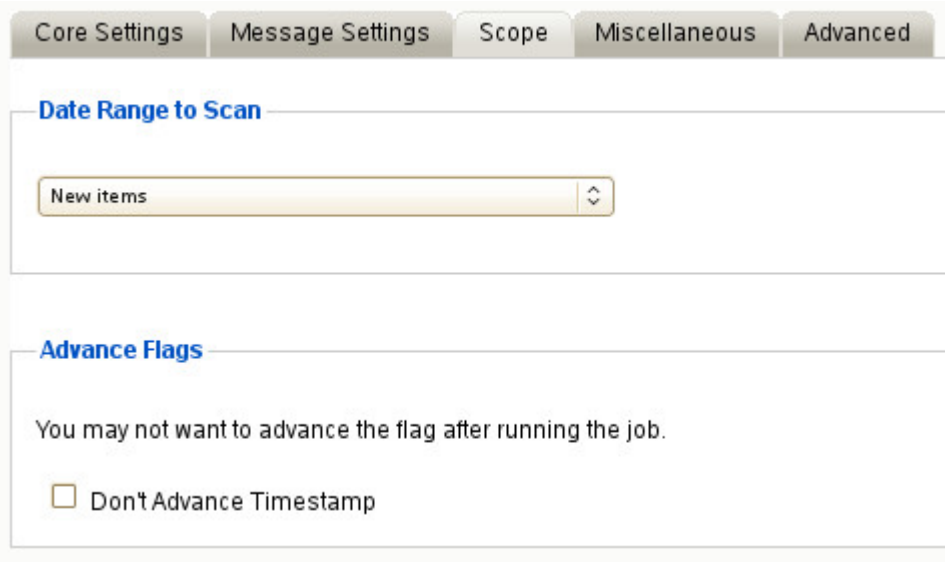
The Item Source option allows administrators to exclude or include messages that have not yet been sent or received, or posted.

The Message Status allows messages which have or have not been read or opened, or marked private or confidential to be archived. The different options in the drop-down menu are as shown.



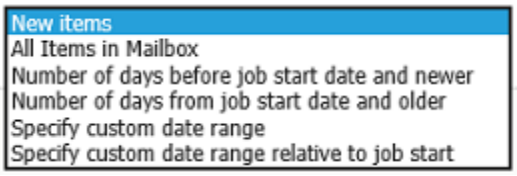
4.4.4 Scope

The Scope tab dictates the date range Retain will scan in the attached archiving jobs.



Date Range to Scan

The Date Range to Scan instructs Retain to scan for, and archive, messages after, or before, a certain date. This is useful if only specific chunks or areas of mail are to be archived.



New Items: All items that have not been archived by Retain since the last time the job ran.

All Items in Mailbox: All items in the mailbox starting from 1/1/1970, duplicates will be processed but not stored if they already exist in Retain's archive.

Number of days before job start date and newer: Only items from the relative number of days from the time the job began will be archived. E.g. messages that came into the email system less than 7 days ago.

Number of days from job start date and older: Only items previous to the relative number of days from the time the job began will be archived. E.g. messages that came into the email system more than 7 days ago.

Specify custom date range: Only items between two absolute dates will be dredged.

Specify custom date range relative to job start: Only items between two relative dates will be dredged. E.g. messages that came into the email system between 7 and 5 days ago.

It is recommended to archive all New items.

Advance Flags

Enabling "Don't Advance Timestamp" will not update the timestamp flag. Items that are dredged will still be considered new by Retain the next time the job runs.

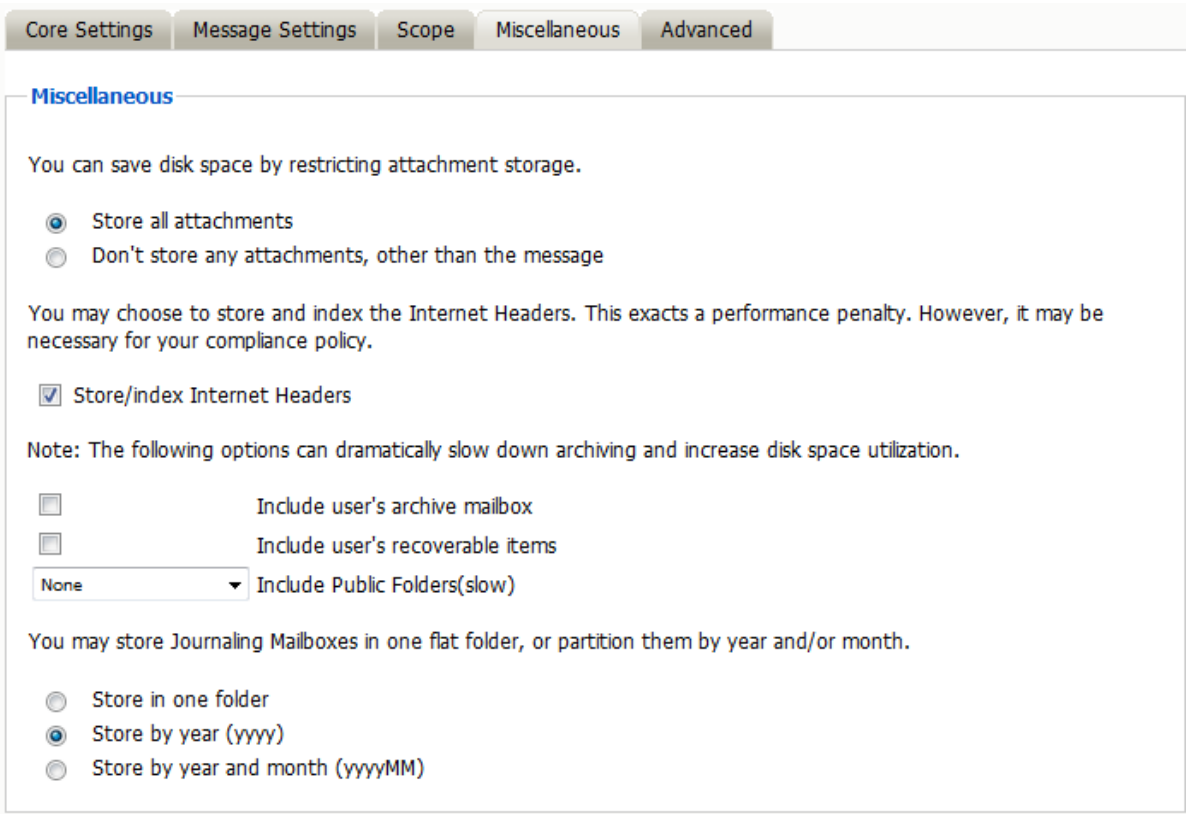
This is useful when troubleshooting, but is generally not used for normal jobs.

NOTE: Unlike GroupWise, Exchange does not ensure any compliance when scanning end user mailboxes; users may freely delete their email. The Item store flag does not prevent mail deletion. Only setting a rolling hold on all mailboxes guarantees all items have been archived.

4.4.5 Miscellaneous

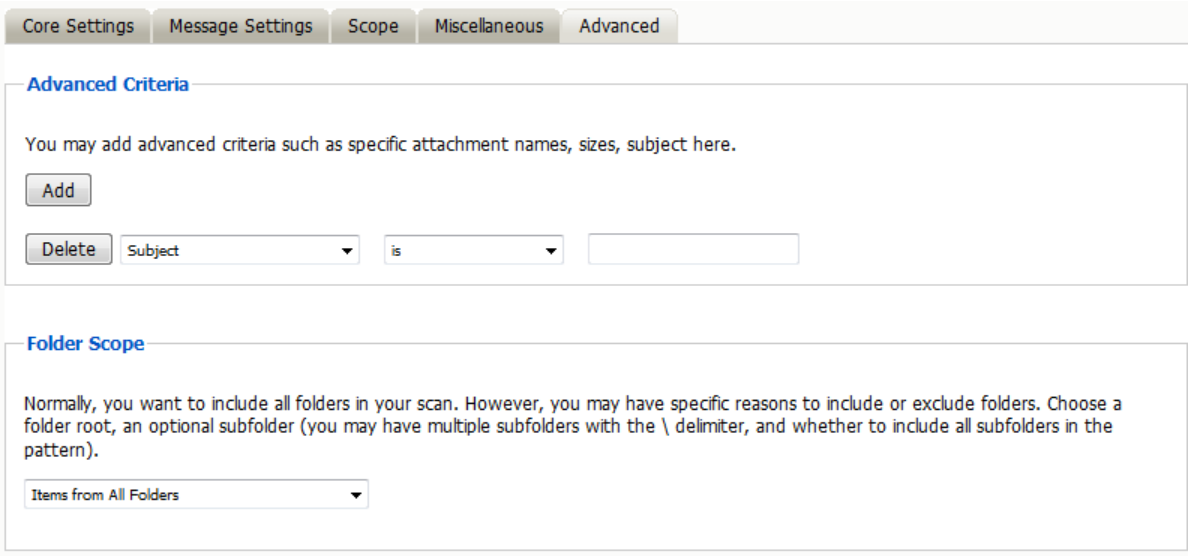
The Miscellaneous tab allows access to settings detailing how messages are stored and what is archived. Attachments, message information such as the internet headers, and how the data is stored and named, (by folders, year, or year and month), dictate not only the message store structure, but affect the storage size.

Miscellaneous options also allow for the archiving of the ‘recoverable items’. To enable checking and archiving of the ‘Recoverable Items’ for compliance reasons, select the checkbox next to the option.



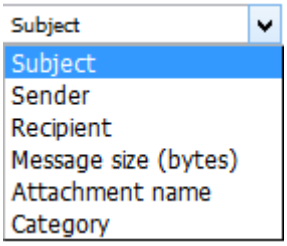
4.4.6 Advanced

The Advanced tab allows you to limit what is stored by Retain. This must be used with caution as this opens holes for data to be lost through. It is recommended to store everything since storage space is inexpensive.



Advanced Criteria

If you want to be more specific as to what to dredge or not to dredge, add the criteria here.



Each line will be logically AND-ed together. Think “Dredge all items where the following is true:”Criteria A AND Criteria B AND Criteria C AND etc.

You may select based on:

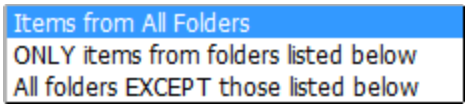
- Subject
- Sender
- Recipient
- Message size (in bytes)
- Attachment name
- Category

is is not contains, does not contain: And whether they are , *is not*, *contains,* or *does not contain* the items specified.

This provides flexibility and granularity. It allows customized dredges and retention for many different groups, or even individuals.

4.4.7 Folder Scope

By default, we dredge items from all folders. You can specify one or more inclusions or exclusions.



Your choices are:

- Dredge everything
- Dredge only these listed folders
- Dredge everything except these listed folders

How to specify the list of folders to dredge/exclude:

1. Specify a System Folder (mandatory). Example: Calendar.
2. You specify a subfolder of that folder (optional).

   Example: entering “old” would mean the folder “old” under “Calendar”.
3. You can have multiple hierarchies under that with the / delimiter.

   Example: “old/mail” would mean the subfolder “mail” under “old” under “Calendar”.
4. You specify if the option includes subfolder.

   Example: If you select “old” and “includes subfolder” is unchecked, “Calendar/mail” is selected. If “includes subfolder” is CHECKED, “Calendar/old/mail” would also be selected.

You may now configure Schedules, Workers and Jobs.

Skype for Business

With O365, Retain will also archive Skype for Business conversations. They will be saved to the Conversation History folder of the user.

4.5 Exchange Module On-Premise Exchange

Retain supports:

- A single forest Active Directory system, (Exchange and standard users)
- An Exchange Resource Forest, (One Exchange Forest linked to one or multiple User Forests)

Retain does NOT support multiple linked Exchange Forests. Ensure that the Exchange Settings have been configured correctly before continuing the Exchange module setup.

4.6 Exchange Prerequisites required for Retain

There are several prerequisites that need to be done in Exchange for Retain to successfully archive the mailbox databases:

- A mailbox user with ApplicationImpersonation rights
- Basic Authentication enabled for Autodiscover and EWS on all Client Access Servers
- A DNS SRV record
- Set the DNS used by the Retain server to be the same as used by Exchange.
- Set a Rolling In-Place Hold to retain data until Retain can archive it.

4.6.1 Create a Retain User

To connect with Exchange, Retain needs a user with appropri-ate rights. This can be accomplished by using an existing user, or by creating a new one. It is recommended to create a new user for Retain archiving. If creating a new user, ensure that the user is an active user account and that the password does not change to ensure Retain will be able to access mail without changing settings. This user is sometimes called a ‘service account’. Retain calls this user the ‘global catalog user’.

The user created or used for Retain must be a “mailbox-enabled user” with read access to see all other users, groups, resources, and Exchange Servers in the Exchange Forest. The user will be utilized by both the Retain Server and Worker for LDAP lookups in Active Directory. The Retain user also must have Exchange impersonation rights to every mailbox user on every server in the organization to be archived. The Retain user MUST NOT be a member of any Exchange Administrator group, as Exchange denies impersonation rights for all admin-istrator accounts.

Additional permissions need to be added to the user created for Retain. The quickest way to add these rights is through the Exchange Management Shell.

After creating the new user in Active Directory, open the Exchange Management Shell.

4.6.2 Grant Impersonation Permissions to the Retain user.

In Exchange 2013 and 2016 Impersonation permissions can be granted in the Exchange Admin Center under Permissions.

Under Admin Roles create a new role (e.g. Retain Impersonation Management). Add the role "ApplicationImpersonation" and add the Retain User as a member.

You can also accomplish this via PowerShell commands using the Exchange Management Shell.

The commands required are different depending on the version of the Exchange Server. Exchange 2010, and 2013 require only one command per Exchange system to be issued, whereas Exchange 2007 requires the commands to be run on every Exchange server in the Exchange system to grant required permissions. If the Exchange system contains mixed 2007, 2010, and 2013 servers, the different commands must be completed on one server of each type.

Exchange 2010, 2013, and 2016 commands

For Exchange 2010, 2013, and 2016 the only command necessary for impersonation permissions is:

New-ManagementRoleAssignment –name ImpersonationAssignmentName –Role ApplicationImpersonation –User ServiceAccount

Where the ‘Name’ is a name chosen by the administrator and the ‘ServiceAccount’ is the name of the Retain user.

For Example:

New-ManagementRoleAssignment –Name impersonation-retain -Role ApplicationImpersonation -UserRetain

If additional Exchange servers are added to the system after running this command to grant rights to the ‘retain’ user, the command must be run again to grant rights to the new server.

4.6.3 Room and Equipment Resources

To archive Room and Equipment Resources, or to restore them, the Retain user, or Service Account, must also have delegation rights. These commands must be issued manually for each Room and Equipment or resource mailbox on every relevant server. This is required for 2013 and 2016.

These commands must be issued:

('Retain' is used here as the name of the Service Account, or Retain user, and the 'Mailbox Database' should be changed to the appropriate name.)

NOTE: (every time a new Room and Equipment or resource mailbox is added, the first command must be re-run.)

Exchange 2013 and 2016 Powershell commands

Get-Mailbox –ResultSize Unlimited –Database "Mailbox Database" | Add-MailboxPermission –User "Retain" –AccessRights FullAccess

Add-ADPermission –Identity "Mailbox Database" –User "Retain" –ExtendedRights Receive-As

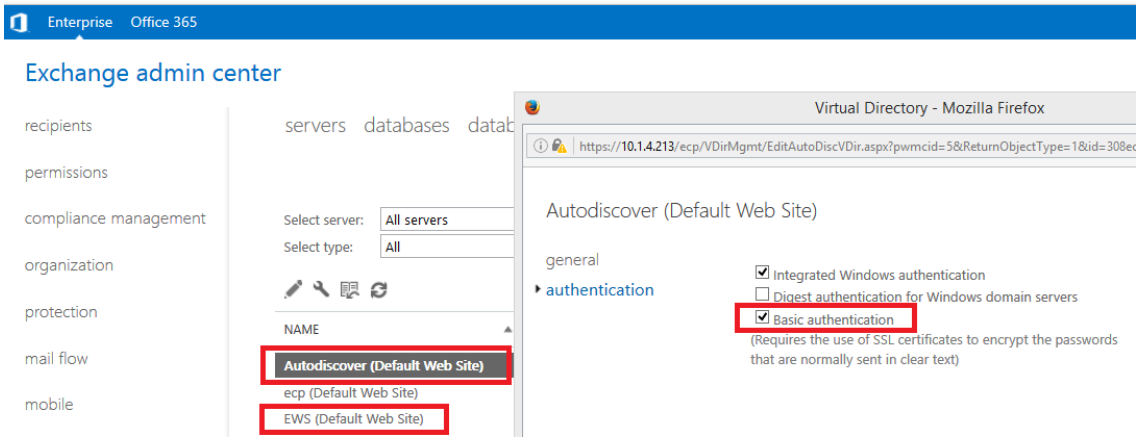
Add-ADPermission –Identity "Mailbox Database" –User "Retain" –ExtendedRights Send-As

4.6.4 Basic Authentication

Retain requires Basic Authentication to be enabled on *each* CAS Exchange server in the system for Autodiscover and EWS.

In Exchange Admin Center, go to Servers, then go to the Virtual Directories tab.

1. Edit Autodiscover and under Authentication enable Basic authentication if it is not enabled.
2. Edit EWS and under Authentication enable Basic authentication if it is not enabled.
3. Do this for each server in the list.



To check if this worked, run the following PowerShell cmdlets:

For EWS:

Get-WebServicesVirtualDirectory | ft server,basicauthentication

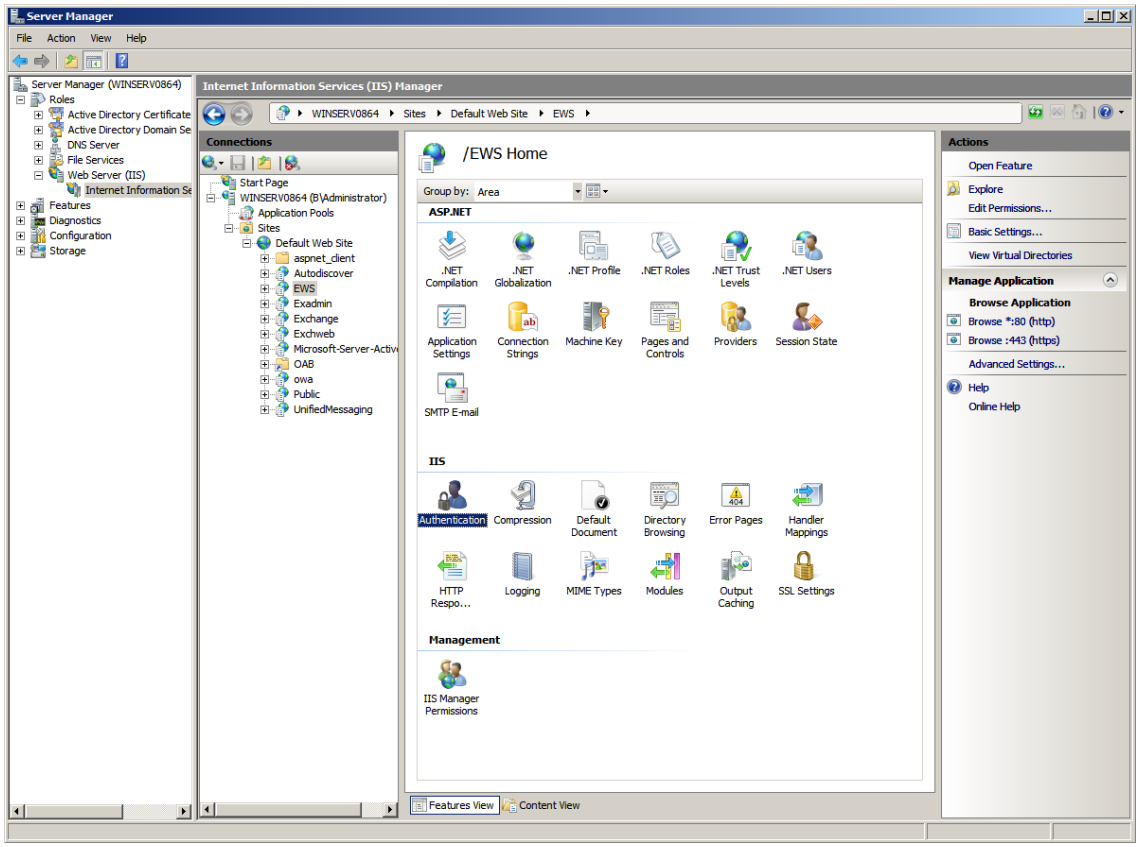
For Autodiscover:

Get-AutoDiscoverVirtualDirectory | ft server,basicauthentication

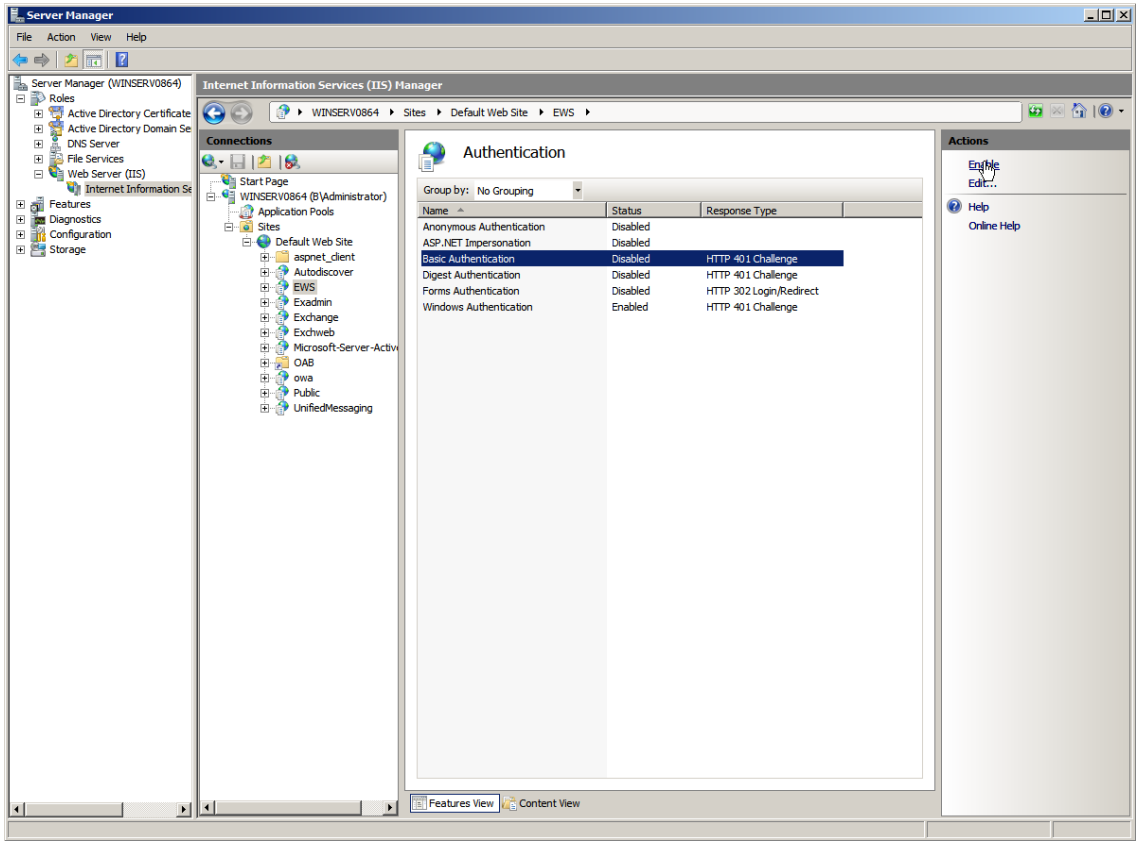
On Exchange systems prior to 2013 you may need to set basic authentication manually.

Open “Server Manager” on Exchange server.

1. In left pane, expand “Roles”, expand “Web Server (IIS)”, select “Internet Information Ser-vices (IIS) Manager”.
2. A new “Connections” pane opens, expand your Exchange server object, expand “Sites”, expand “Default Web Site (Multiple Protocols)”, select “EWS”.



3. Under heading “IIS”, open “Authentication” icon



4. Select “Basic Authentication”, click “Enable” in right pane.

You can now close “Server Manager”.

DNS SRV Record

Microsoft has an article describing how to set up a DNS SRV record titled "A new feature is available that enables Outlook 2007 to use DNS Service Location (SRV) records to locate the Exchange Autodiscover service".

In general, you will need to:

1. Go to the DNS Manager
2. Expand Forward Lookup Zones
3. Locate and right-click on the external DNS zone and choose Other New Records
4. Click Service Location (SRV) and enter:

   Service: _autodiscover

   Protocol: _tcp

   Port Number: 443

   Host: [your mail host, e.g. mail.gwava.net, usually the AD domain forest found in AD Domains and Trusts on the MS AD server]
5. Click OK

The Microsoft autodiscover library in Retain expects a URL along the lines of https://autodiscover.[your domain]/Autodiscover/Autodiscover.xml (e.g., https://autodiscover.xyzcompany.com/Autodiscover/Autodiscover.xml), which can be found in the worker log as it attempts to login by searching for "Discovered endpoint:" or "AutoDiscover".

Server DNS Setting

Retain has best performance when the server's network setting is using the same DNS as the Exchange servers.

Set Rolling In-Place Hold

To prevent data loss, it is highly recommended that a rolling in-place hold be set so users are unable to remove items from disk before Retain has a chance to archive them.

In Retain, set Profile/Miscellaneous to Include user's recoverable items.

In Exchange Admin Console, set up a distribution list, for example All_Mailboxes, that contains all mailboxes. It will be best to create a policy to add new users to this distribution list by default.

Place the distribution list under a 90-day hold.

In the Exchange Management Shell:

An In-Place Hold can be set up for all mailboxes for 90 days:

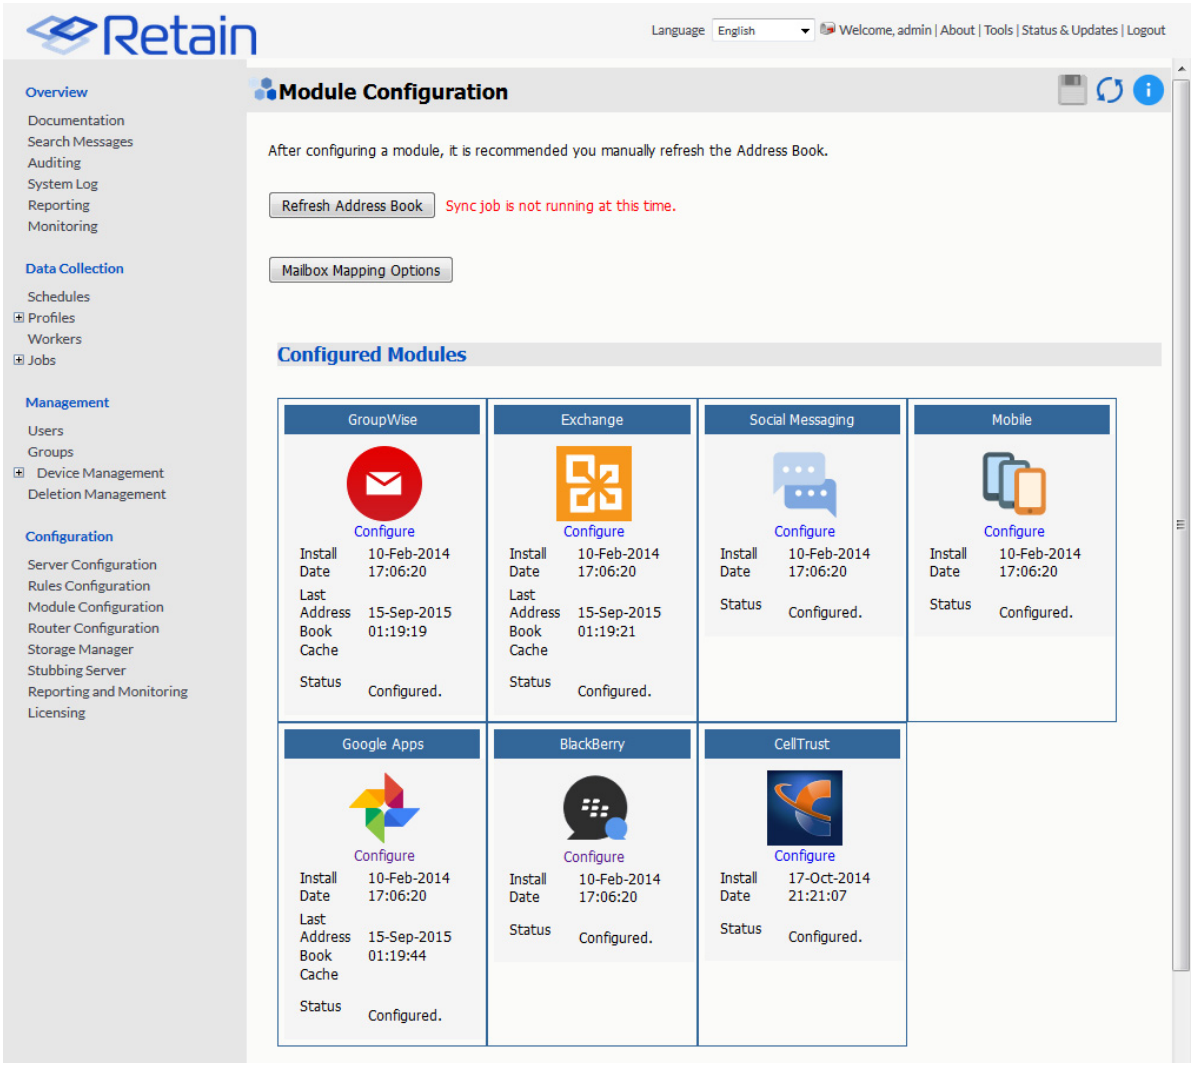
New-MailboxSearch "Retain90DayHold" -ItemHoldPeriod 90 -InPlaceHoldEnabled $true -SourceMailboxes All_Mailboxes

It will take time for the hold to take effect. You can determine how many mailboxes were placed under hold with the script:

((Get-Mailbox).InPlaceHolds).Count

4.7 Configure Exchange Module for On-Premise Exchange

The Exchange module must be configured in the Retain Server before any communication between Retain and an existing Exchange message system can occur. Open the Retain management page on the Retain Server, and select Module Configuration.



Select the ‘Configure’ option in the Exchange module. A new window or tab will open with the module configuration.

NOTE: Ensure that your Retain Server DNS is set to the same DNS server that your Exchange server uses. The Exchange module uses these DNS setting to auto discover critical information about Exchange stored in Active Directory and will not function correctly unless both systems are pointed to the same DNS server.

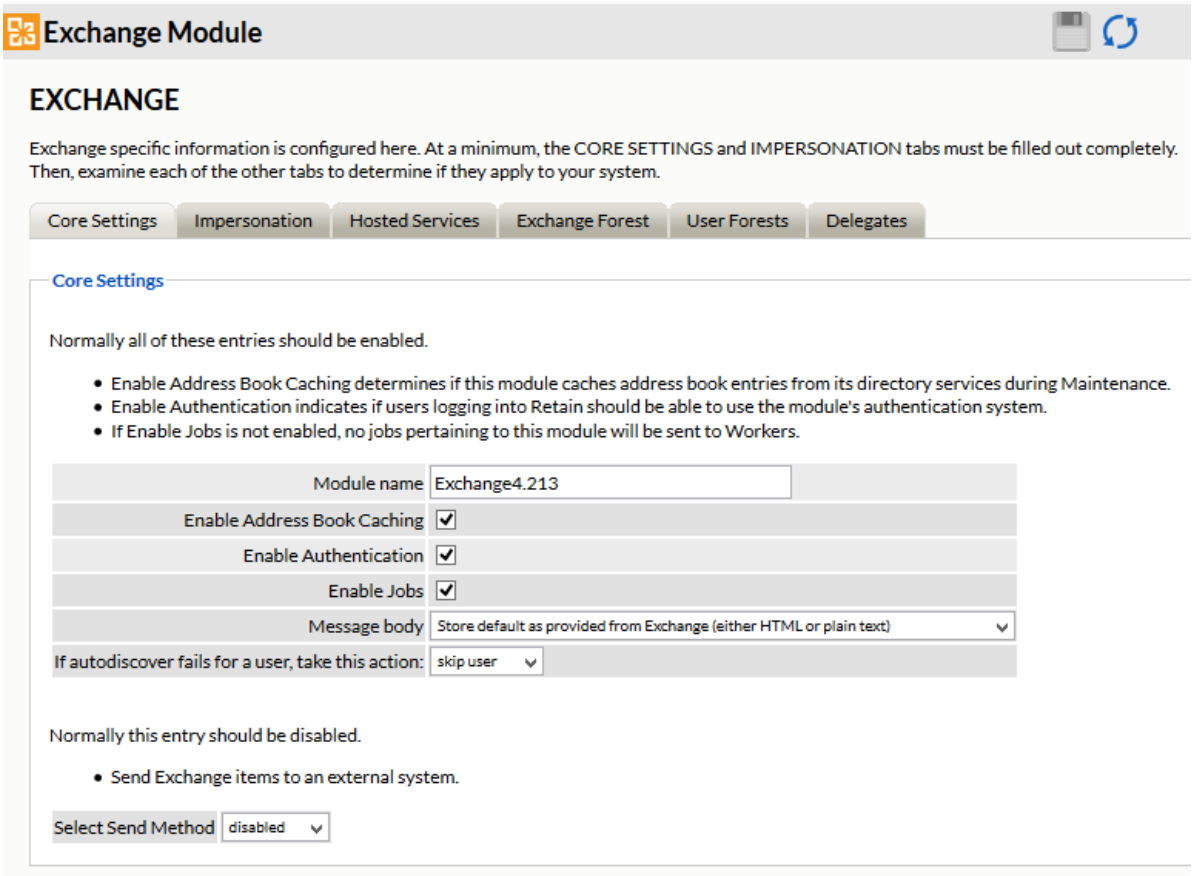
4.7.1 Core Settings

The module needs to be enabled on this page to make it active in the Retain system.

The module can be given a name.

The Send Method option enables either the SMTP Forwarding or FTP features. For either feature to appear and function, the Module Forwarding tab must be configured on the Server Configuration page. See that section for more information.

Normally all the checkboxes on this tab are always left selected. It is rare that you would ever deselect any of them. Two cases where you might, would be: troubleshooting (as instructed by Technical Support), or retrieving an old email system.



The Enable Address Book Caching function allows Retain to regularly cache the online email systems address book and synchronize it with Retain. This is critical for administration, authentication, and archiving purposes. It is recommended to cache the Address Book once every 24 hours to keep the Retain storage system up to date. By default, maintenance is set to cache the Address Book once every 24 hours.

The Enable Authentication checkbox determines if end-user authentication is performed when the user logs into Retain. If it is deselected, the Retain system will NOT authenticate the user against the email sys-tem and the user will not be able to log in unless another authentication method is enabled.

The Enable Jobs checkbox determines if configured data retrieval jobs are ever passed to the Worker. Even if the individual job is fully configured and enabled, if this checkbox is switched off, no jobs config-ured for this module will be processed.

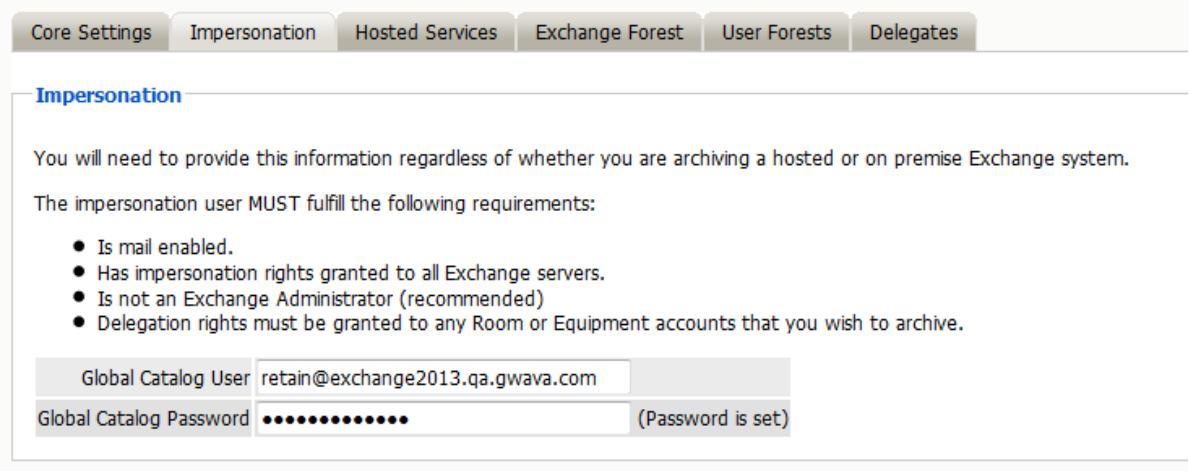
The Message body allows the administrator to decide whether to store either the HTML or plain text message body, or both.

Send Method

The Send Method option enables either the SMTP Forwarding or FTP features. For either feature to appear and function, the Module Forwarding tab must be configured on the Server Configuration page. See that section for more information.

4.7.2 Impersonation

If the Impersonation and Core Settings tabs are not completely configured with the correct information, the hosted system will not be archived correctly.

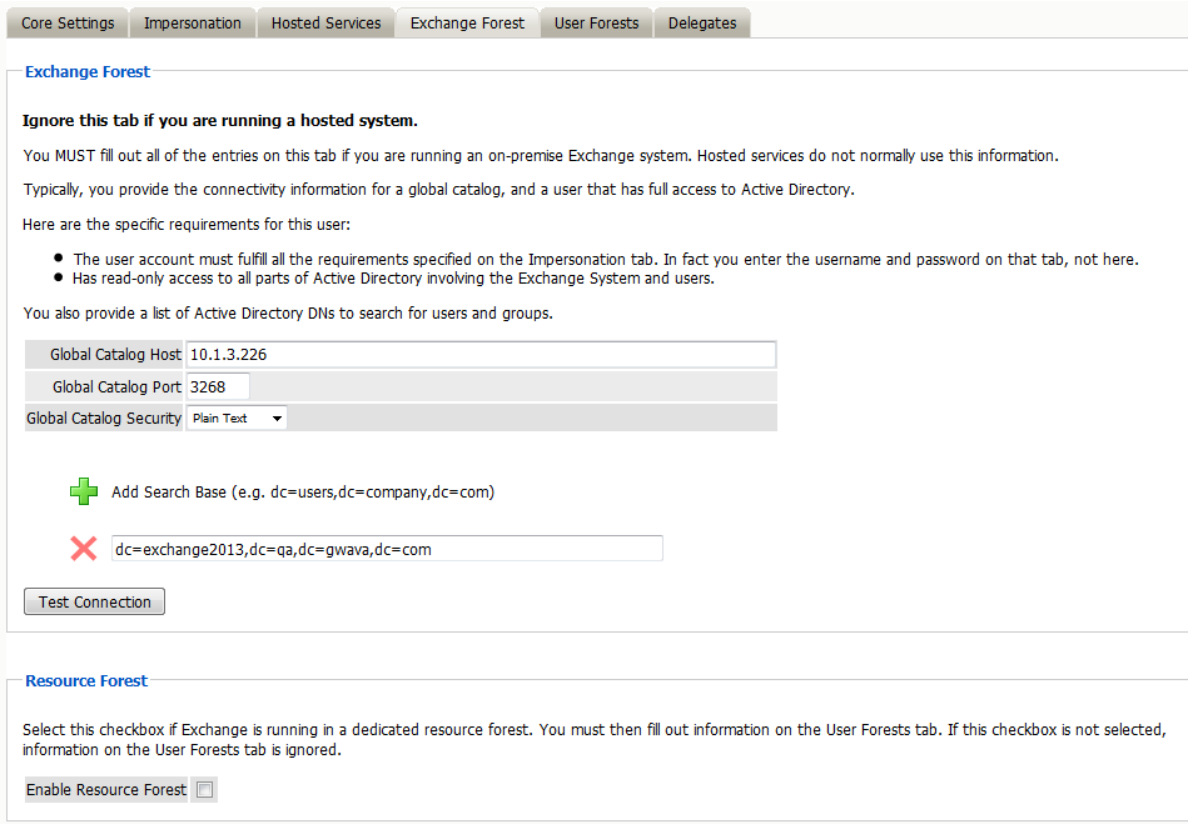


4.7.3 Hosted Services

This tab is not used with an On-premise Exchange system.

4.7.4 Exchange Forest

Retain needs to know where to access the Global Catalog Host and existing domains before any archiving can be accomplished.



Open the “Exchange Forest” tab and enter the IP address or hostname of the Global Catalog Host.

Click on the Green Plus sign to add a search base. This should be set to the highest level of the LDAP domain so the entire address book can be found. For example: DC=exchange2013,DC=qa,DC=gwava,DC=com

Retain uses Active Directory extensively when integrating with Exchange. Its uses include: populating the address book, authentication, and access to the Exchange System.

There are settings required for Exchange, see the Exchange Permissions required for Retain section.

On the Exchange Forest tab, you configure all the Active Directory information you need for an Exchange forest. There is no need to fill out any information on the User Forest tab unless the users exist in a separate forest from the Exchange Forest.

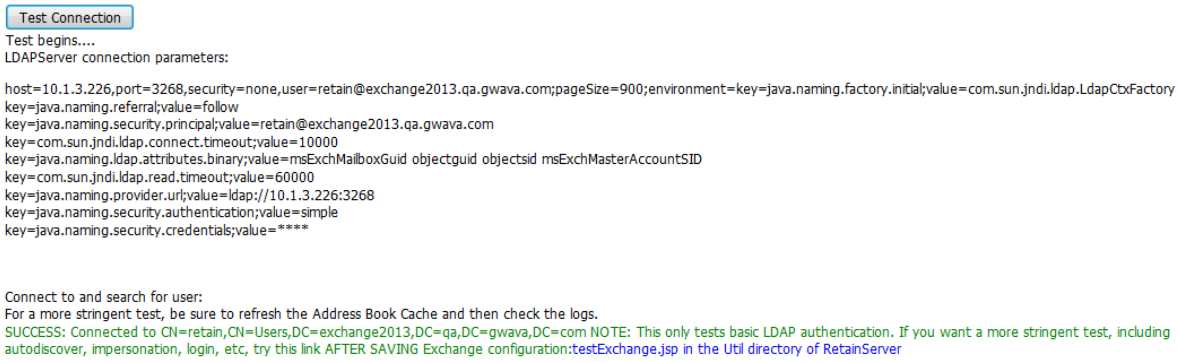
On the Exchange Forest tab, specify whether to use SSL or not for the Global Catalog Security and the search base, (use of SSL with the Global Catalog Security and search base is highly recommended). The search base is the LDAP path to the base of where Retain will start searching for valid Exchange users.

The Global Catalog Port defaults depend on whether SSL is used for security or not. (SSL is strongly rec-ommended. Default ports are 3268 for plain text, and 3269 for SSL.) Adjust as appropriate for your sys-tem.

You also must provide the credentials of an Active Directory user. This user is special It must have full read rights to Active Directory, be a mailbox-enabled, user, and be granted various Impersonation and Delegation rights. More on this is discussed in the Exchange Permissions required for Retain section. The username must be in UPN format, (user principal name).

This search base, in LDAP form, must be high enough in the tree to include ALL users, groups, and servers. Multiple search bases can be specified, though it often results in a less efficient interface. These are LDAP search bases which allow Retain to resolve all users, groups, and servers of interest in the forest.

After the Search Base has been added, test the connection to ensure information and connection works. The test performs a simple login to confirm that the user exists, the Exchange Server is reachable, and that the credentials are accepted. The test does not confirm impersonation or delegation rights neces-sary for the Service Account.

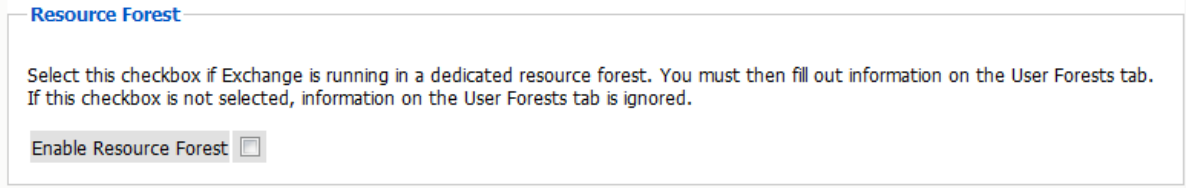


If the test results in an error stating: “FAILURE: User doesn't exist or is not mail enabled,” It indicates that the user’s mailbox is unavailable. A mailbox is not required for Retain to utilize the specified user. If the user Retain utilizes does not have a mailbox, this error may be ignored. However, if the user specified does have a mailbox, this may indicate connection issues.

If the test results in an error with an LDAP error code 49 it is an authentication error. The important bit of information is what comes after the *data* field. That is the LDAP connection error code that applies to this case.

```
525 user not found
52e invalid credentials
530 not permitted to logon at this time
531 not permitted to logon at this workstation
532 password expired
533 account disabled
701 account expired
773 user must reset password
775 user account locked
```

The Exchange Forest tab is the only tab required by the Server and the Worker to archive mail from the Exchange system. The User Forest tab, however, is required for Exchange systems utilizing a resource for-est, to allow the end user to log into Retain.

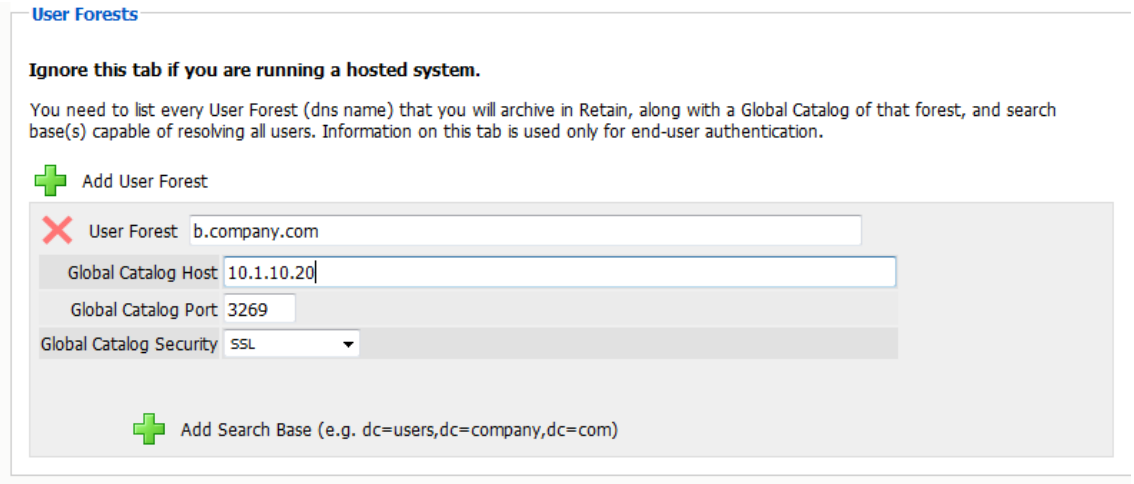


If the system contains a Resource Forest, enable the checkbox on the Exchange Forest tab and save changes. If the Resource Forest checkbox is not enabled, the User Forests tab will be non-functional and all settings contained on that tab will be ignored. The checkbox must be unchecked in a single forest Active Directory deployment, but must be checked in a multiple forest Active Directory deployment.

Check all information to ensure that it is correct and save changes, and then configure the User Forest if required.

User Forest

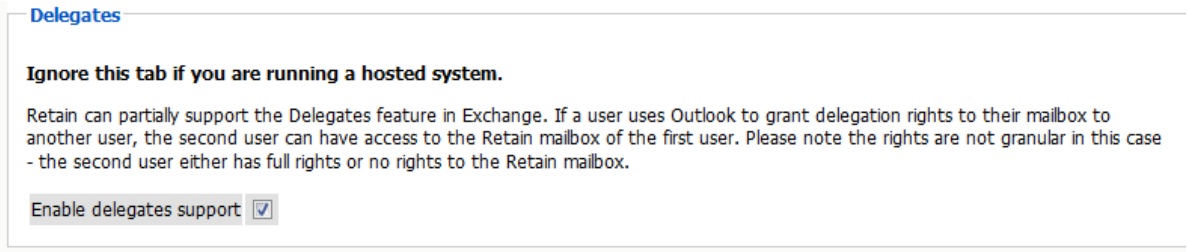
- The User Forest must have an entry for each user forest attached to the system.



- Select the green '+' but-ton and input the LDAP information required by the Forests' Global Catalog server: IP address or host-name, port, security, (SSL is strongly recommended), and all search bases to include all the users. No administrative credentials are required. Each end user's provided credentials will be used on login.

Delegates

You can set Retain to use delegate rights with On-Premise Exchange.

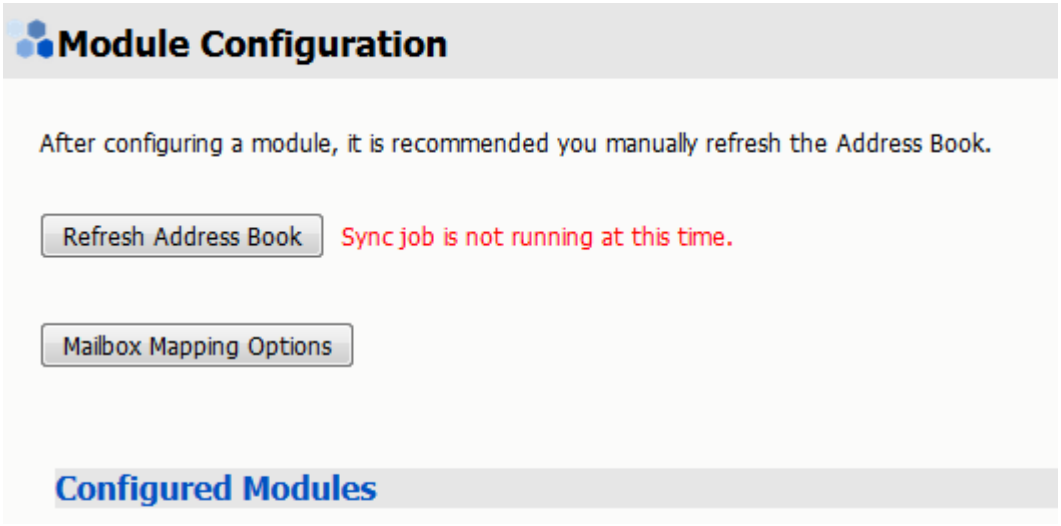


4.8 Finishing On-Premise Exchange

Save all changes before closing the Exchange Module page.

4.8.1 Refresh Address Book

After saving changes, return to the Retain Server's Module Configuration page, and trigger a refresh of the Address Book.



Depending on the size of the address book, it may take several minutes to return with information, but a successful configuration will return a correct address book cache date and no errors. The date should reflect the date of when the address book refresh was triggered.

The Status may show “Address Book Cache Never Run” or may list commonly misconfigured or missed items if the Refresh job fails.

Once the status is configured and the Address Book has been cached, Retain can connect to and archive messages from the Exchange server. The system is ready to have workers, schedules, profiles, and jobs configured, and those options will now appear on the main administrative interface.

4.9 Exchange Profile

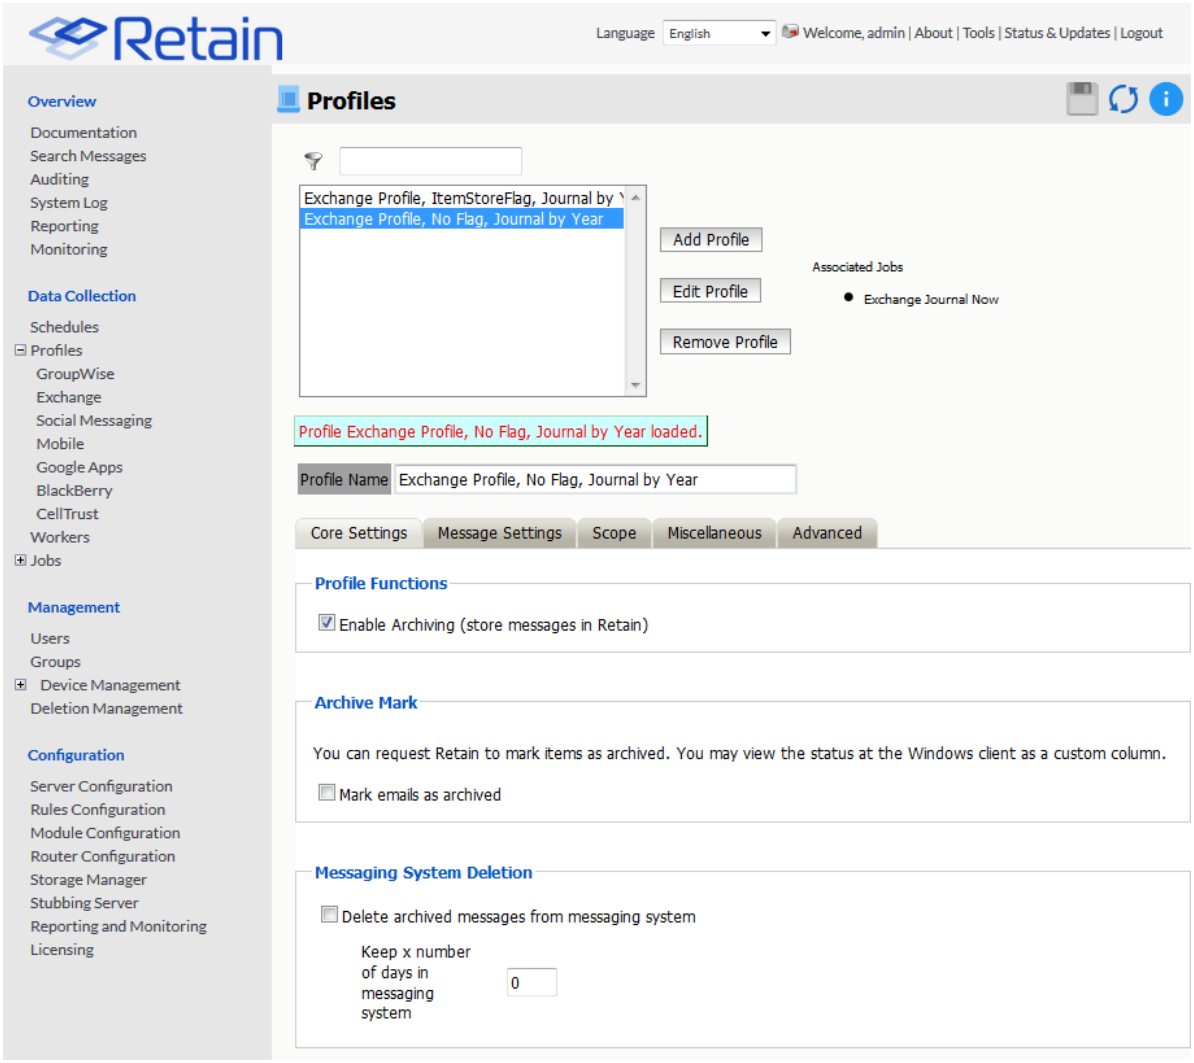
The job will need an Exchange profile setup to connect to the email system properly.

After the Exchange Module has been configured, the Exchange Profile will be available for configuration. If an Exchange Profile is not configured, jobs cannot be run against the Exchange system.

Click on “Add Profile” and provide a profile name, or select an already existing profile to access the con-figuration tabs. All changes made on this page must be saved by selecting the “save changes”, disk icon, at the top right of the page. Tabs may be changed and navigated through without affecting new settings, but any move to another page will require saving, or abandoning the changes made.

4.9.1 Core Settings

The core settings consist of an enabled/disabled option which must be enabled for any jobs based on this profile to archive anything.

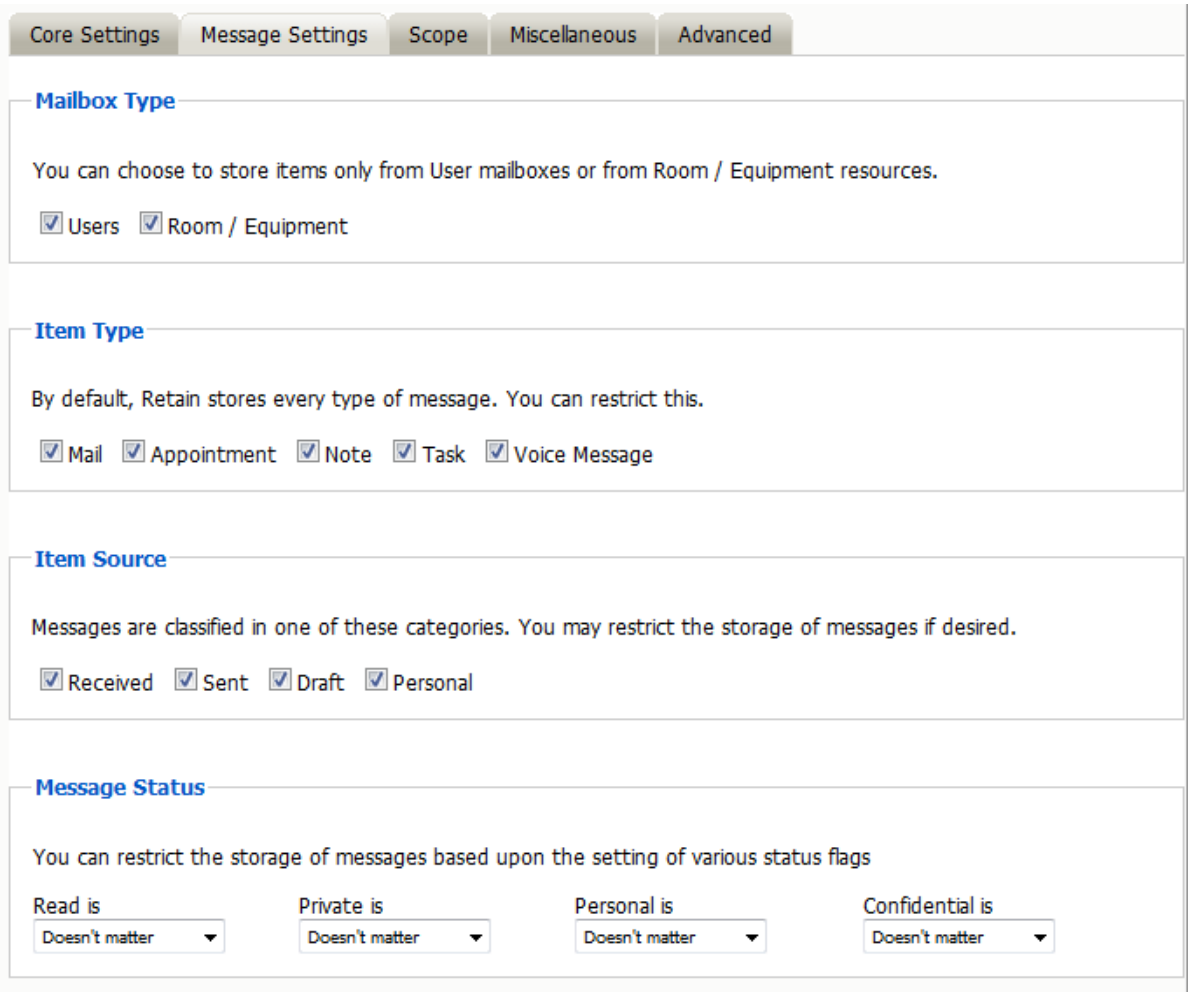


4.9.2 Messaging System Deletion

For systems where the administrator wishes to have archived messages removed from the system auto-matically, the Messaging System Deletion option may be used. Messaging System Deletion will remove messages from a mailbox after they are archived, according to the time frame specified in the settings. The amount of time to keep messages is specified in days. The recommended setting depends on the archiving scheme in the system. For instance, if messages are to persist in the system for 30 days, then the system deletion setting should be set to 30 and enabled. A setting of 0 will remove messages from the system as soon as they are archived. Be sure to configure the system before enabling the setting in the profile.

4.9.3 Message Settings

Retain can archive and select specific types of mail and Exchange system items to be archived. The Man-age Settings tab provides access to manage those settings.

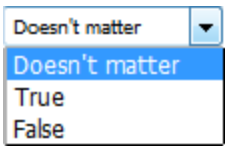


The Mailbox type specifies whether to include or exclude the available types of mailboxes. Because there can be multiple profiles and jobs, it may be advantageous to archive the Users and Room / Equipment mailboxes separately as needed and appropriate for the system.

The Item Type option specifies the different types of messages found in Exchange that can be archived, and allows the exclusion of or inclusion of the different individual types.

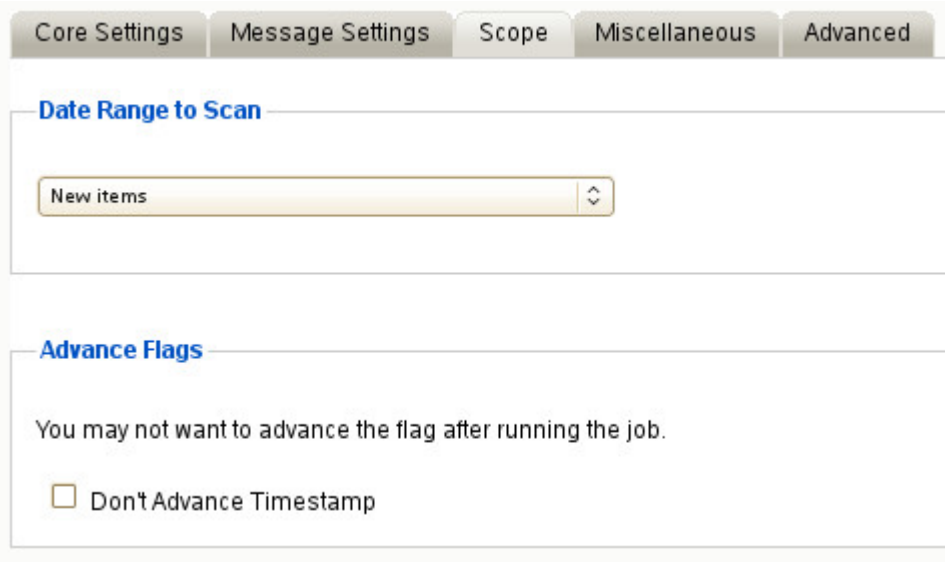
The Item Source option allows administrators to exclude or include messages that have not yet been sent or received, or posted.

The Message Status allows messages which have or have not been read or opened, or marked private or confidential to be archived. The different options in the drop-down menu are as shown.



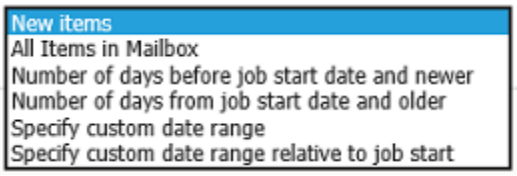
4.9.4 Scope

The Scope tab dictates the date range Retain will scan in the attached archiving jobs.



Date Range to Scan

The Date Range to Scan instructs Retain to scan for, and archive, messages after, or before, a certain date. This is useful if only specific chunks or areas of mail are to be archived.



New Items: All items that have not been archived by Retain since the last time the job ran.

All Items in Mailbox: All items in the mailbox starting from 1/1/1970, duplicates will be processed but not stored if they already exist in Retain's archive.

Number of days before job start date and newer: Only items from the relative number of days from the time the job began will be archived. E.g. messages that came into the email system less than 7 days ago.

Number of days from job start date and older: Only items before the relative number of days from the time the job began will be archived. E.g. messages that came into the email system more than 7 days ago.

Specify custom date range: Only items between two absolute dates will be dredged.

Specify custom date range relative to job start: Only items between two relative dates will be dredged. E.g. messages that came into the email system between 7 and 5 days ago.

It is recommended to archive all New items.

Advance Flags

Enabling "Don't Advance Timestamp" will not update the timestamp flag. Items that are dredged will still be considered new by Retain the next time the job runs.

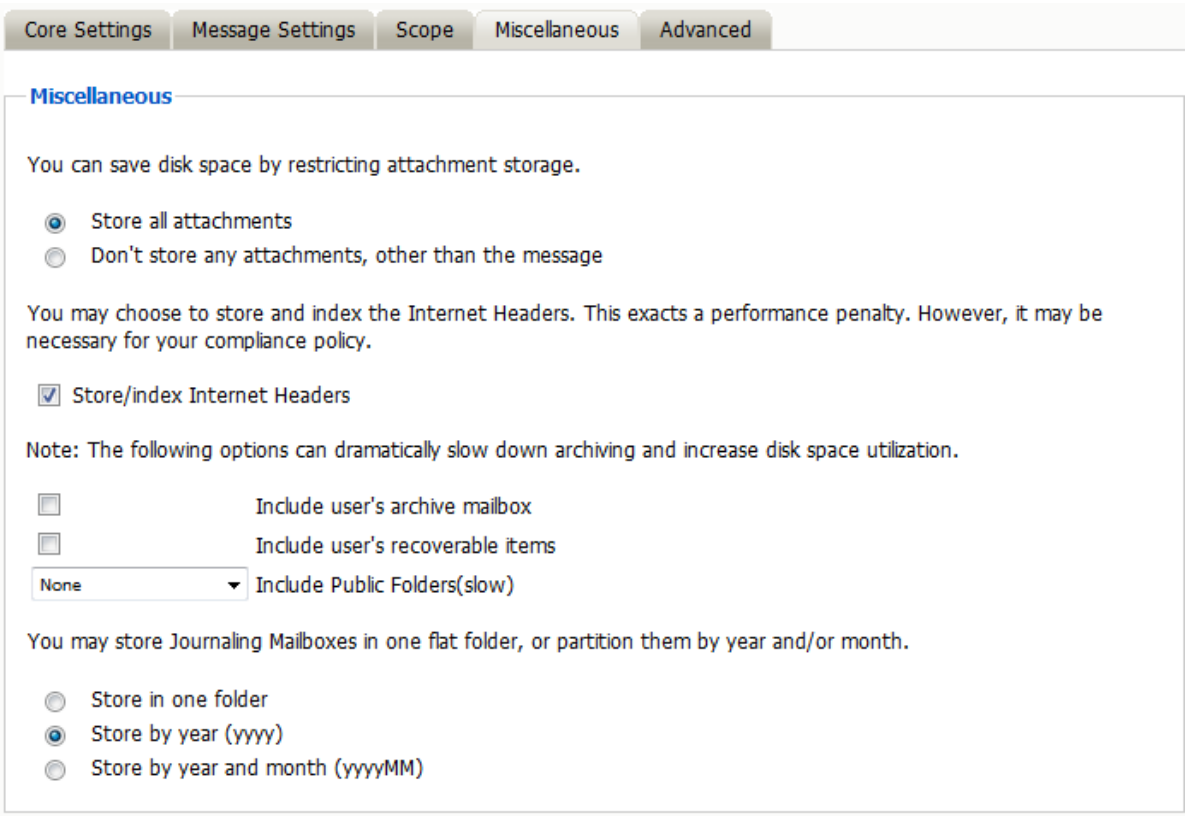
This is useful when troubleshooting, but is generally not used for normal jobs.

NOTE: Unlike GroupWise, Exchange does not ensure any compliance when scanning end user mailboxes; users may freely delete their email. The Item store flag does not prevent mail deletion. Only setting a rolling hold on all mailboxes, or journaling and archiving a journaling mailbox guarantees all items have been archived.

4.9.5 Miscellaneous

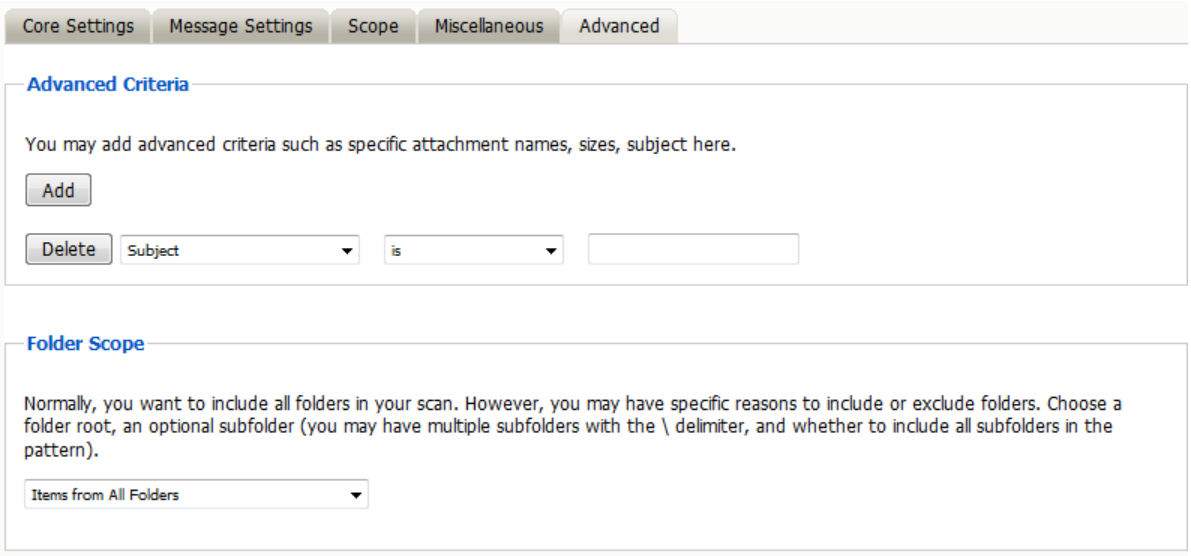
The Miscellaneous tab allows access to settings detailing how messages are stored and what is archived. Attachments, message information such as the internet headers, and how the data is stored and named, (by folders, year, or year and month), dictate not only the message store structure, but affect the storage size.

Miscellaneous options also allow for the archiving of the ‘recoverable items’. To enable checking and archiving of the ‘Recoverable Items’ for compliance reasons, select the checkbox next to the option.



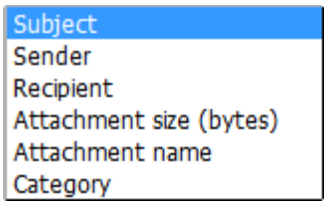
4.9.6 Advanced

The Advanced tab allows you to limit what is stored by Retain. This must be used with caution as this opens holes for data to be lost through. It is recommended to store everything since storage space is inexpensive.



Advanced Criteria

If you want to be more specific as to what to dredge or not to dredge, add the criteria here. Each line will be logically AND-ed together. Think “Dredge all items where the following is true:”Criteria A AND Criteria B AND Criteria C AND etc.



You may select based on:

- Subject
- Sender
- Recipient
- Attachment Size (in bytes)
- Attachment Name
- Category

And whether they are equal to, not equal to, contain or do not contain the item you specify.

This gives you great flexibility and granularity. It allows you to customize dredges and retention for many different groups, or even individuals.

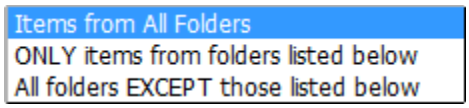
4.9.7 Folder Scope

By default, we dredge items from all folders. You can specify one or more inclusions or exclusions.

Your choices are:

- Dredge everything
- Dredge only these listed folders

- Dredge everything except these listed folders



How to specify the list of folders to dredge/exclude:

1. Specify a System Folder (mandatory). Example: Calendar.
2. You specify a subfolder of that folder (optional).

   Example: entering “old” would mean the folder “old” under “Calendar”.
3. You can have multiple hierarchies under that with the / delimiter.

   Example: “old/mail” would mean the subfolder “mail” under “old” under “Calendar”.
4. You specify if the option includes subfolder.

   Example: If you select “old” and “includes subfolder” is unchecked, “Calendar/mail” is selected. If “includes subfolder” is CHECKED, “Calendar/old/mail” would also be selected.

You may now configure Schedules, Workers and Jobs.

4.9.8 Distributions Lists

You can create distribution list in Exchange Admin Center to manage information dissemination. Retain will query Exchange for a list of users in each distribution list.

While you can create a distribution list in Active Directory Users and Computers these changes will not be reflected in Exchange therefore Retain will not see them.

If you wish to rename a distribution group it needs to be done in Exchange or Retain will not see it either.

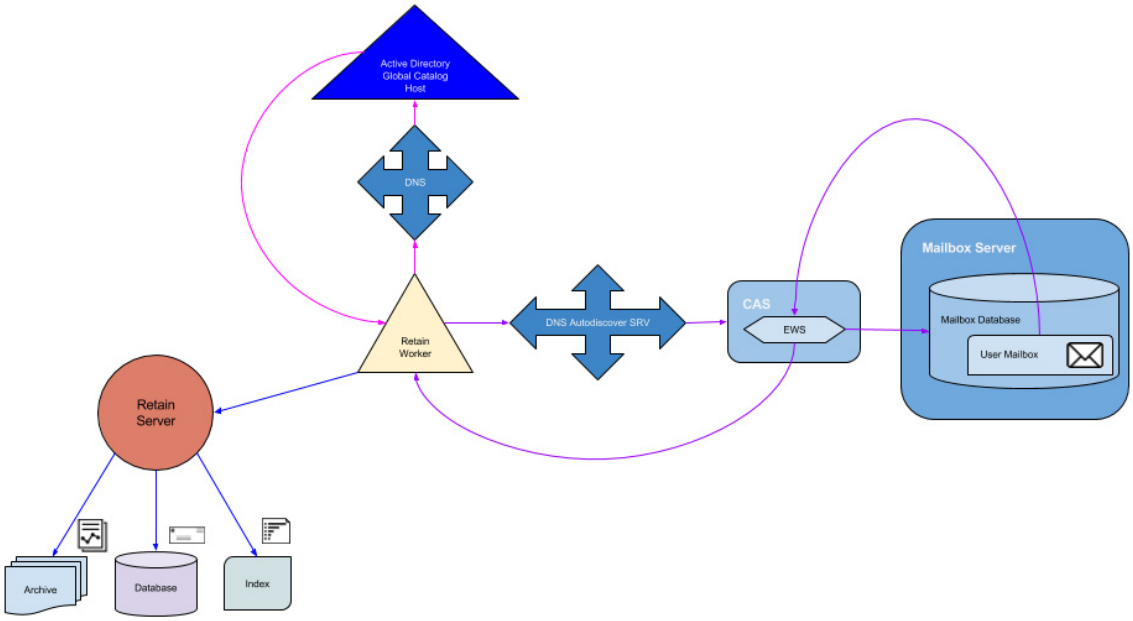
Distribution lists can be hidden in Exchange. If a distribution list is hidden, Retain will not be able to see the users associated with the distribution list and will not be able to archive the distribution list. The distribution list will be marked as (hidden) in Job | Mailboxes | Distribution Lists.

Dynamic Distribution Lists cannot be seen by Retain as they only create a user list at the time the message is sent. So, it is more of a filter then a list.

Remember to refresh the address book if you wish to see the latest list changes.

4.9.9 Exchange Message Dredging Process Overview

How does Retain get messages from Exchange?



1. When a job starts, the Retain Worker will query the DNS for the SCP record to the URL of the Active Directory Global Catalog Host.
2. Then the worker queries Active Directory for the Autodiscover SCP Records and Active Directory returns the Autodiscover URLs. The URLs tell Retain where to connect to autodiscover. There are also some default autodiscover URLs that Retain uses to connect to autodiscover.
3. Retain then uses autodiscover to connect to the Client Access Server. It is helpful to have an autodiscover SRV record on the DNS to speed up this process.
4. Once Retain has connected to the Client Access Server (CAS), the CAS uses EWS to connect Retain to the correct Mailbox Server.
5. Retain uses the impersonation user credentials to enter the mailbox of the user we are attempting to dredge messages from. Retain queries Exchange for messages that meet the criteria set in the job.
6. Exchange then serves the oldest message that meets the criteria back to the Retain Worker through EWS on the CAS.
7. The Retain Worker receives the message and opens it to query the Retain Server if the message body or attachments already exists.
    a. If the Retain Server determines that the message is new, then the body and attachments are stored in the archive, the header information and hash is saved in the database with links to the archive and the contents of the message are indexed.
    b. If the message already exists, the database is updated with the header data and linked to the existing data, and the existing message body or attachment is dropped by the worker and the next message is retrieved from the email system.

4.10 BlackBerry Module

The BlackBerry module allows the collection of BlackBerry phone data from the Enterprise server into the Retain system. BES 5.x, 10.x, and 12.x are supported.

To archive the BlackBerry information, the Black-Berry Enterprise Server must be configured to log phone calls, PIN, BBM and SMS data.

Archiving BBM Protected data is also supported through the REST API. (Archiving BBM Protected will require two addi-tional licenses: REST API license and the BBM Protected License. See the rest API section.)

4.10.1 Prerequisites

The server log files are what Retain uses to collect the information about phone activity, and as such no software or modification needs to be made to any system phones. If the BlackBerry server has not been configured to log the data, Retain will be unable to archive the information.

For BES 5.x, instructions to modify the logging in the BlackBerry Enterprise Server can be found in the BlackBerry Enterprise Server help file. (Found under Start | Programs | BlackBerry Enterprise Server | Help | BlackBerry Manager Help)

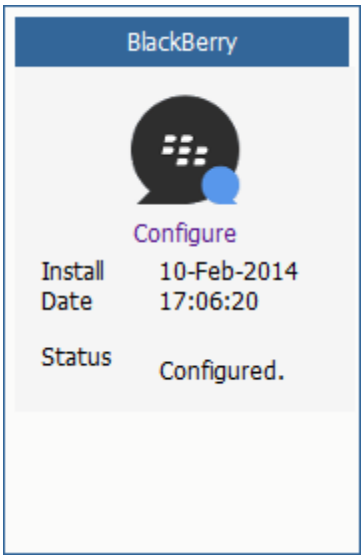
For BES 10.x and 12.x, users must be set on EMM – Regulated mode and have the logging set to ‘yes’ for all types desired. Balanced mode users will not have logs created for them in the BES system and as a result will not be archived.

4.10.2 BlackBerry Module Setup

Installing the BlackBerry Module is incredibly simple. Essentially, the module needs to be enabled and configured with the BES log file location. Afterwards, the data is treated the same as any other data entering the Retain archive. Data is archived in the database under the user/phone it originated from and can be searched and viewed through the Search Messages interface.

4.10.3 Configuring the Module

To enable the BlackBerry Module, select the ‘BlackBerry Module’ from the module page and click ‘config-ure’.



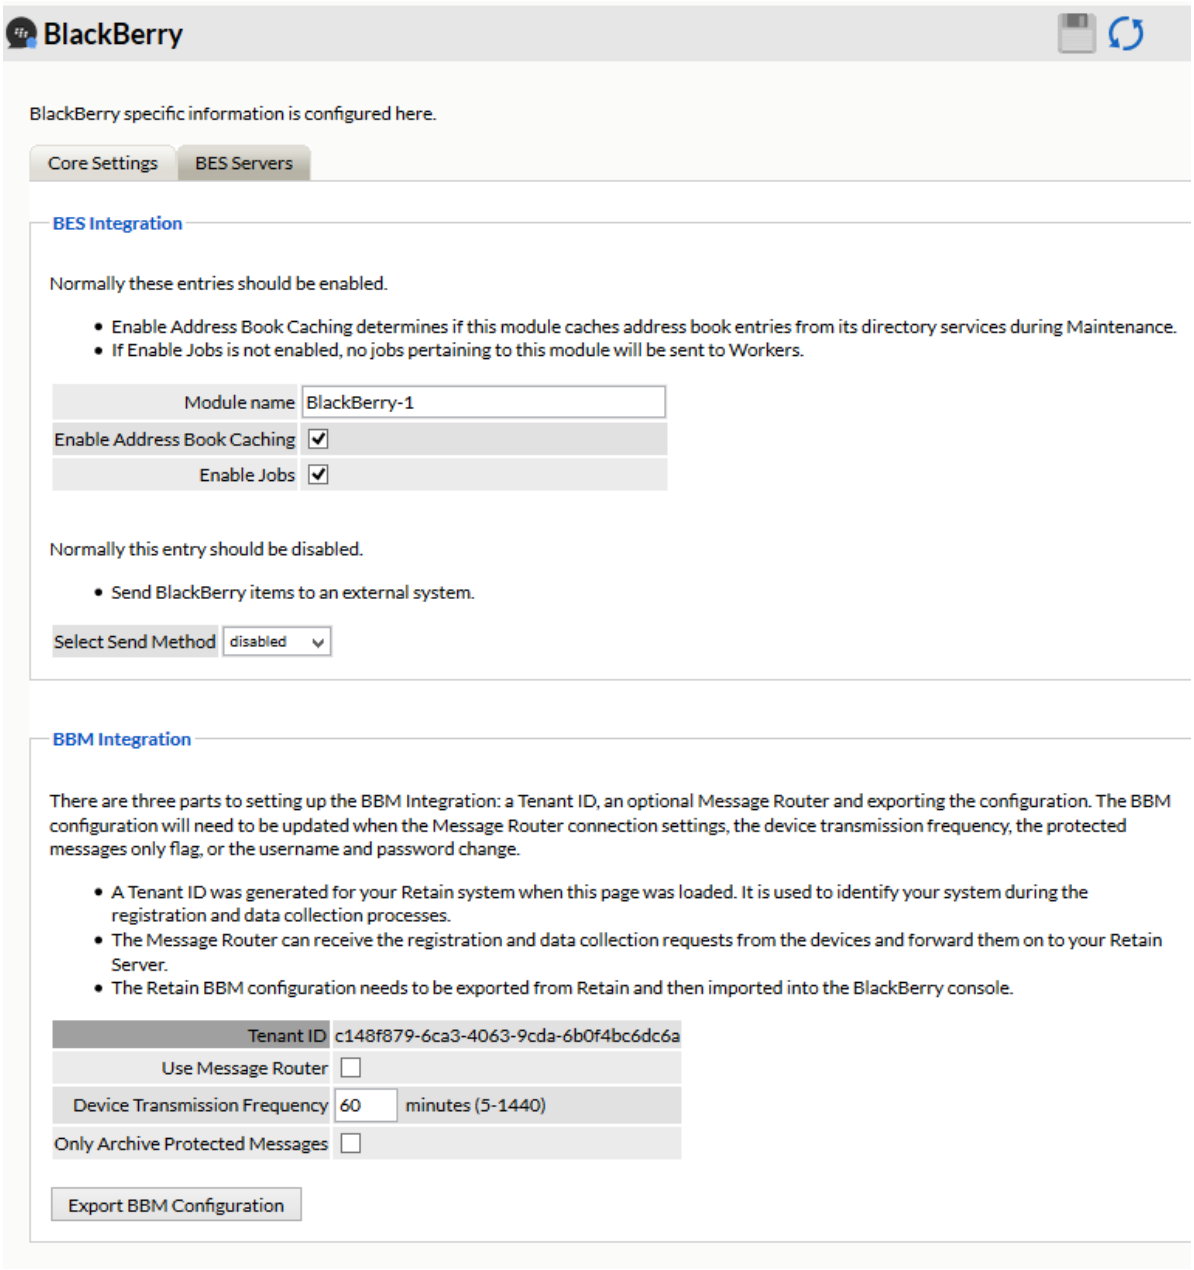
Core Settings tab

The module needs to be enabled on this page to make it active in the Retain system.

The module can be given a name.

The Send Method option enables either the SMTP Forwarding or FTP features. For either feature to appear and function, the Module Forwarding tab must be configured on the Server Configuration page. See that section for more information.

The BlackBerry module needs to be enabled on this page to make it active in the Retain system.

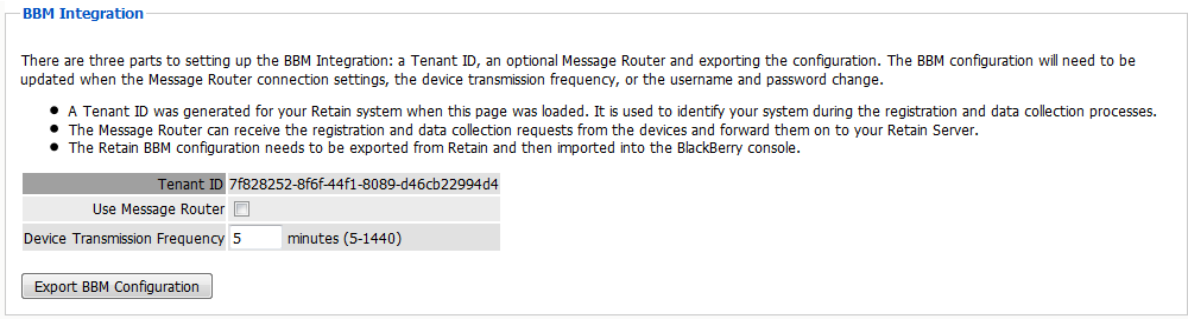


The Send Method option enables either the SMTP Forwarding or FTP features. For either feature to appear and function, the Module Forwarding tab must be configured on the Server Configuration page. See that section for more information.

For address book caching to occur, it must be enabled on the Core Settings screen. For more information on caching, see below.

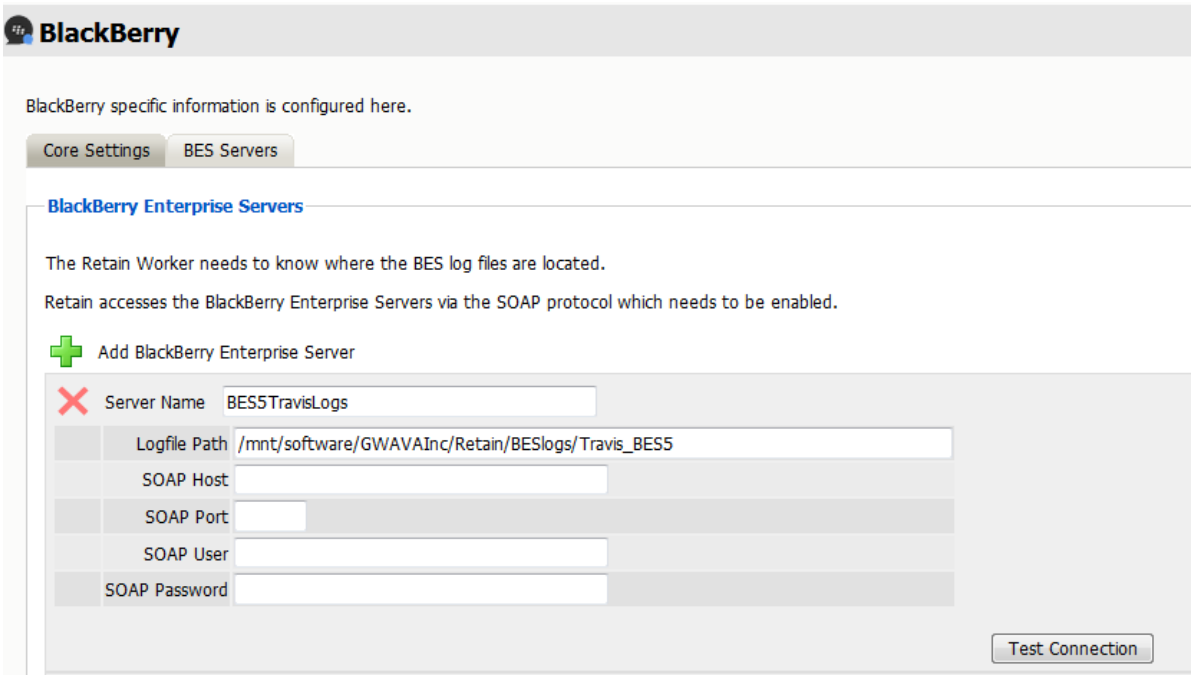
Before data can be archived, the BES Servers must be configured and a BES log file path must be pro-vided.

If BBM Protected integration is to be used, no BES servers will be listed and no other configuration in the BlackBerry Module is required. However, the BBM Configuration file needs to be exported and uploaded to the BBM server. Please see the BBM setup and install checklist for specific information on how to con-figure the BBM server.



4.10.4 BES Servers

The module needs the location of each BES server. Click the green plus sign to add a BES server.



The log file location must be the address to the Blackberry Server log files, from the perspective of the Retain Worker. Regardless of where the worker is installed, local or remote to the Blackberry Server, it must have direct file access to the logs. This access can be accomplished through a shared folder or loca-tion, but it must be available to the Retain worker. The path must be an open and active path from the Worker to the log file location. Appropriate access through firewalls or network security must be made available if the worker is not installed on the BES system.

It is important to keep in mind that the log path needs to be the path relative to where the Retain Worker is installed. If the Retain Worker is installed on a system that is local to the log directory than a local path can be used here. If the Retain Worker is on a remote system than a UNC path must be used. If using a UNC path to a Windows share it is necessary to ensure that the Retain windows service (tomcat) be run-ning under a user with the proper rights and permissions to read the log data being shared.

The SOAP data must also be filled-out. The SOAP connection is how the Retain Server connects to the BAS service's SOAP interface in order to cache the address list. SOAP must be enabled on every BES listed. List every BES server which the Retain system will be archiving data from. The SOAP address or host name and port which is configured must be specified for each Server. The SOAP user name and password are the credentials to the BES admin web console.

4.10.5 Caching

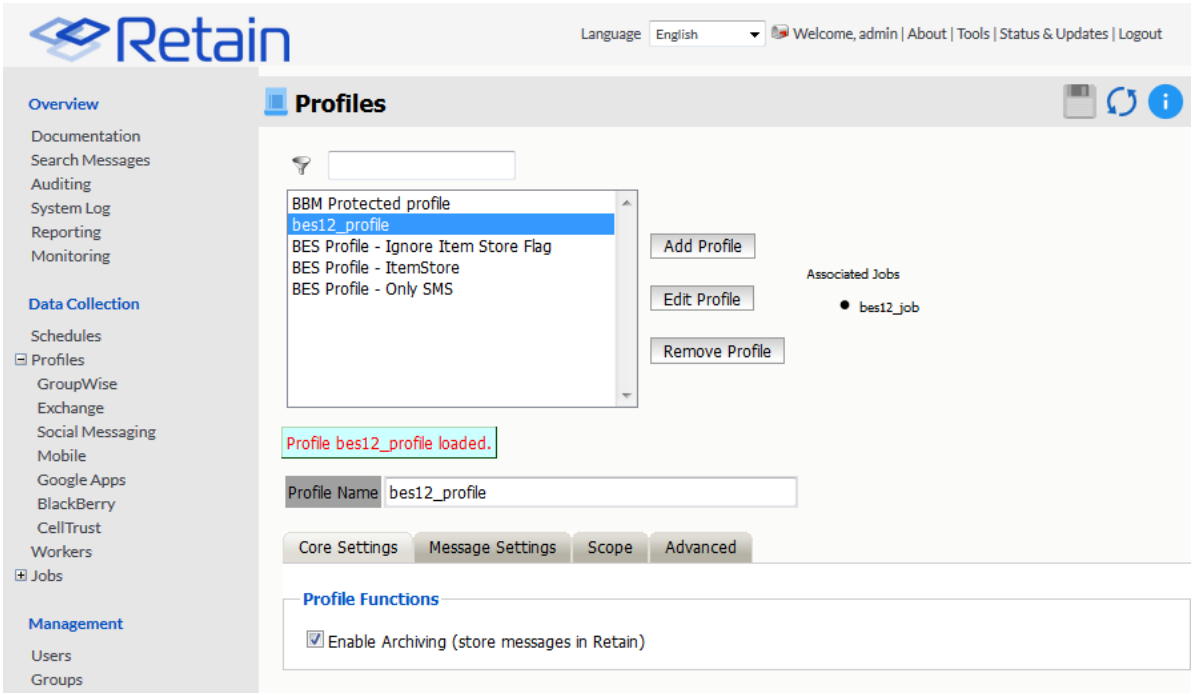
Retain will attempt to cache the address book from each server listed if address book caching is selected. Because the Blackberry Server logs, and especially the BBM Protected data does not always contain user names attached to the device ID's, Retain works to cache the address book to fill out data for messages; user name, email address, etc. Without Address Book caching, stored message data may only be identi-fied by a device ID or phone number. Each server's configuration has a test button to help ensure that Retain can successfully connect to the server. Selecting the 'Test Connection' button will initiate a con-nection attempt by Retain for the server. The results will be displayed after the test has been completed.

4.10.6 BlackBerry Profile

To create a new profile to archive BlackBerry data, first select the 'add profile' button and name the pro-file. Then continue configuring as desired.

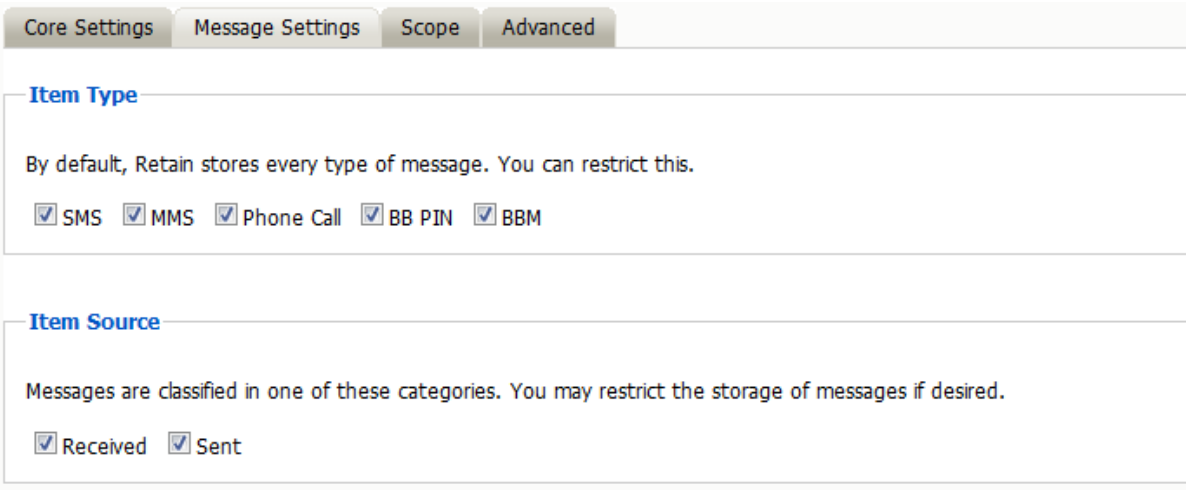
Core Settings tab

The Core Settings tab holds the enabling setting, which allows jobs connected to this profile to run. This must be checked before jobs will run.



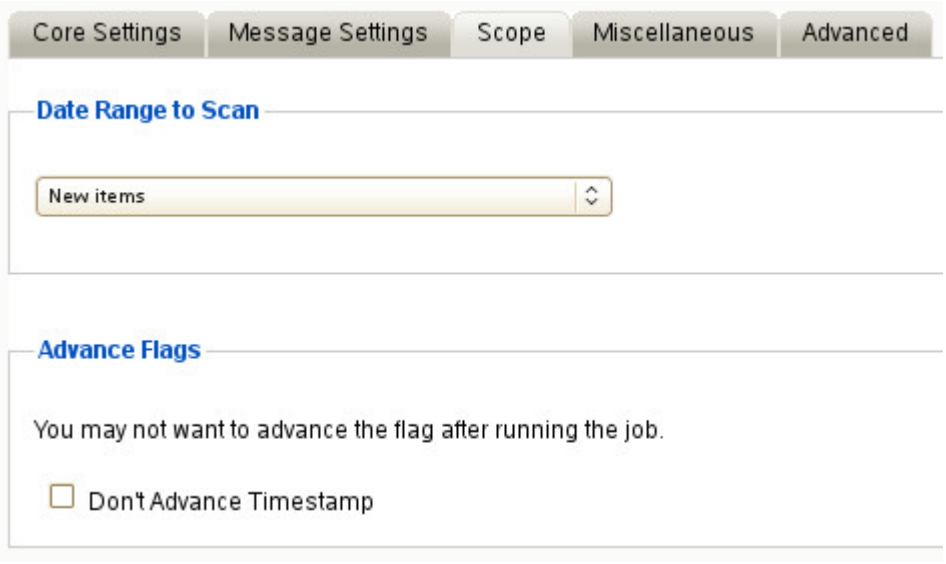
Message Settings tab

The message settings tab can be used to exclude or include specific data types for this profile. By default, all types of messages will be archived. All BlackBerry data which is selected to be archived must be logged in the BlackBerry Enterprise Server or the data will not be available for Retain to archive it.



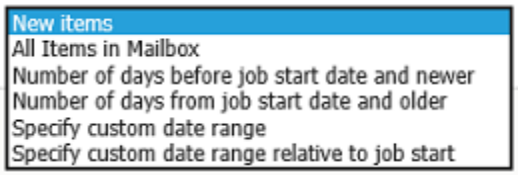
4.10.7 Scope

This is the most critical tab to fill out as it sets the limits on how much to archive.



Date Range to Scan

The Date Range determines which message items are collected, depending on the date of the message.



New Items: All items that have not been archived by Retain since the last time the job ran.

All Items in Mailbox: All items in the mailbox starting from 1/1/1970, duplicates will be processed but not stored if they already exist in the Retain archive.

Number of days before job start date and newer: Only items from the relative number of days from the time the job began will be archived. E.g. messages that came into the email system 7 days ago or less.

Number of days from job start date and older: Only items previous to the relative number of days from the time the job began will be archived. E.g. messages that came into the email system 7 days ago or more.

Specify custom date range: Only items between two absolute dates will be dredged.

Specify custom date range relative to job start: Only items between two relative dates will be dredged. E.g. messages that came into the email system between 7 and 5 days ago.

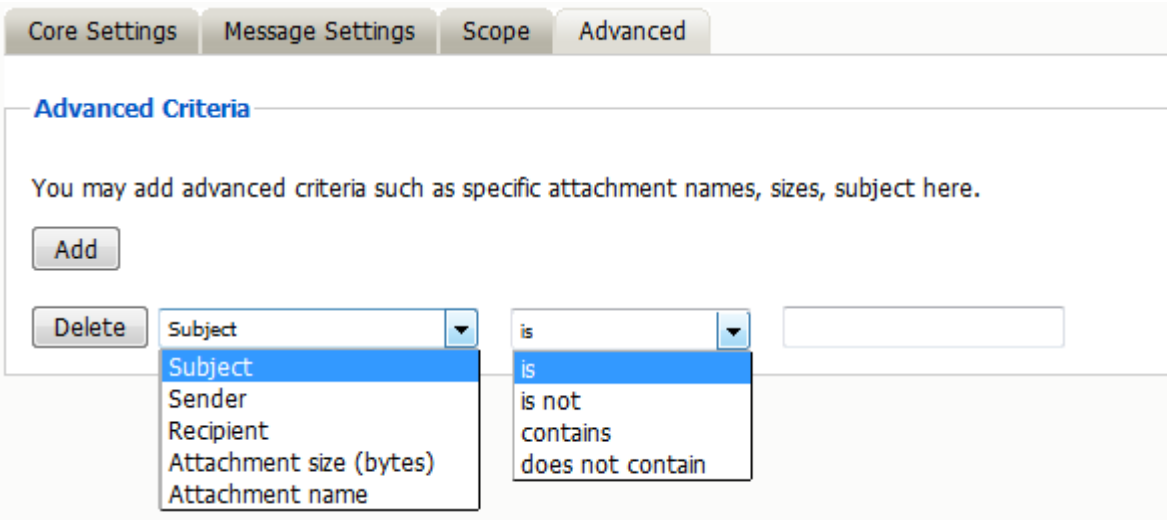
Advance Flags

Enabling "Don't Advance Timestamp" will not update the timestamp flag. Items that are dredged will still be considered new by Retain the next time the job runs.

This is useful when troubleshooting, but is generally not used for normal jobs.

4.10.8 Advanced

If you want to be more specific as to what to dredge or not to dredge, add the criteria here. Each line will be logically AND-ed together.



Think “Dredge all items where the following is true:” Criteria A AND Criteria B AND Criteria C AND etc.

You may select based on:

- Subject
- Sender
- Recipient
- Size
- Attachment Name

Whether they are equal to, not equal to, whether they contain or do not contain the item may be speci-fied.

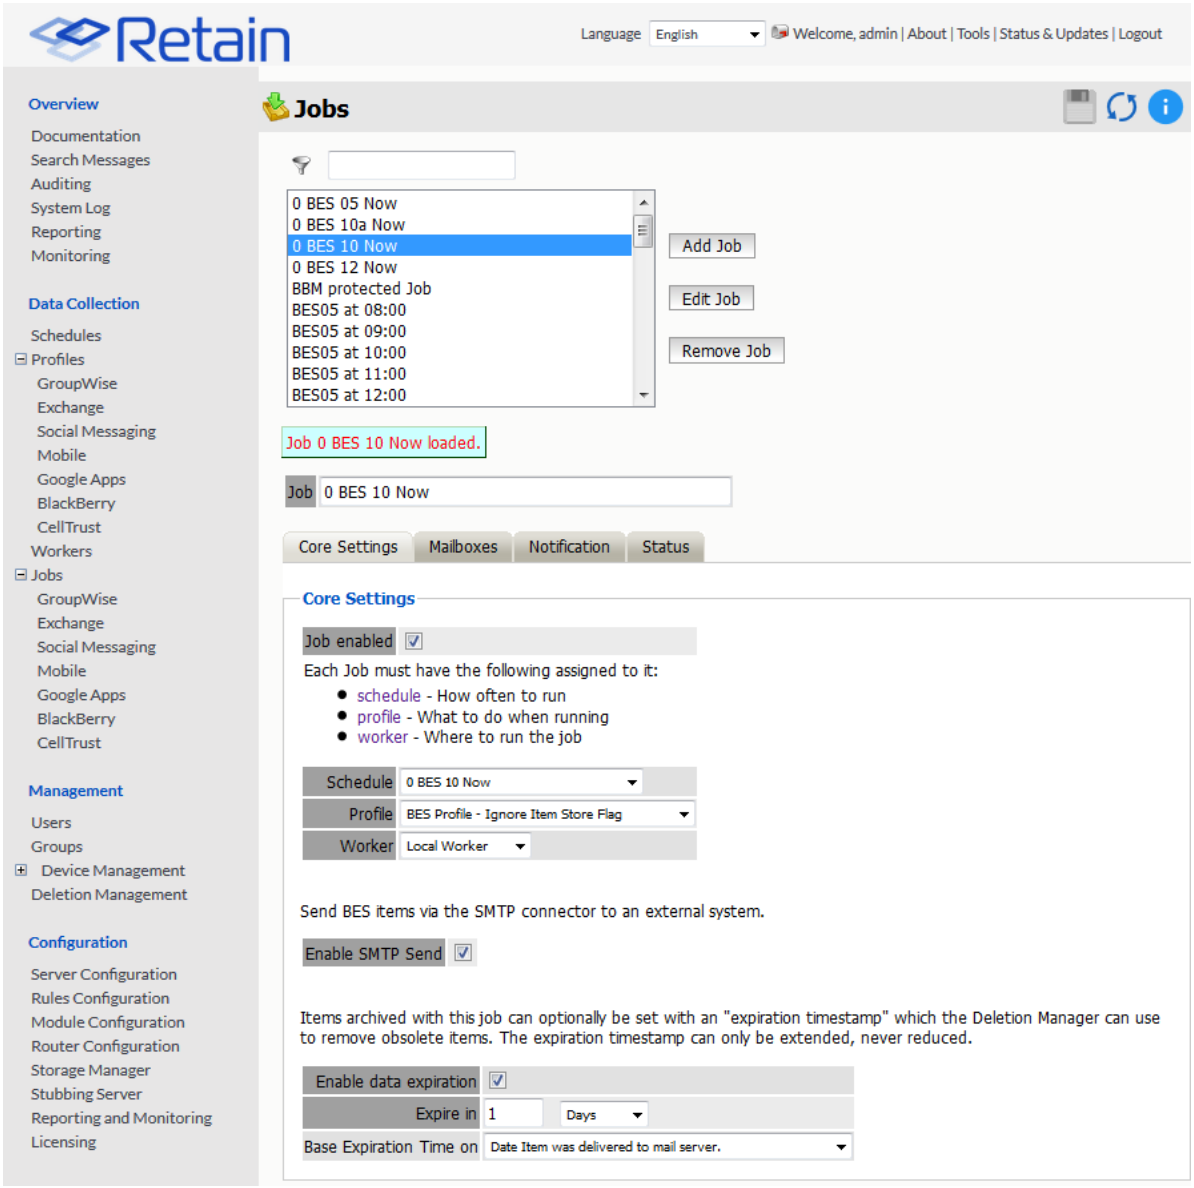
This provides great flexibility and allows you to customize dredges and retention for many different groups, or even individuals.

4.10.9 Blackberry Jobs

All data collection for the Blackberry module is configured through the Blackberry job interface. The jobs here combine the Profile, the Schedule, and a Worker together to archive the specified server logs. A Job must have a Profile, Schedule, Worker, and a selected server, before it can be saved or run. To create a new job, click, ‘Add Job’, and configure and save.

4.10.10 Core Settings

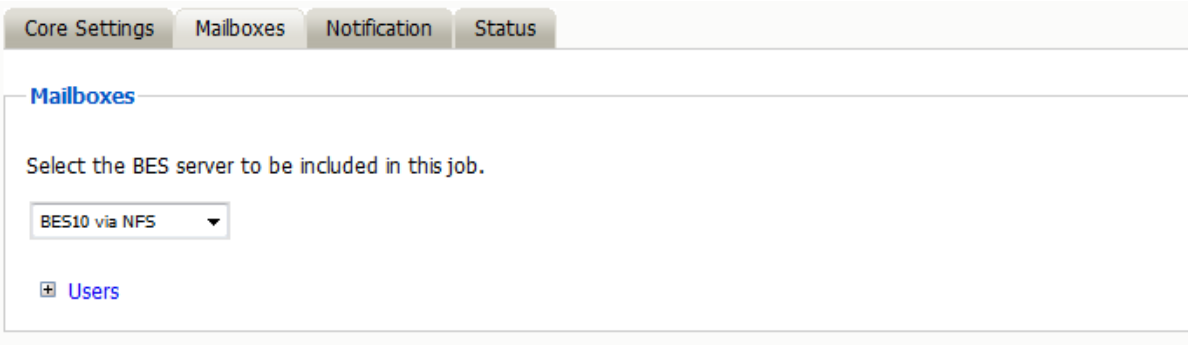
The core settings tab contains the basic configuration. Select the ‘Job enabled’ option to enable archiving. Also, select a Schedule, Profile, and Worker to complete the job. If the SMTP Send is enabled, the worker will forward BES message data.



The Data Expiration is an option to place a timestamp on the data in the Retain database, which allows for ease of automation for the deletion manager. In addition, devices such as NetApp, Centera, and Hitachi HCAP may use this number to enforce hardware level protection of the stored item so that no one (including Retain) may delete the item before its expiration date. Job Expiration is not retroactive for mail in the database, and only applies to mail archived by the job that it is active for. The base folders and criteria are specified under the custom expiration dates accessed through the ‘Add’ button at the bottom. All messages included in any specified folders will have a different date or be exempted from the standard expiration date. To have messages with custom job or folder expiration dates properly expire, the deletion management date scope must be set to delete messages with an Expiration Date older than 1 day.

4.10.11 Mailboxes

The mailboxes tab contains the target users, and BES servers to be archived by the selected job.



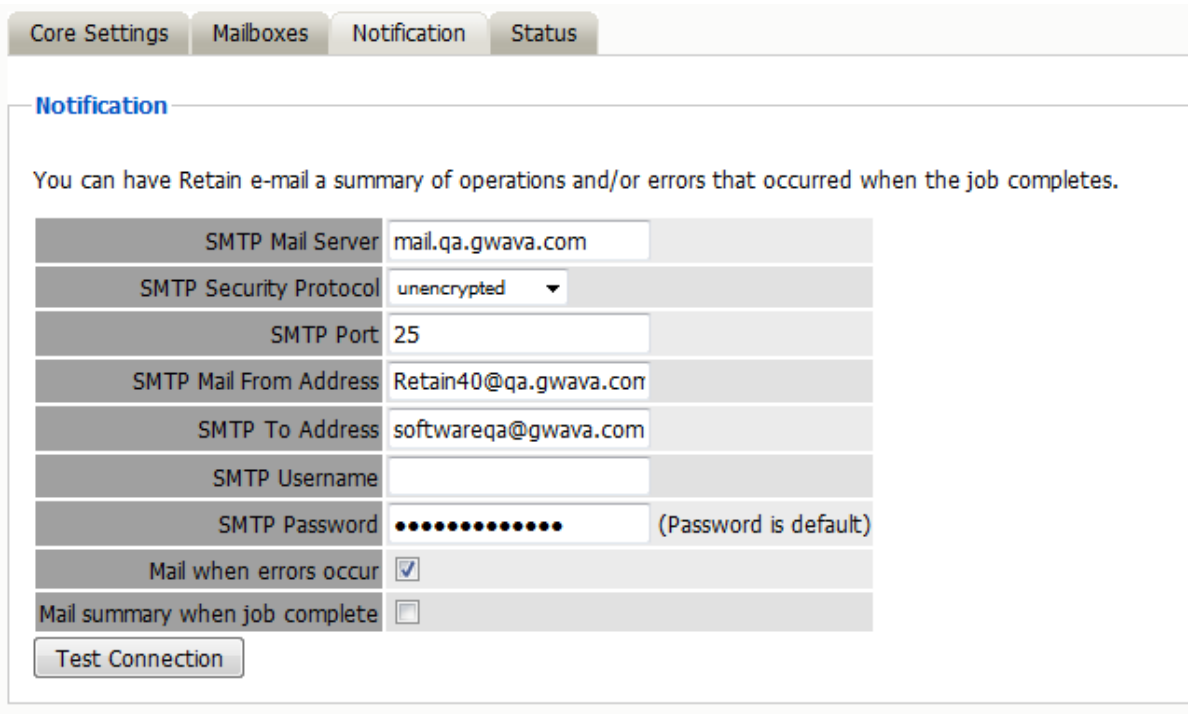
Under the Mailbox tab, select the BES server to be included in the Job. The list of available BES servers is configured from the Modules page. Once selected, all users associated or present on the BES server are included.

Once a job has been run, the user list will be populated and further configuration is possible. The job may also be restricted to specific users, specified on the ‘Include Users’ list, or restricted to exclude users listed on the ‘Exclude’ list. To add users, expand the User’s option box and select the users to be archived. Users may be selected through the user search interface, spawned through selecting the ‘add user’ but-ton.

If there are users listed, then the users listed in the ‘Include Users’ box are the ONLY included users, while those listed in the ‘Exclude’ window are excluded from the jobs.

4.10.12 Notification

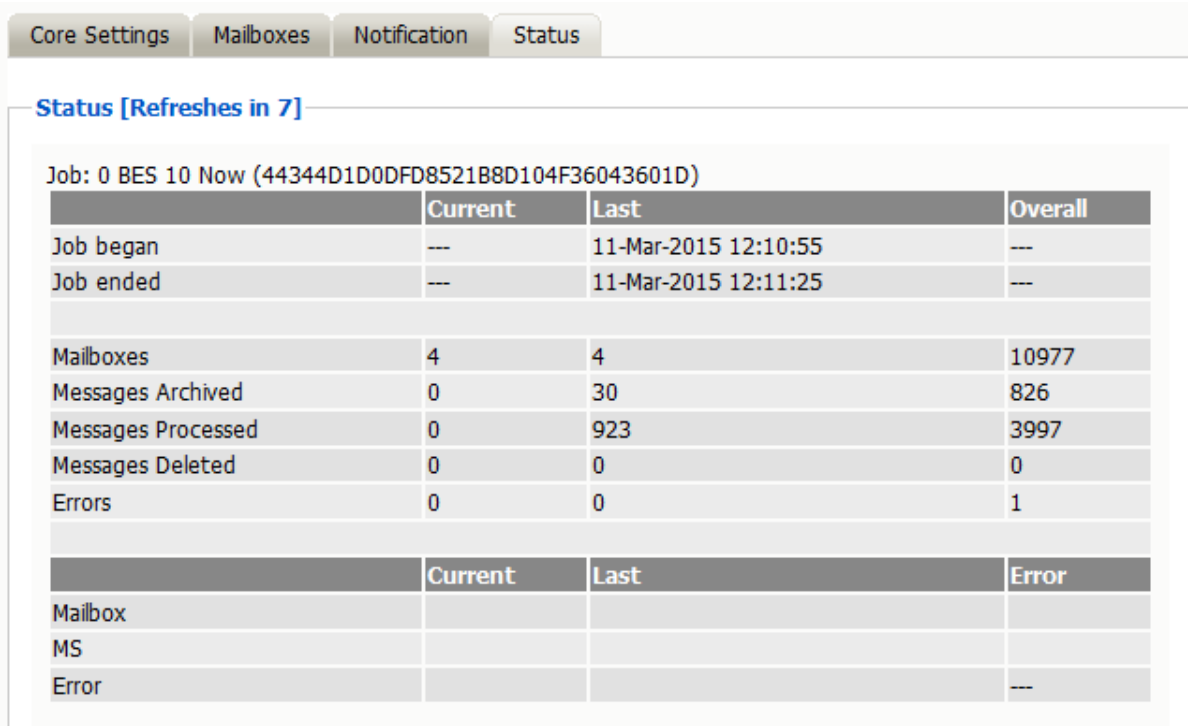
All errors, summaries, and statistics may be sent in an email notification to a desired address on job com-pletion. If it is desired to be notified, configure the notifications under this tab. At least one of the two notification options must be selected for an email notification to be created.



4.10.13 Status

On initial creation, there will be nothing of interest shown under the ‘Status’ tab, as the job is under cre-ation or has not yet run.

However, after a job has been run, or has begun, the ‘Status’ tab becomes a monitoring tool to report on exactly how much mail and mailboxes have been archived, along with any errors encountered.



	Current	Last	Overall
Job began	---	11-Mar-2015 12:10:55	---
Job ended	---	11-Mar-2015 12:11:25	---
Mailboxes	4	4	10977
Messages Archived	0	30	826
Messages Processed	0	923	3997
Messages Deleted	0	0	0
Errors	0	0	1

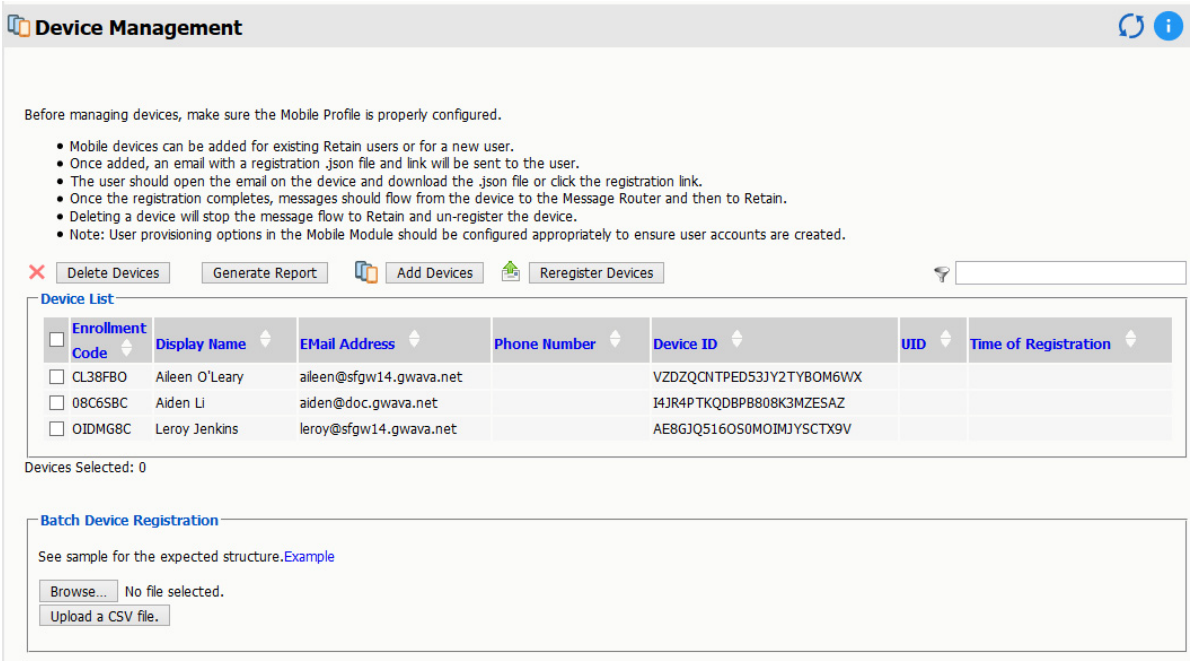
	Current	Last	Error
Mailbox			
MS			
Error			---

4.10.14 Device Management

Devices are deleted from Retain through the Device Management interface. Device Management is split into two sections: Mobile and Blackberry. In both sections, multiple devices may be selected and acted on. In addition, all selected devices remain selected until either an action is taken or the administrator has navigated away from the page.

4.10.15 BlackBerry

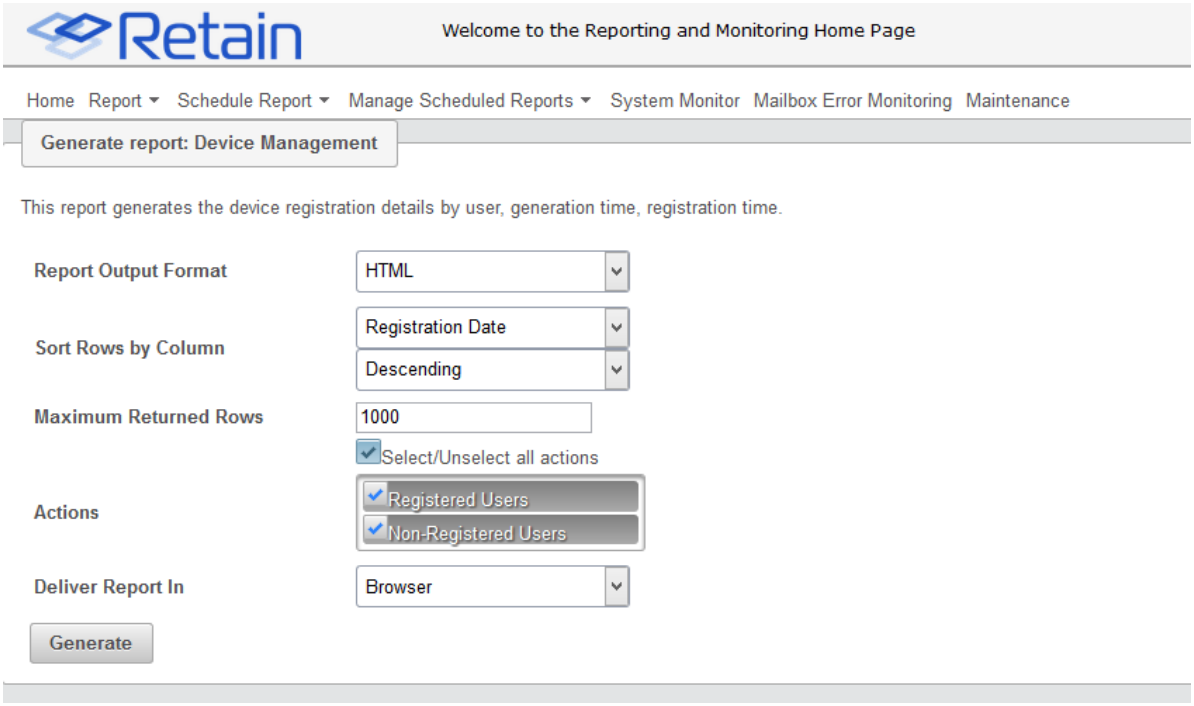
Retain supports archiving data from both BlackBerry devices and devices running the BBM Protected sys-tem. Devices running the BBM Protected system can have their communication data archived through the BBM app, obtained from BlackBerry.



The BlackBerry device page displays all devices which have been registered into the Retain system. Native BlackBerry and BBM Protected devices are displayed together. Retain registers BlackBerry devices through log files from the BES server or through data sent by the device running the BBM Protected appli-cation. If a device is not yet registered, it is either not in the BES address book or no data has been archived for the device. The device list may be filtered for any specified content desired. In addition, devices may be selected for removal.

The *Delete Device* button will remove the device from device management but will not remove the data.

The *Generate Report* button will open a new Reporting and Monitoring: Generate report window, where a report can be created.



You may now configure Schedules, Workers and Jobs.

4.11 BlackBerry Auditing and Archiving Service for BBM

The setup for Retain BlackBerry Auditing and Archiving Service for BBM Protected (BAAS) is quite simple for Retain, as the BBM Protected application directly inputs data to the Retain BAAS system. The BBM Protected app can forward data to either the Retain Router or the Retain Server, but it will only archive to one or the other. Because the BBM Protected app communicates directly to the Retain BAAS system, Retain needs to have an open connection to the Internet. Due to this requirement, the desired connec-tion should be placed in the network's DMZ. To protect the Retain Server, the Retain Router has been designed to be placed in the DMZ. When the Router cannot be installed in the system, the Retain Server can handle all communication, but will need to be installed in the DMZ. Whichever will be used, Router or Server, the configuration and a security certificate must be obtained before setup can be completed.

This supports the BBM App for Desktop (Windows and macOS).

Item and Requirements Checklist:

- Enterprise ID (From Blackberry)
- Base API license (In licensing tab)
- BlackBerry license (In Licensing tab)
- REST API base License (In REST API tab)
- BBM REST API license (In REST API tab)
- SSL security certificate obtained from a trusted third party certificate authority for either Server or Router (whichever is to be used).

Pre-setup tasks:

- Network structure plan (Where to place the Retain Server, and or to use the Router.)
- Retain Server setup and configured – or – Server and router configured
- Router or Server connected to, accessible in DMZ

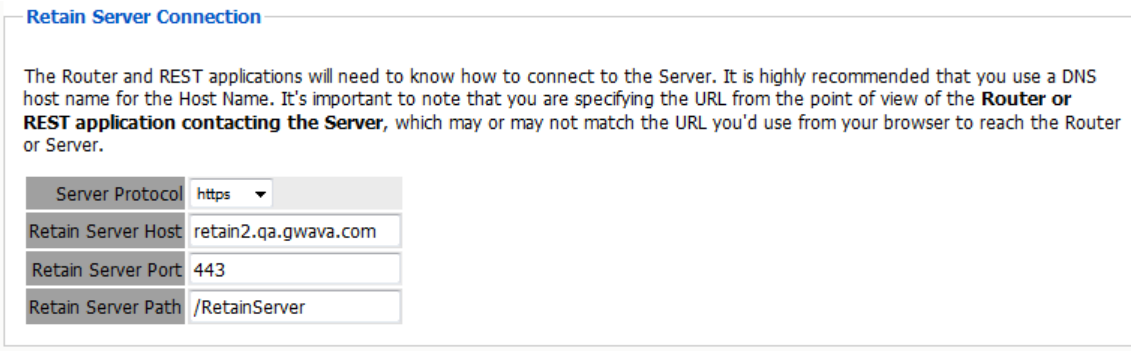
Setting up the Retain environment:

1. License Retain.

   Retain requires a total of 4 licenses to enable Retain BAAS. The first two licenses, base Retain license and the Blackberry module license, are uploaded to the Licenses page in the Retain Server management console. The second two licenses are REST API licenses, REST API base license and the Blackberry REST API license, are installed under the Server Configuration | REST API tab.

2. Configure Retain server connection.

   Found under Server Configuration | Communications tab, the Retain server connection info at the bottom of the page. This connection setting is the information which will be used by the REST applications, (in this case BBM Protected application), or the Retain Router, to connect to the Retain Server. If this is not correct, the Router or the REST applications will not be able to connect the Retain Server. This must be configured before moving on to step

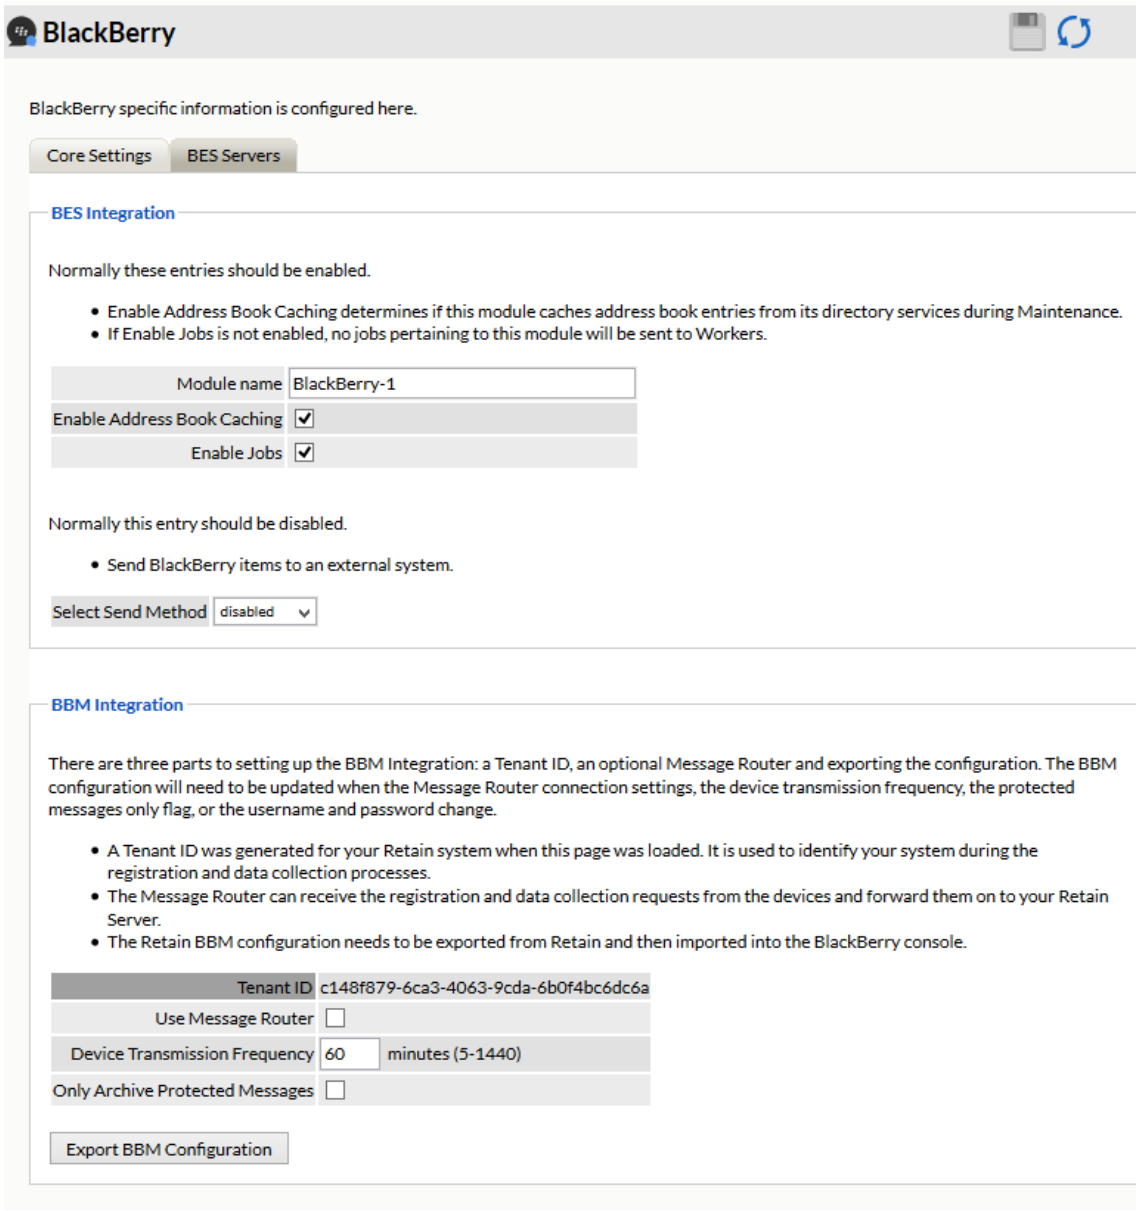


3. From Blackberry module, Export BBM configuration file.

   The BBM Configuration file is generated through the Blackberry module configuration page. Open the module configuration and look at the Core Settings tab. Look at the BBM Integration section. Here the decision is made whether to use the Retain Router or to use the Retain Server. If the Retain Router is to be utilized, select the 'Use Message Router' checkbox.

   The Device Transmission Frequency setting is the setting which determines how often the BBM Protected application will upload archive data to the Retain Server. If there is no information to be archived, the BBM Protected application will wait until there is. The setting is in minutes. Minimum is every 5 minutes, maximum is 1440 minutes (24 hours). Once configured, select the 'Export BBM Configuration' button and save the file; you will need it later.

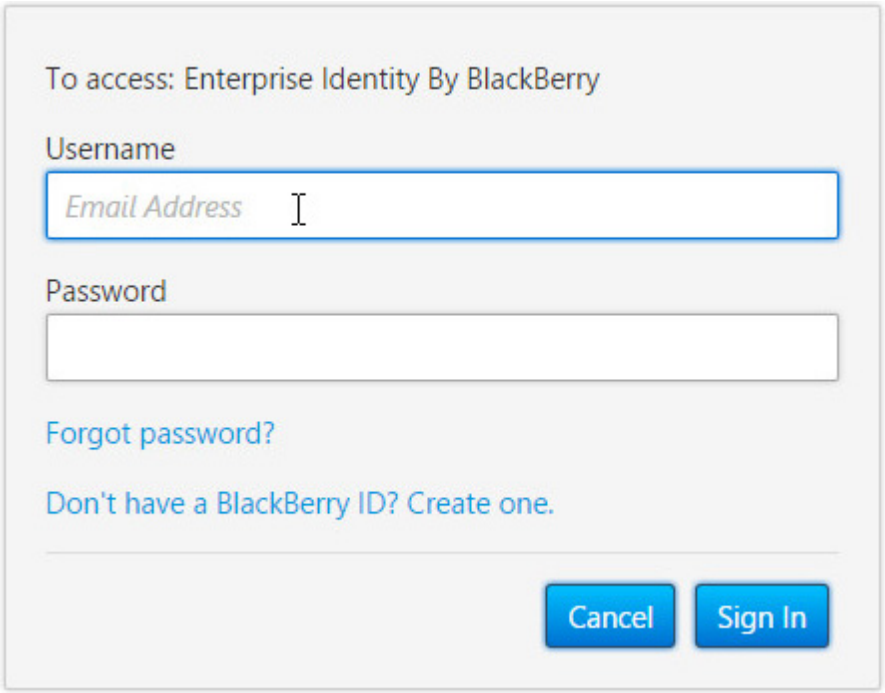
   The module may be given a name. The Send Method option enables either the SMTP Forwarding or FTP features. For either feature to appear and function, the Module Forwarding tab must be configured on the Server Configuration page. See that section for more information.



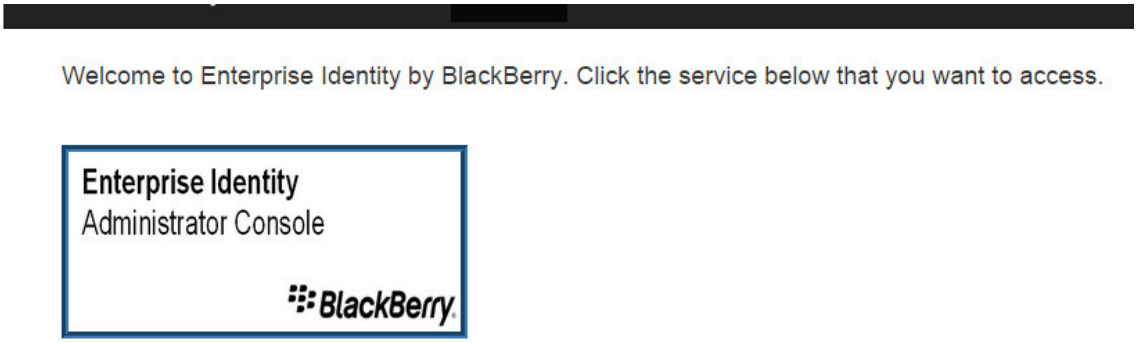
4. Obtain a SSL security certificate for destination server/or router (if not already obtained)

   The security certificate must be obtained to configure the Blackberry Identity console. The security certificate needs to be in base 64 format.

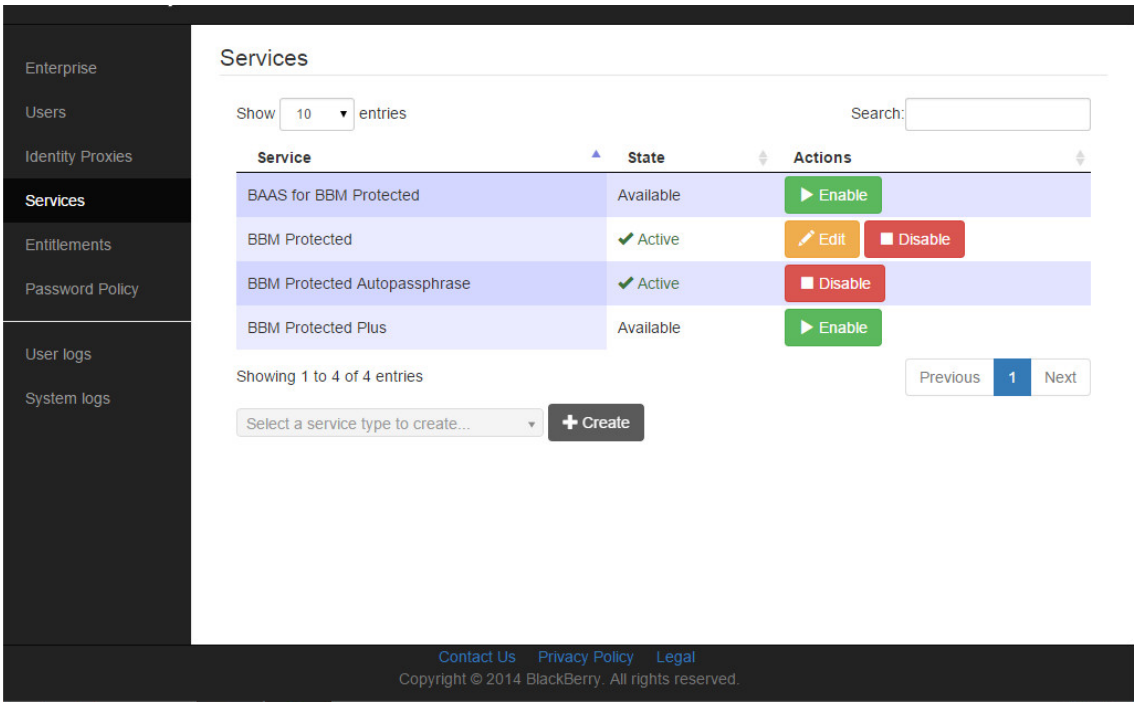
5. Go to the Blackberry Enterprise Identity Console. (https://idp.blackberry.com/enterprise



6. Enter the credentials provided by Blackberry when account was created.
7. Enter administrator console.



8. Select the ‘Services’ page.
9. BBM protected must be enabled.



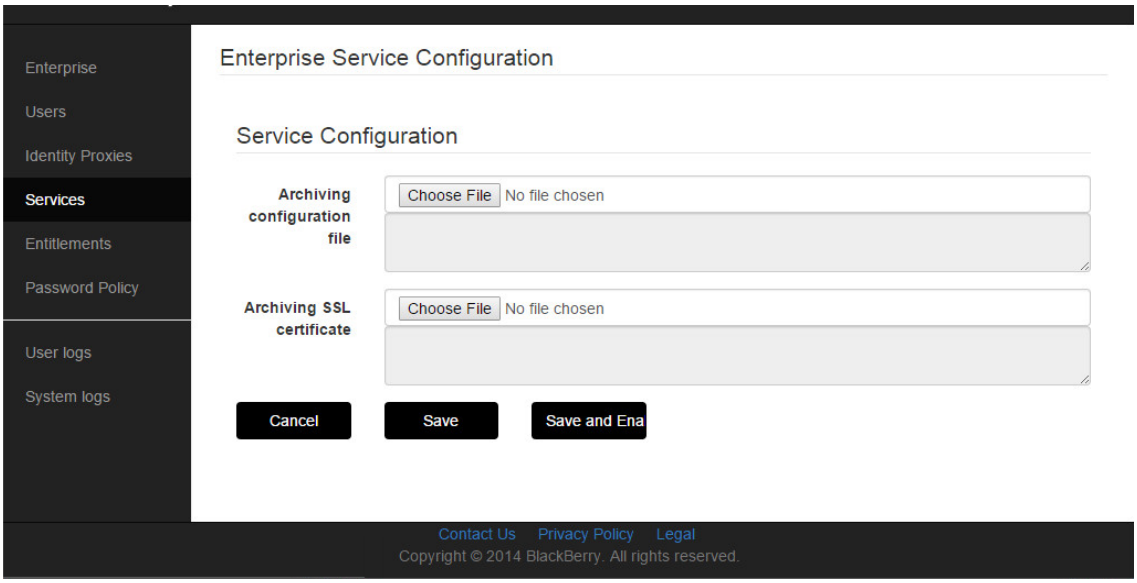
10. Enter ‘entitlements’ and invite users/devices to be added to the service (In this case BBM Protected)

    Users must first be added to the BBM Protected system before they can be added to Retain BAAS.



11. Enable BAAS for Protected.

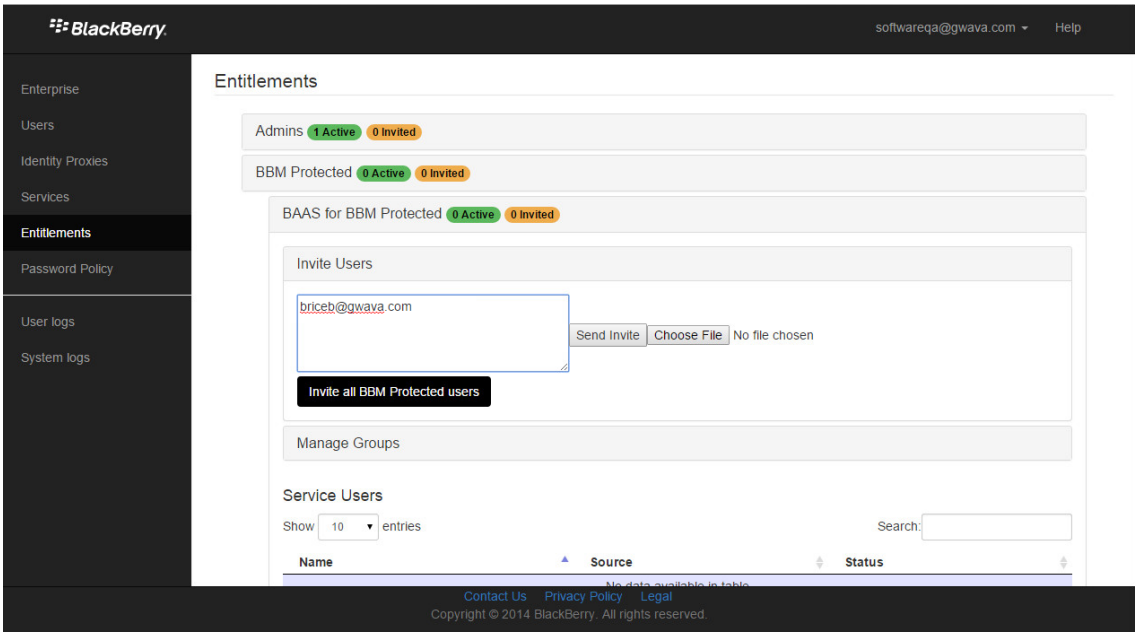
    To enable BAAS for protected, from the ‘Services’ tab, select the ‘enable’ button for BAAS. There are two files which must be uploaded to the Blackberry Identity Console, these are the two files saved earlier: the full chain SSL security certificate for the Retain Server or the Retain router, (whichever is going to be used), and the BBM Configuration file created in step 3.



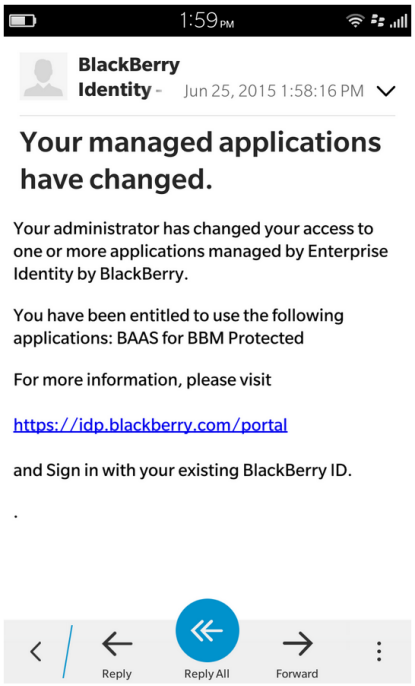
The Archiving configuration file is the BBM Configuration File, and the full chain SSL certificate needs to be uploaded to the ‘Archiving SSL Certificate’ section. Select the ‘Choose file’ button and browse to the appropriate file to upload it. Once both files have been uploaded, select the ‘Save’ button.

12. Enter Entitlements and invite users

    Users invited to BAAS MUST already be users of BBM Protected. If users are not registered to use BBM Protected, BAAS will not work. First invite them to BBM Protected, and then invite to BAAS after the device has been registered with the BBM server for BBM Protected service.



13. Users will be notified that they have been added to the archiving



Once the devices and users have been invited to join the BAAS program, they will be sent a notification that they are now part of the BAAS archive.

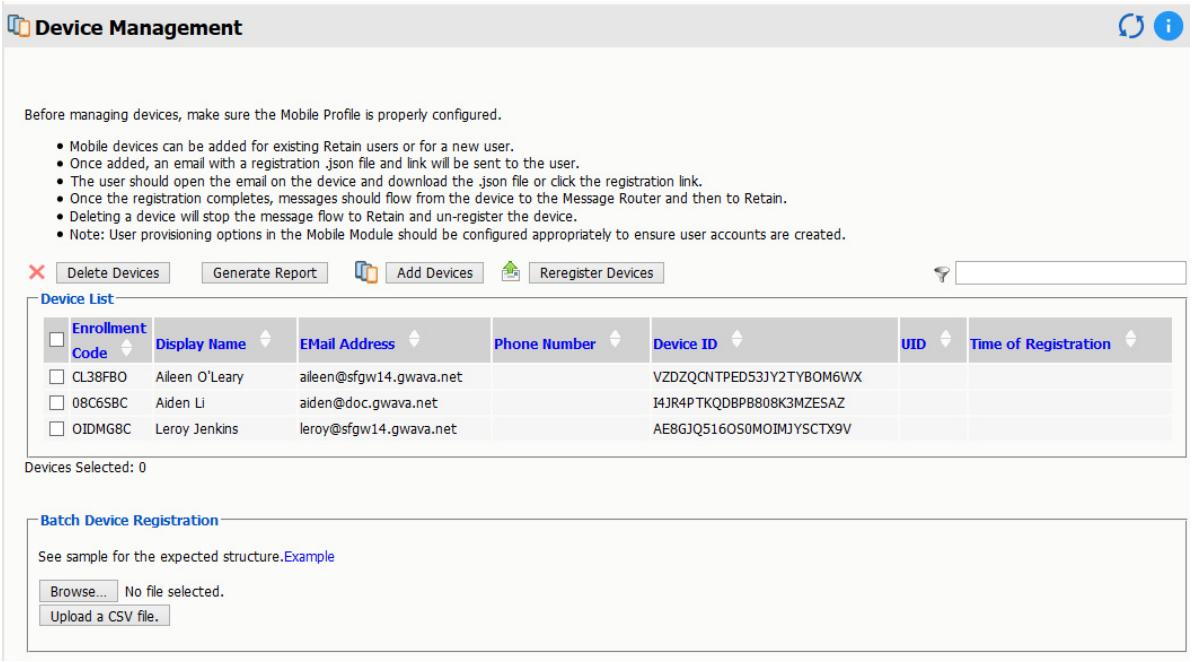
They are now part of the BAAS system and their BBM Protected communication will be added to the Retain BAAS archive.

4.11.1 Device Management

Devices are deleted from Retain through the Device Management interface. Device Management is split into two sections: Mobile and Blackberry. In both sections, multiple devices may be selected and acted on. In addition, all selected devices remain selected until either an action is taken or the administrator has navigated away from the page.

4.11.2 BlackBerry

Retain supports archiving data from both BlackBerry devices and devices running the BBM Protected sys-tem. Devices running the BBM Protected system can have their communication data archived through the BBM app, obtained from BlackBerry.



The BlackBerry device page displays all devices which have been registered into the Retain system. Native BlackBerry and BBM Protected devices are displayed together. Retain registers BlackBerry devices through log files from the BES server or through data sent by the device running the BBM Protected appli-cation. If a device is not yet registered, it is either not in the BES address book or no data has been archived for the device. The device list may be filtered for any specified content desired. In addition, devices may be selected for removal.

4.12 Google Apps Module

The Google Apps module for G Suite allows Retain to archive Gmail data items. To configure Retain for Gmail archiving, Retain needs Gmail to be configured to allow Retain access, and the appropriate information entered into Retain.

Google Apps requires that a project be created, an OAuth key created and a Service Account specified and enabled before the Retain system can connect and archive mail.

To archive from Google Apps:

- Create a project
- Create a new Client ID key
- Record client secret
- Create a Service Account and Manage API Client Access
- Generate a P12 key certificate

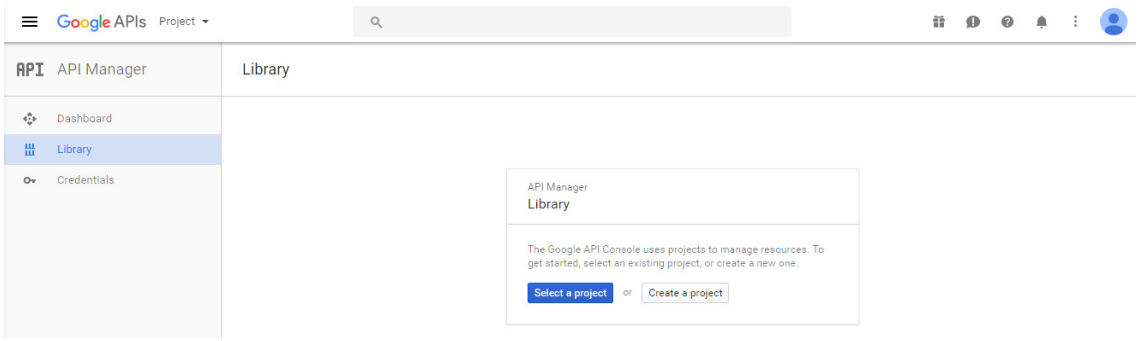
4.12.1 Prerequisite

Retain supports two-factor authentication with OpenId Connect for OAuth 2.0. To utilize OpenId the Retain Server needs to be accessible from the Internet. The URL through which the Retain Server is accessible from the internet must be specified. For this, you must create a Project and Client ID.

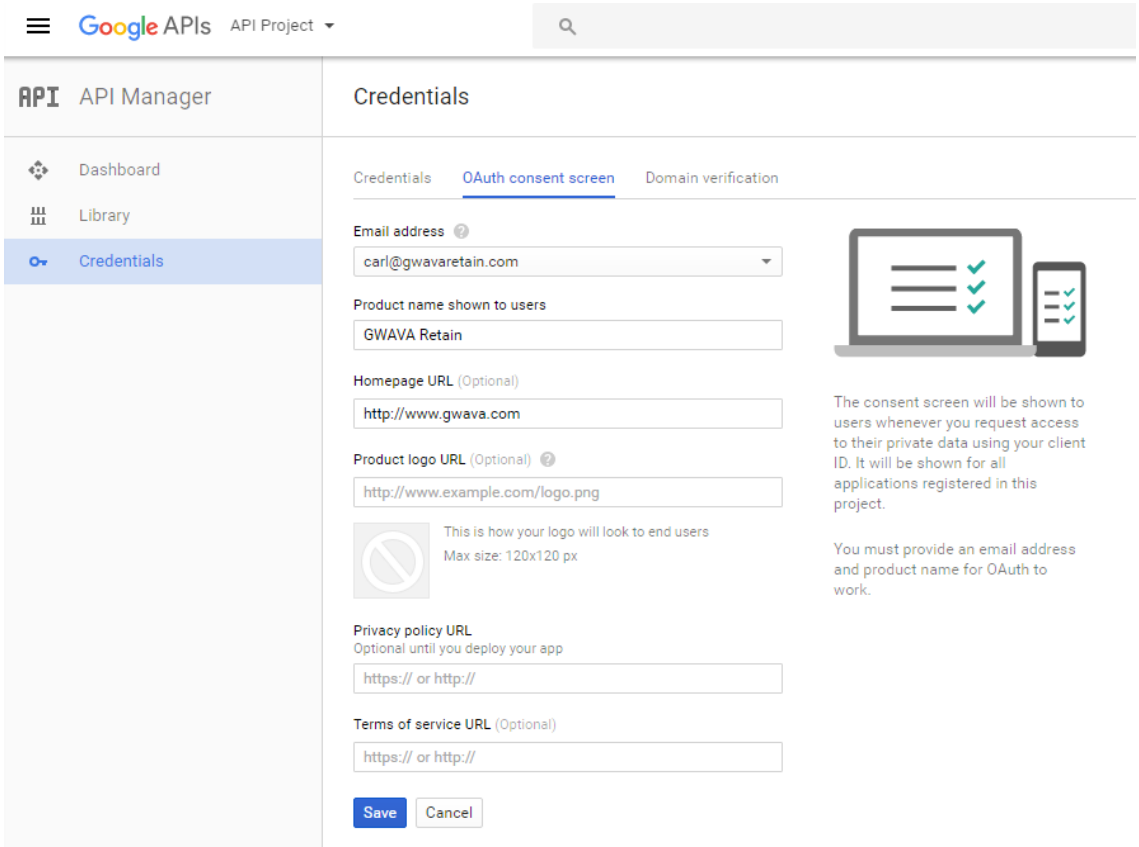
NOTE: As a cloud service these screens may change at anytime.

To configure OpenID for Retain use:

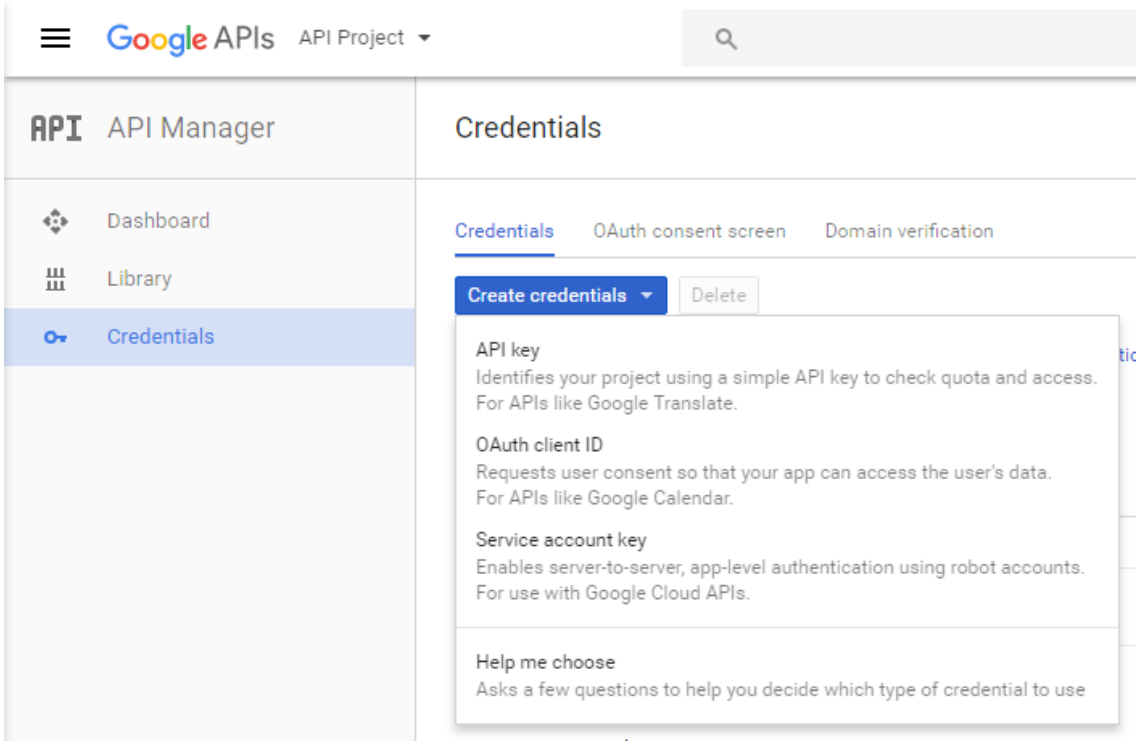
1. Go to the Google API Console, and select 'Create a project'



2. In the sidebar under "API Manager", select 'Credentials' and then select 'OAuth consent screen tab
3. Choose an 'Email Address', specify a 'Product Name', and press 'Save'



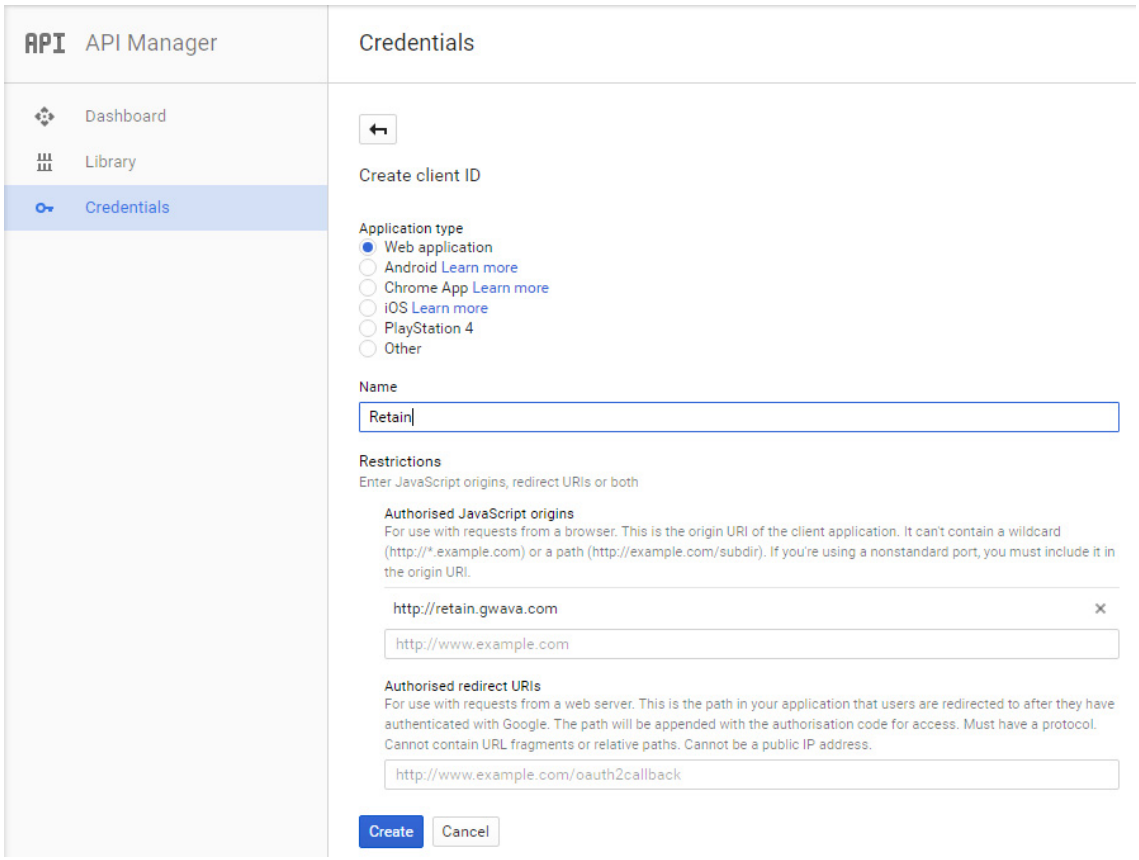
4. In the 'Credentials' tab, select the 'Create credentials' drop-down list, and select OAuth client ID



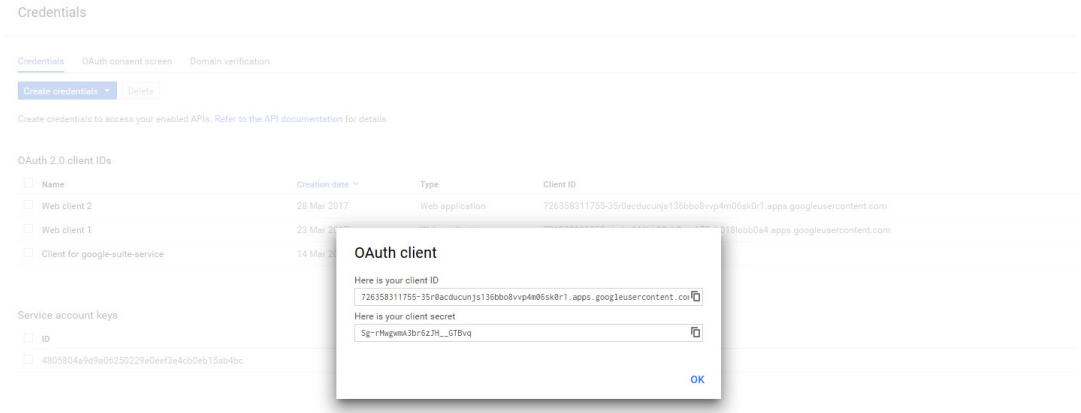
5. Under 'Application type', select 'Web application' and Specify a name. The origin field should be the Retain Server's URL.

   If using two-factor authentication for GMail login, enter an Authorized redirect URL (e.g. http://retain.gwava.com/RetainServer/Server/openIdConnect.jsp?).

   When complete, select the 'Create' button.

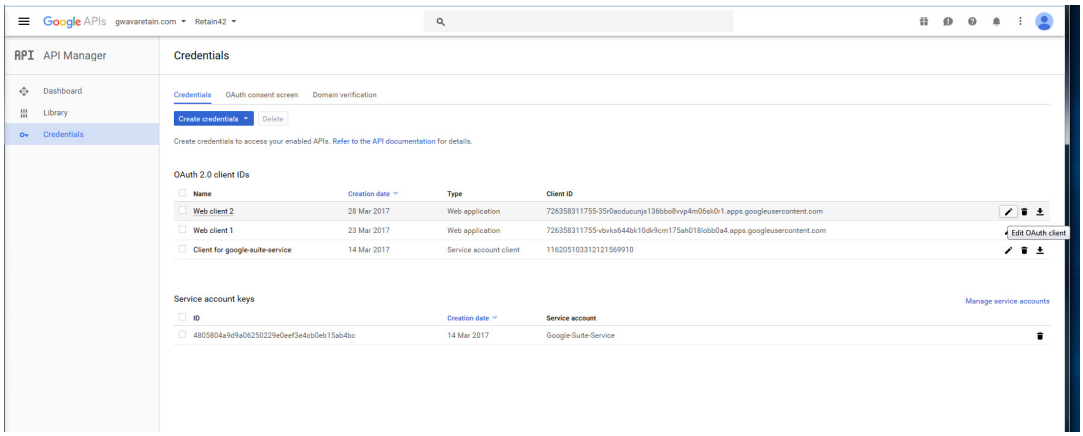


6. The OAuth client dialog box should be displayed. This dialog box contains the Client ID and Client secret.
   - Copy the Client ID and paste it into the Retain Client ID field in the Google Apps module configuration. Do not lose the client secret.
   - OAuth access requires the secret and ID. *This is the only time the secret will be displayed*.

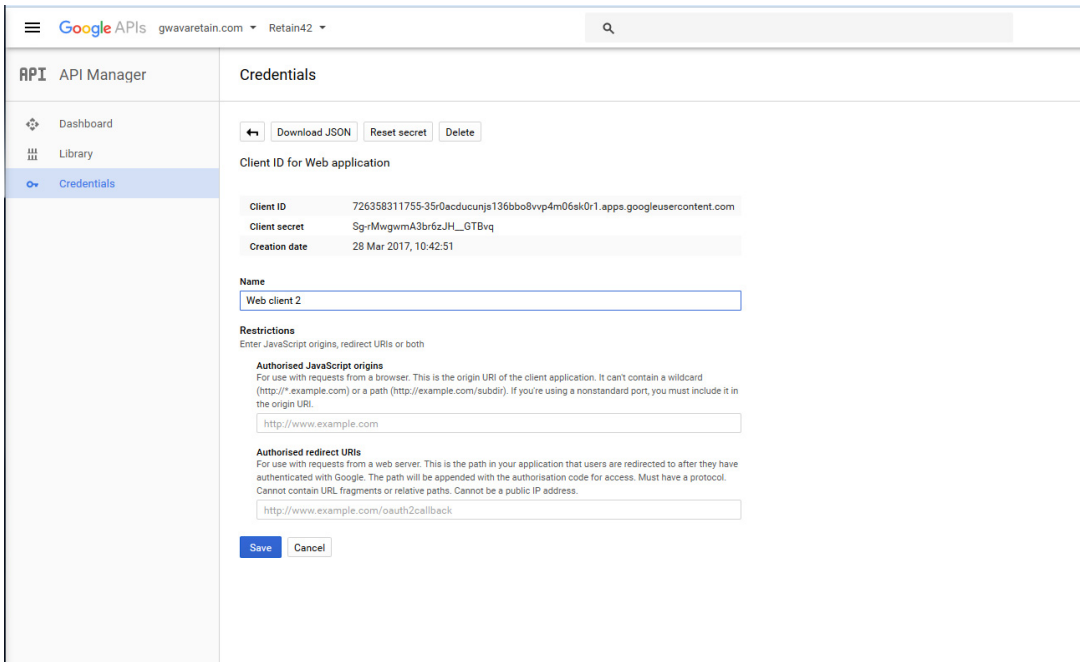


6a. To attain the client secret for an existing project:

a. Select the Web Client and click the *Edit OAuth Client* button

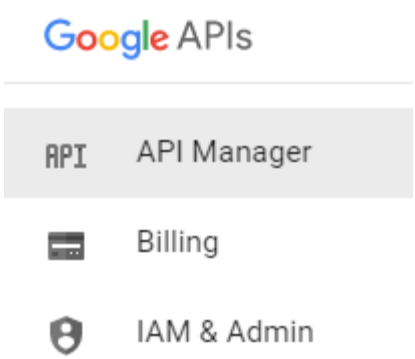


b. Copy the Client ID and Client Secret which is required to configure OpenId tab for Google apps module.

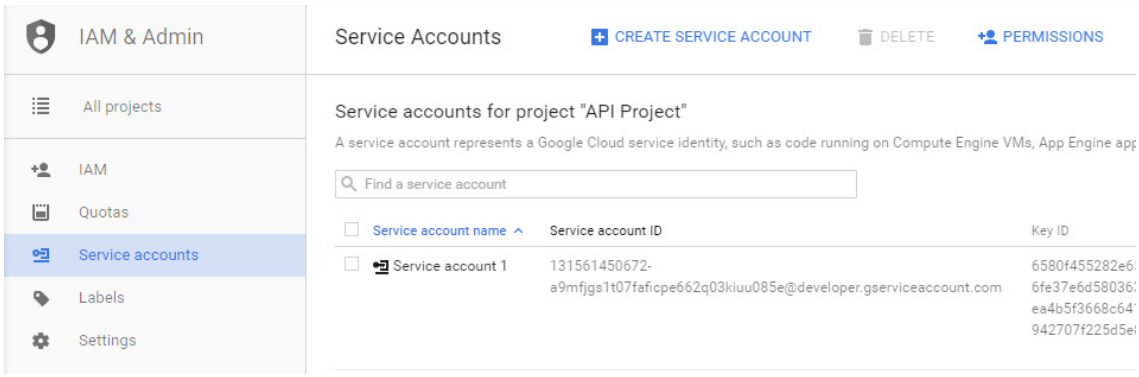


The OAuth Client needs to be enabled for domain-wide delegation to function. To enable domain-wide delegation, Retain requires a service account.

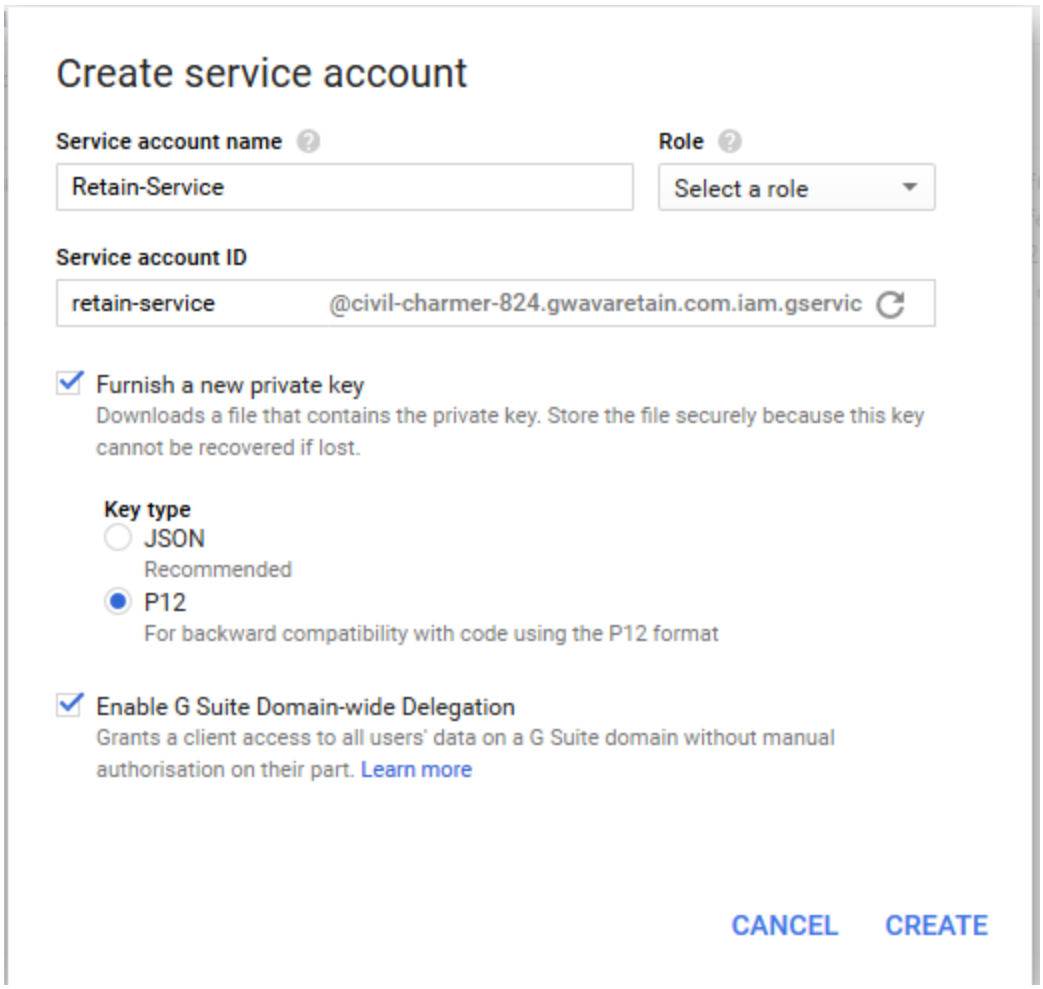
7. From the Products and Services hamburger menu at the top-left: Select IAM & Admin



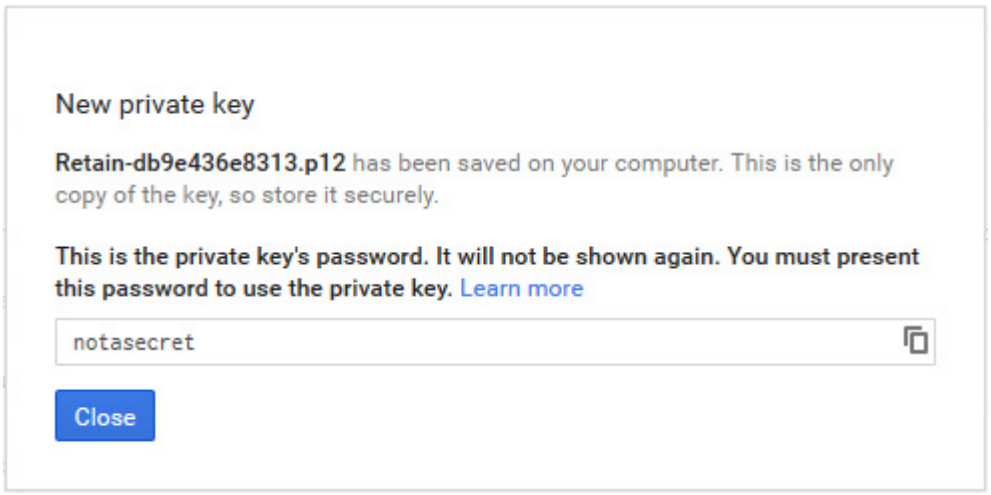
8. Select the 'Service Accounts' and click the 'Create service account' button



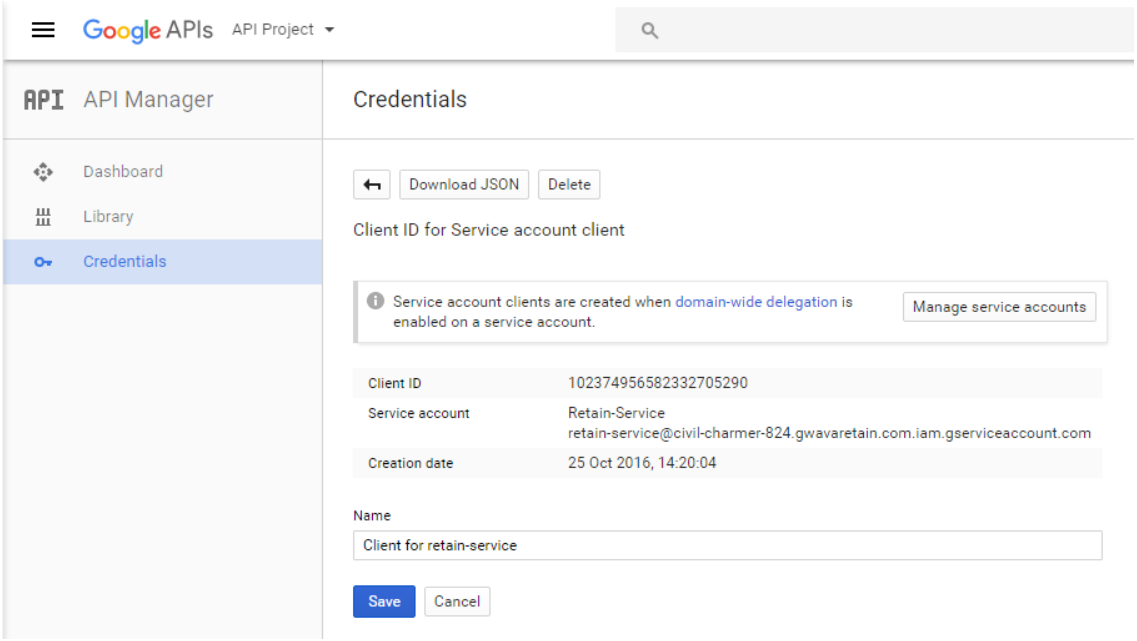
9. Configure the Service Account Name and ID. No Role is needed.
   - Select the option to Enable Google Apps Domain-wide Delegation.
   - Enable "Furnish a new private key" and select the P12 file.
   - Copy the Service Account name and ID. These are required by Retain.
   - Click 'Create'.



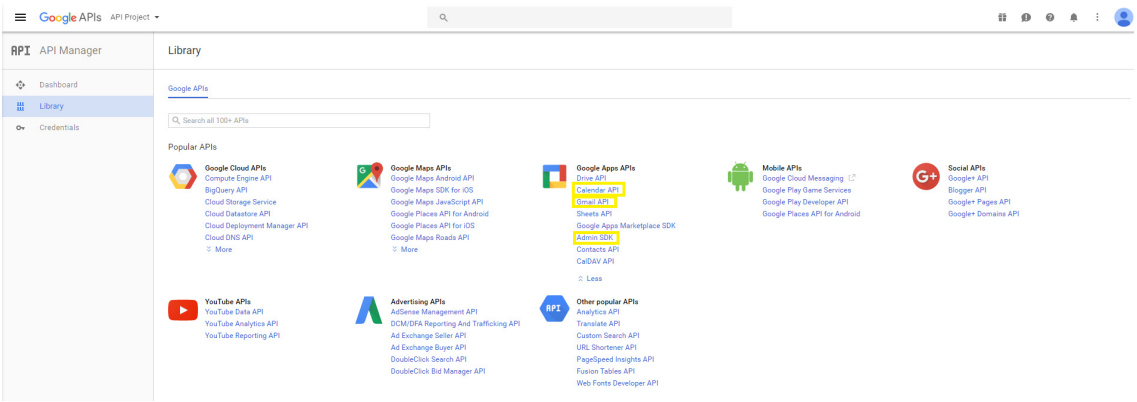
The key will be automatically created and downloaded to the local machine. Do not lose the P12 file. This is the only copy and it is required to configure Retain.



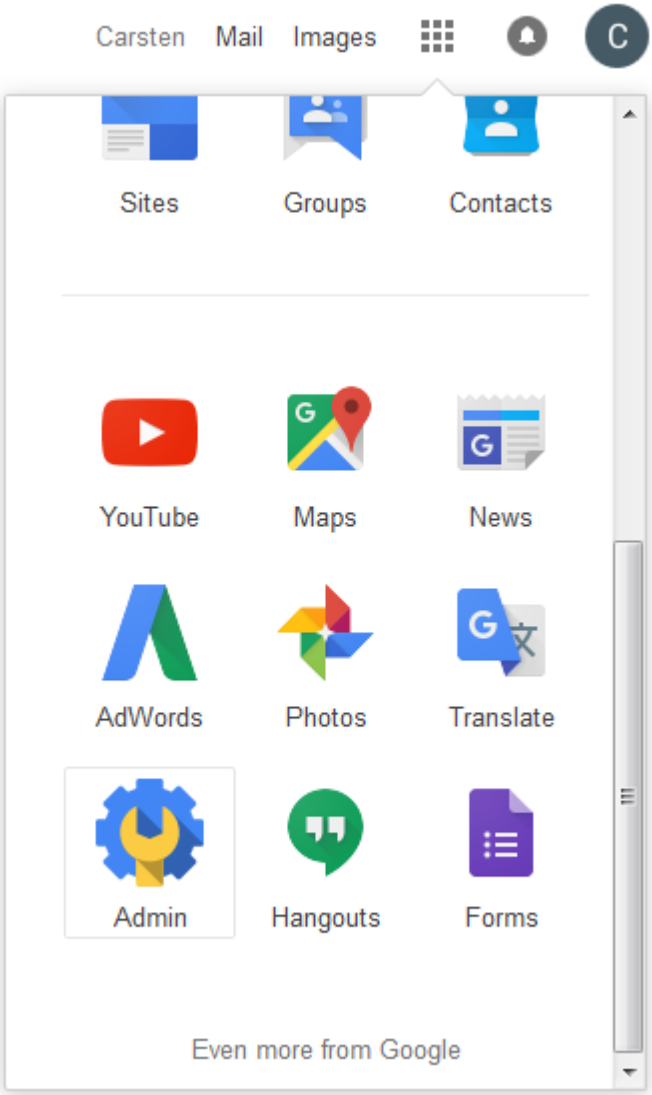
10. You will need the Client ID from the Service Account



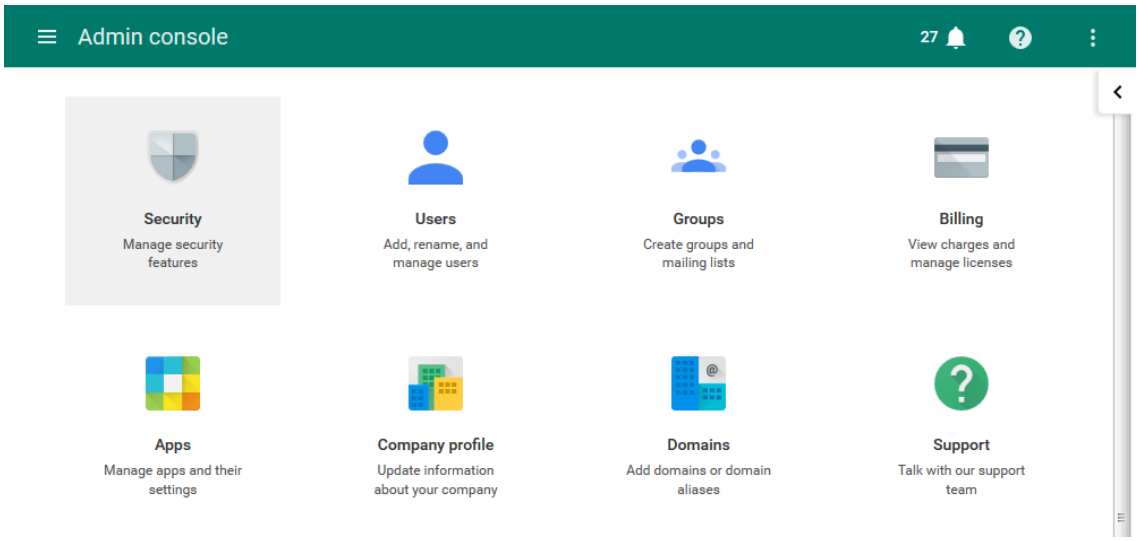
11. Go back to the API manager and enable: Calendar API, Gmail API and Admin SDK. Click on each one and click Enable on top.



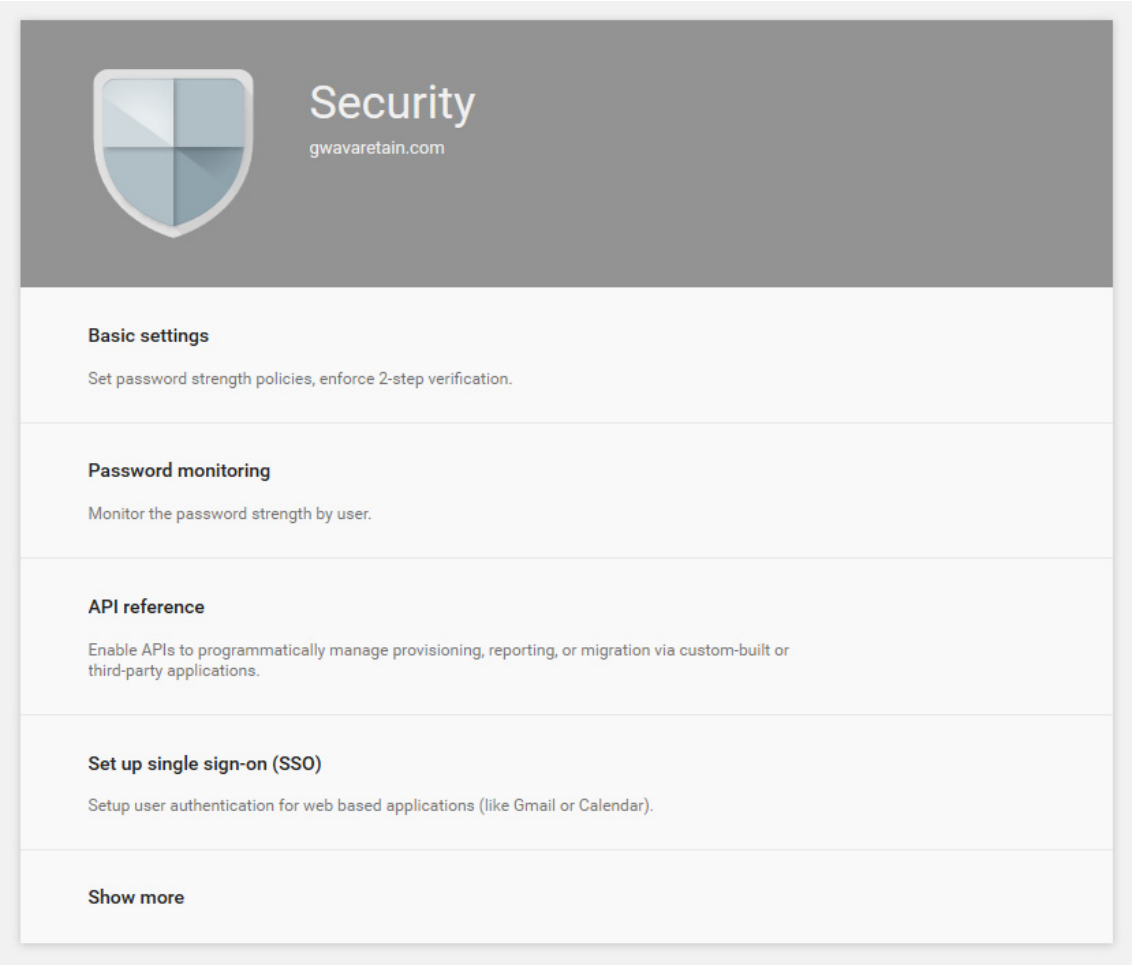
12. The last step is to authorize the domain. Browse to the Google home page and select 'Admin' from the drop down menu at the top right.



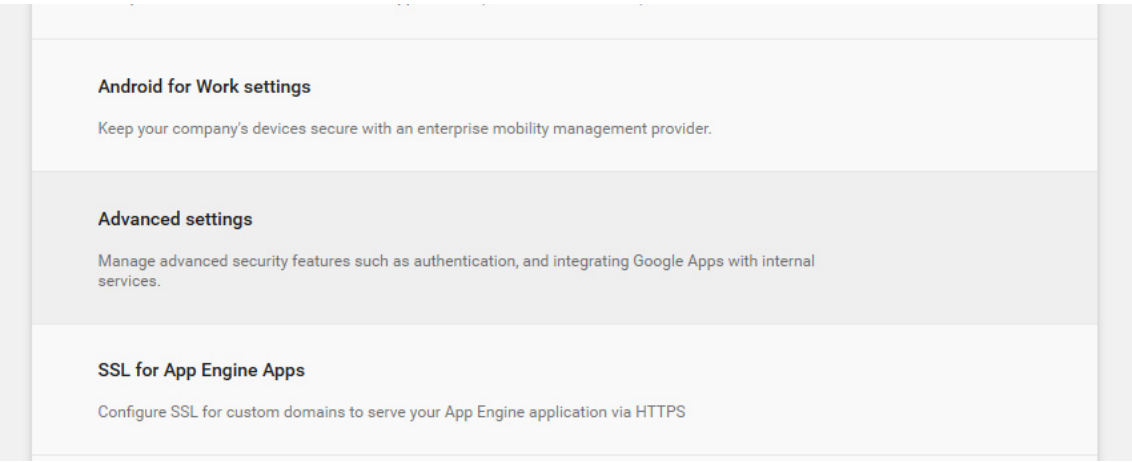
13. Select 'Security' from the administration menu.



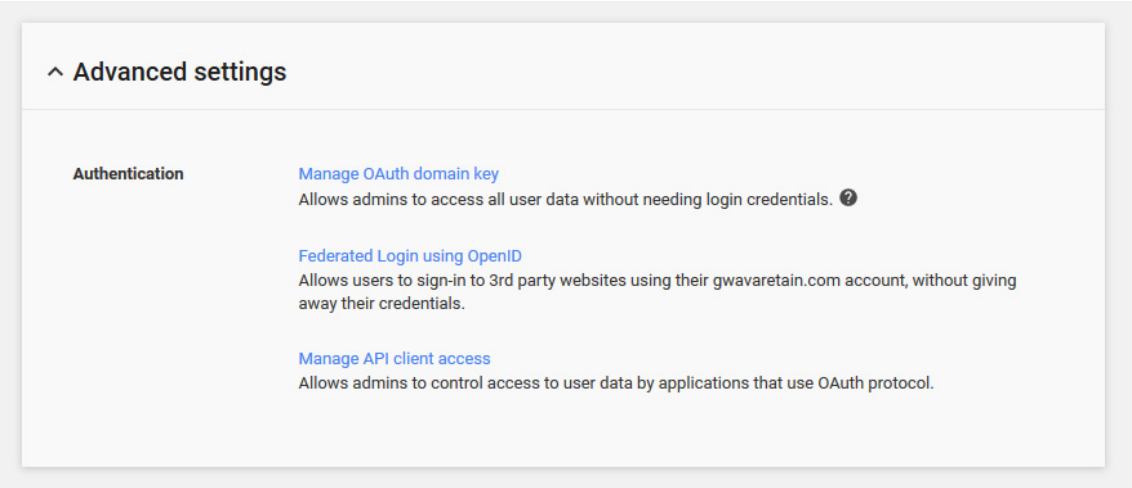
14. From the Security menu, click 'Show More'



15. Select the 'Advanced Settings' option.



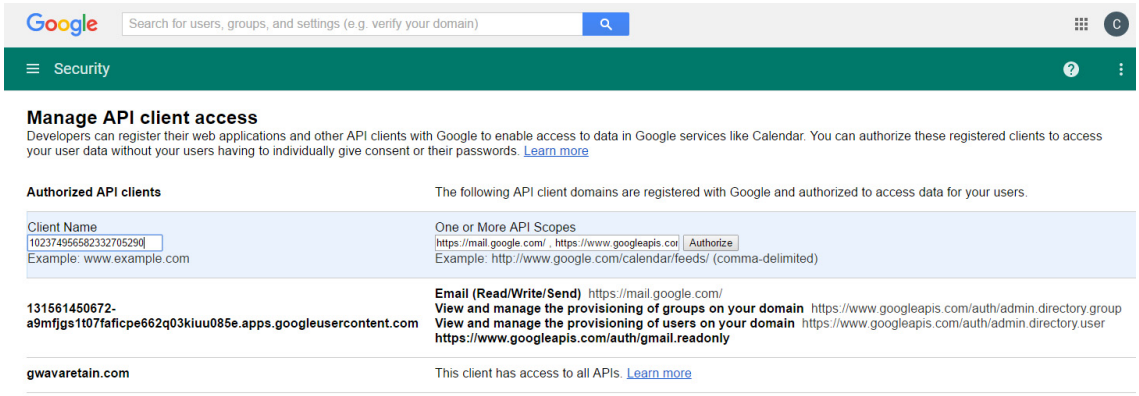
16. On Advanced Settings, select 'Manage API client access'



17. In this window, enter the client ID created with the service account, and then input the entire domain the client will be accessing.

    The Service Account name would be like Retain-Service as above and the API Scopes (listed in the module) would be:

    https://mail.google.com/ , https://www.googleapis.com/auth/admin.directory.group ,https://www.googleapis.com/auth/admin.directory.user ,https://www.googleapis.com/auth/gmail.readonly



18. Once entered, click 'Authorize'.

4.13 Google Apps Module Setup

All configuration required is now complete in Google Apps, and the information and files keys are available for configuration of Retain.

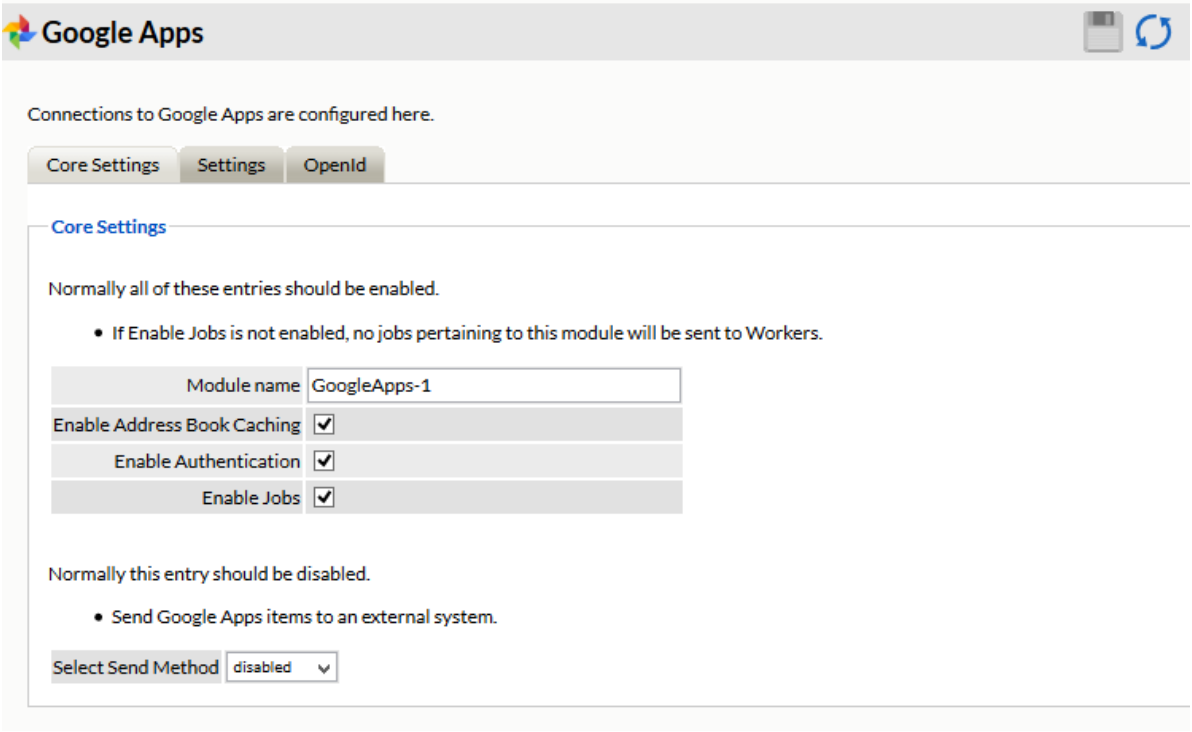
4.13.1 Core Settings Tab

Once Google has been configured to allow access to Retain, the Google Apps Module may be configured.

The module needs to be enabled on this page to make it active in the Retain system.

The module can be given a name.

The Send Method option enables either the SMTP Forwarding or FTP features. For either feature to appear and function, the Module Forwarding tab must be configured on the Server Configuration page. See that section for more information.

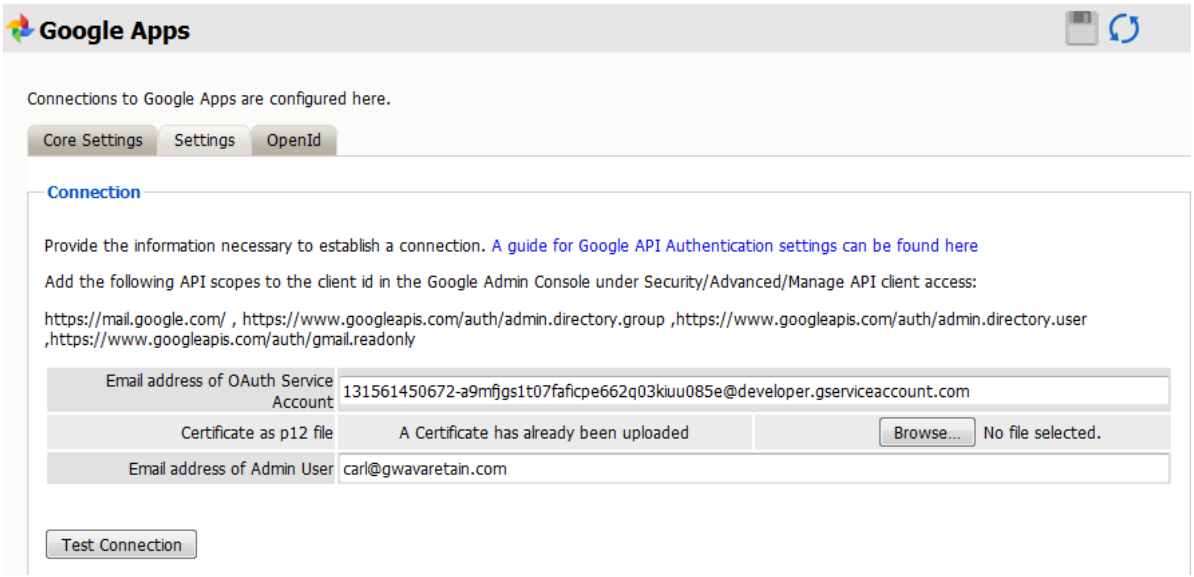


The Send Method option enables either the SMTP Forwarding or FTP features. For either feature to appear and function, the Module Forwarding tab must be configured on the Server Configuration page. See that section for more information.

Address book caching must be enabled to gather and maintain an updated list of users. Authentication is used to allow access to the Retain message store for users based on their existing Gmail account login. If the Enable Jobs option is not enabled, no jobs may be completed with the Google Apps module.

4.13.2 Settings Tab

In order to connect to the Google system, Retain requires the email address of the Admin user OAuth Service Account and a p12 Certificate for authentication. Retain archives the Gmail system through IMAP, and will login and download the message data to the Retain data store.

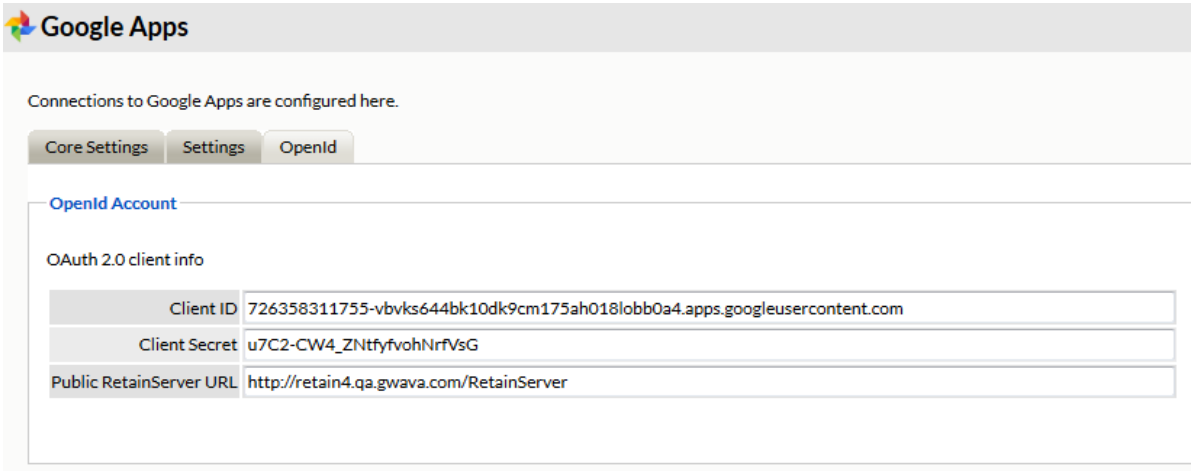


Test the connection to ensure that the configuration has been completed correctly.

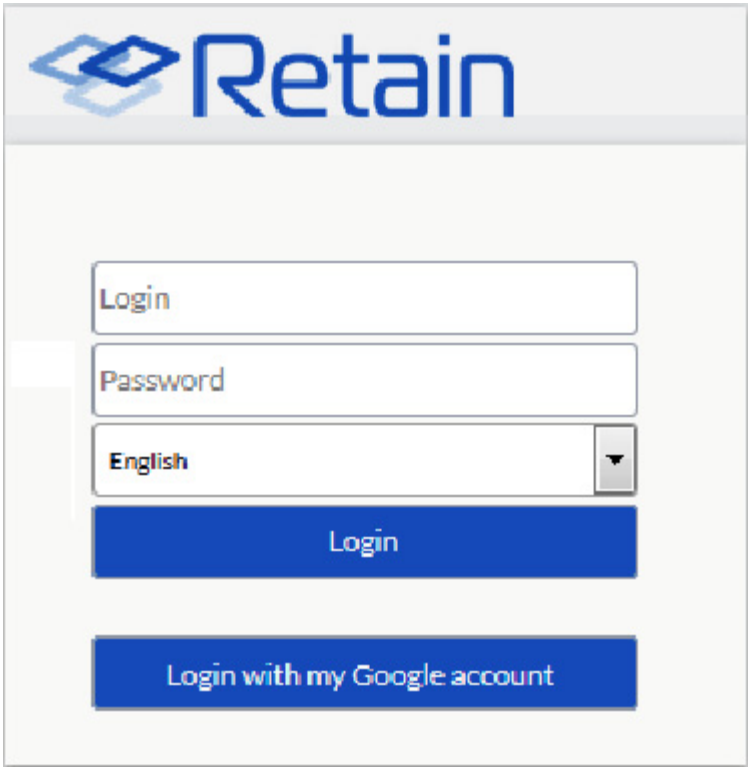
Jobs and profiles for Gmail will not be visible until the address book has been cached. After the module has been cached, all configuration options for profiles, workers, schedules, jobs and data storage will be enabled and visible.

4.13.3 OpenID Tab

Configure the OpenId Connect tab in Retain by inputting the Client ID, the Client Secret, and the Public RetainServer URL and saving changes. NOTE: The public RetainServer URL should look something like http://<yourdomain.com>/RetainServer. Only specify to the ".../RetainServer" portion of the URL, the rest is automatically filled-in. Specifying the complete URL will result in a connection error.



If Retain has been configured with OpenId credentials, the login page will display an option to login with Google credentials. If a user is currently logged-in to their Gmail account, simply clicking the "Login with my Google account" button will automatically log them into Retain.

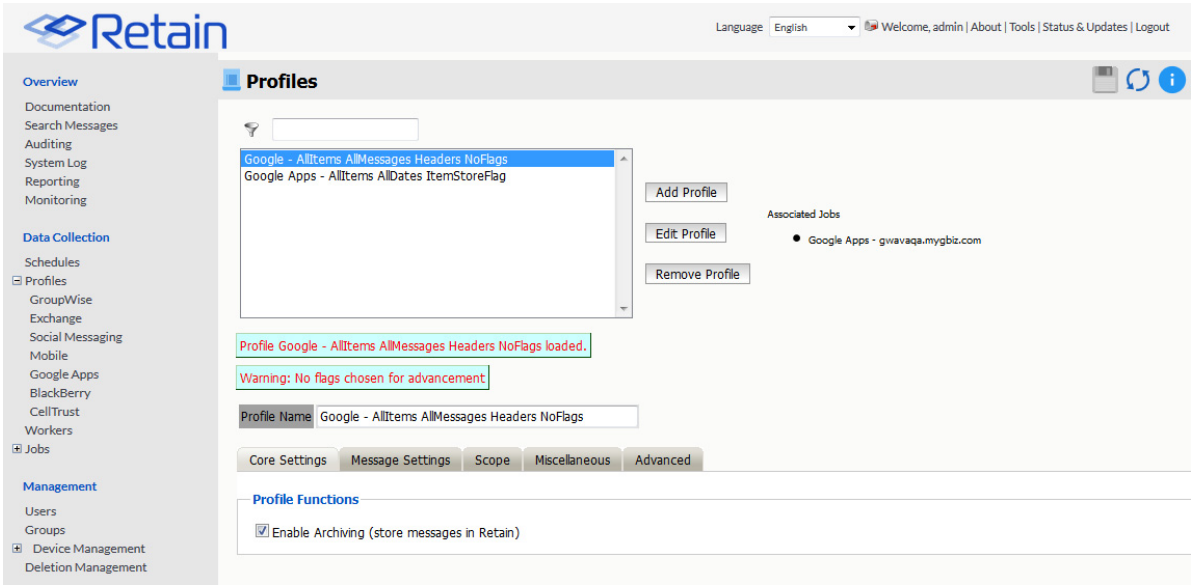


4.13.4 Google Apps Profile

To create a profile to archive email and data from Google Apps, select the add profiles button and name the profile then continue configuring as desired. The Google Apps profile must be configured and all set-tings saved before a job can be created and run.

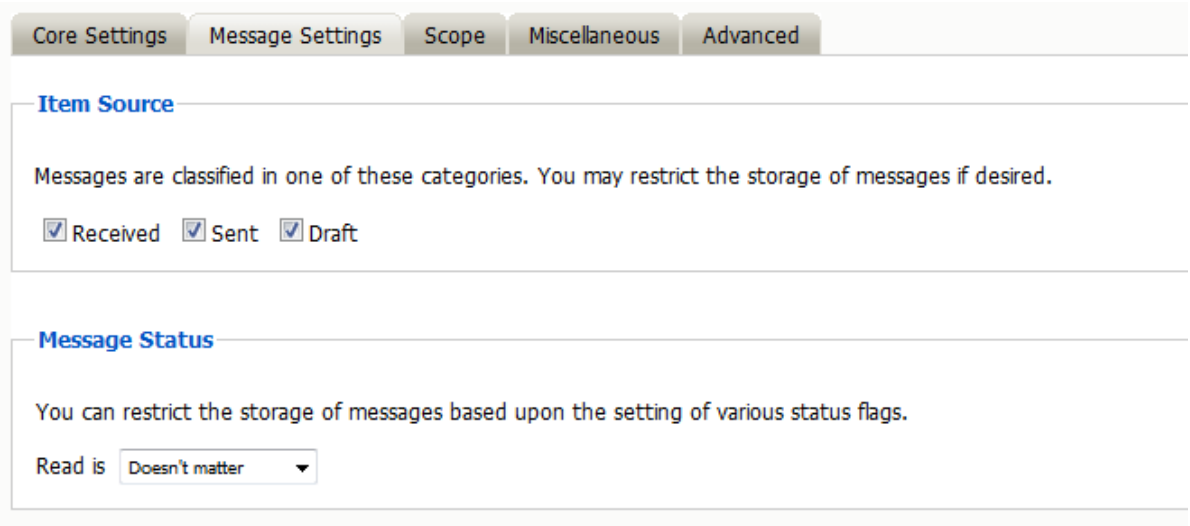
4.13.5 Core Settings

The profile will not become active and will not allow jobs to be run unless the profile is enabled. Enable the profile by placing a check mark in the ‘Enable Archiving’ checkbox.



4.13.6 Message Settings

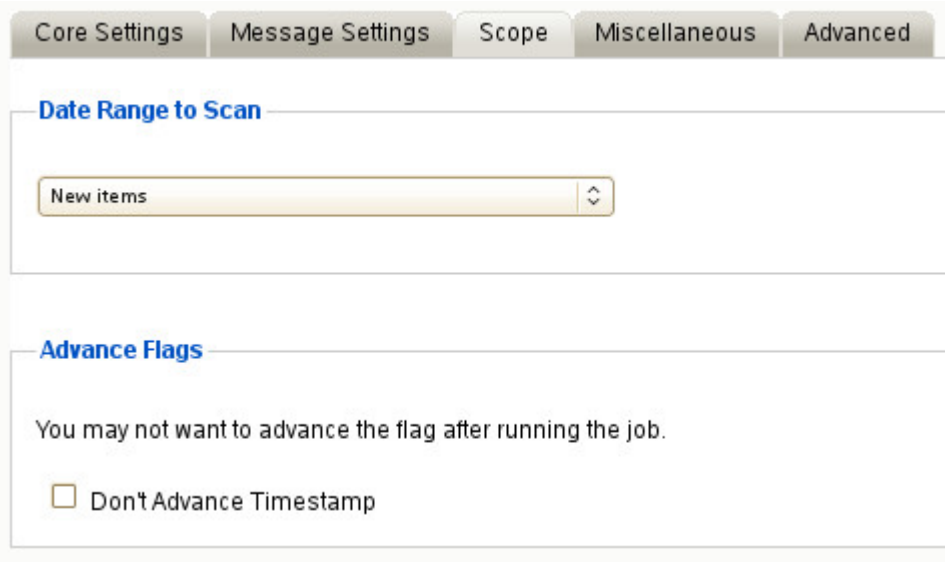
The Message Settings tab contains the source and status settings for the messages to be archived. Mes-sage types and sources which are checked will be archived in this profile.



Any item type or status which is not selected or specified will exclude qualifying items from being archived in Retain.

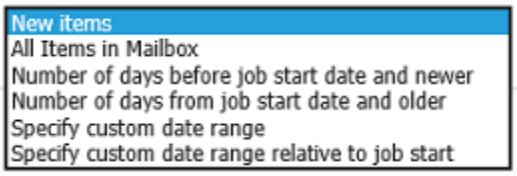
4.13.7 Scope tab

This is the most critical tab to fill out as it sets the limits on how much to archive.



Date Range to Scan

The Date Range determines which message items are collected, depending on the date of the message.



- **New Items:** All items that have not been archived by Retain since the last time the job ran.
- **All Items in Mailbox:** All items in the mailbox starting from 1/1/1970, duplicates will be processed but not stored if they already exist in the Retain archive.
- **Number of days before job start date and newer:** Only items from the relative number of days from the time the job began will be archived. E.g. messages that came into the email system 7 days ago or less.
- **Number of days from job start date and older:** Only items previous to the relative number of days from the time the job began will be archived. E.g. messages that came into the email system 7 days ago or more.
- **Specify custom date range:** Only items between two absolute dates will be dredged.
- **Specify custom date range relative to job start:** Only items between two relative dates will be dredged. E.g. messages that came into the email system between 7 and 5 days ago.

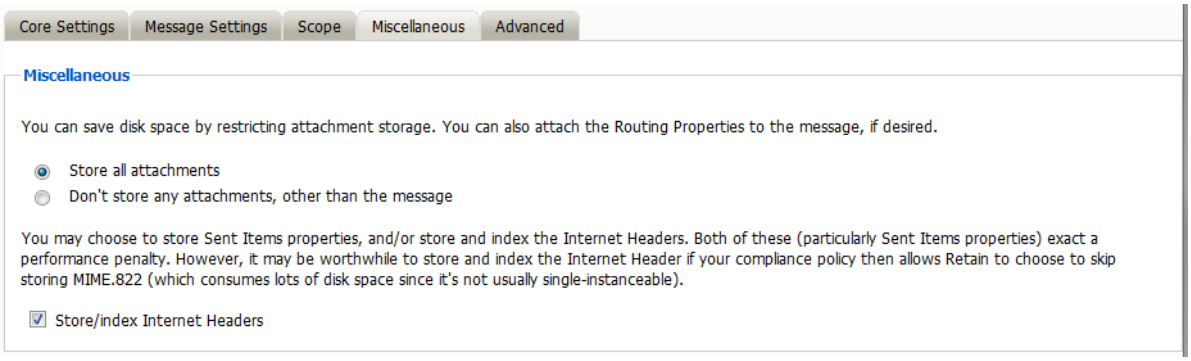
Advance Flags

Enabling "Don't Advance Timestamp" will not update the timestamp flag. Items that are dredged will still be considered new by Retain the next time the job runs.

This is useful when troubleshooting, but is generally not used for normal jobs.

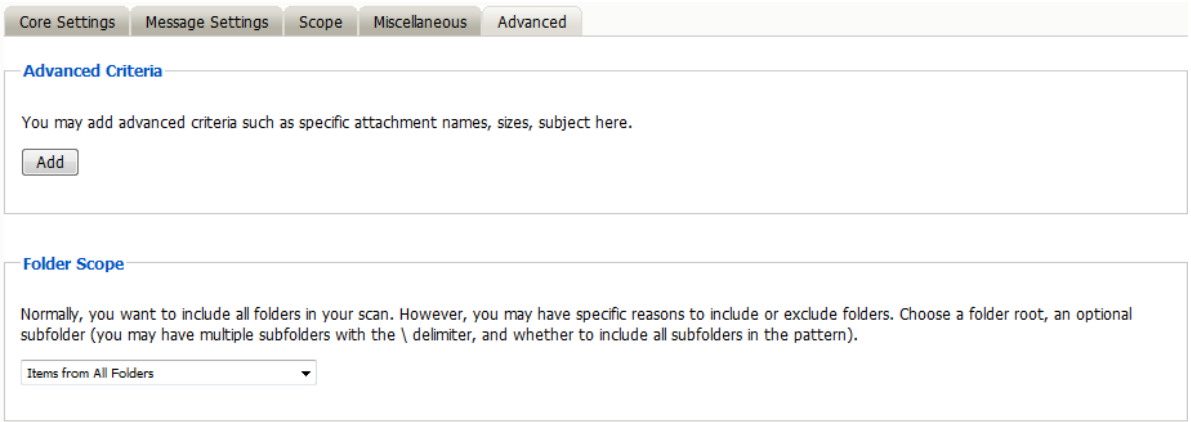
4.13.8 Miscellaneous tab

The miscellaneous tab allows the configuration to allow or deny archiving and indexing of attachments in Retain. If attachments are to be archived, they may also be indexed to provide searching capability in the browse messages interface.



4.13.9 Advanced tab

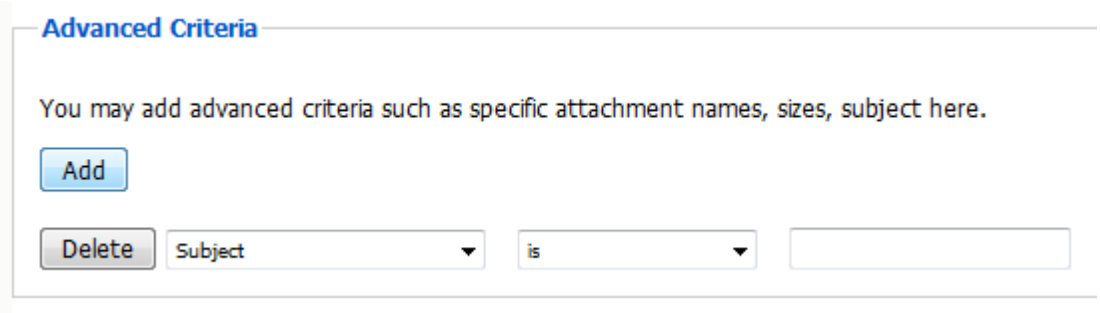
The advanced tab allows the administrator to be even more specific in what to archive and what to exclude. The criteria added under the ‘advanced’ tab will limit associated jobs to only the items matching the criteria. Up to 6 lines of advanced criteria may be added to each profile. Each additional line will be logically AND-ed together. For example, the system will archive all items where the following is true: Cri-teria A AND Criteria B AND Criteria C AND etc.



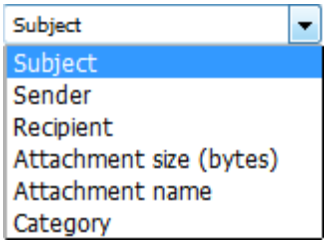
You may select based on:

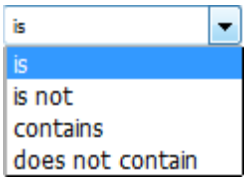
- Subject
- Sender
- Recipient
- Size

- Attachment Name



Specify the criteria, whether it is: equal to or not equal to, whether they contain or do not contain the item.





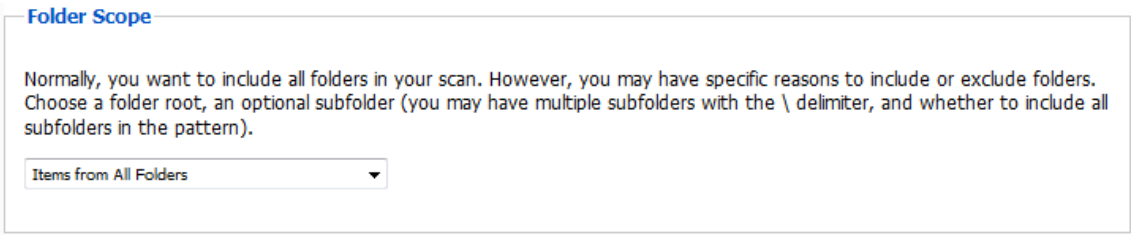
This gives you great flexibility and granularity. It allows you to customize archival jobs and retention for many different groups, or even individuals.

4.13.10 Folder Scope

By default, we dredge items from all folders. You can specify one or more inclusions or exclusions.

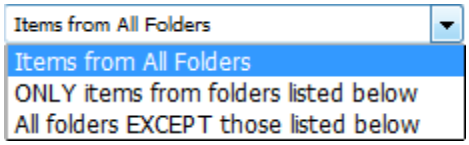
Your choices are:

- Dredge everything
- Dredge only these listed folders
- Dredge everything except these listed folders

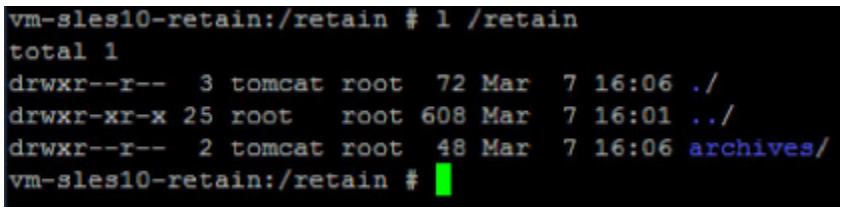


How to specify the list of folders to dredge/exclude:

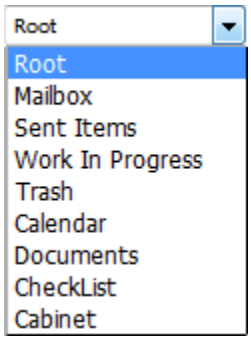
1. Select to only get items from the list, or all except items from the list, as desired



2. Select ‘Add’ to open a new selection



3. Specify a System Folder (mandatory). Example: Calendar



4. You may specify a subfolder of that folder (optional)

   Example: entering “old” would mean the folder “old” under “Calendar”

5. You can have multiple hierarchies under that with the / delimiter

   Example “old/mail” would mean the subfolder “mail” under “old” under “Calendar”

6. You specify if the option includes subfolder

   Example: If you select “old” and “includes subfolder” is unchecked, “Calendar/mail” is selected. If “includes subfolder” is checked, “Calendar/old/mail” would also be selected

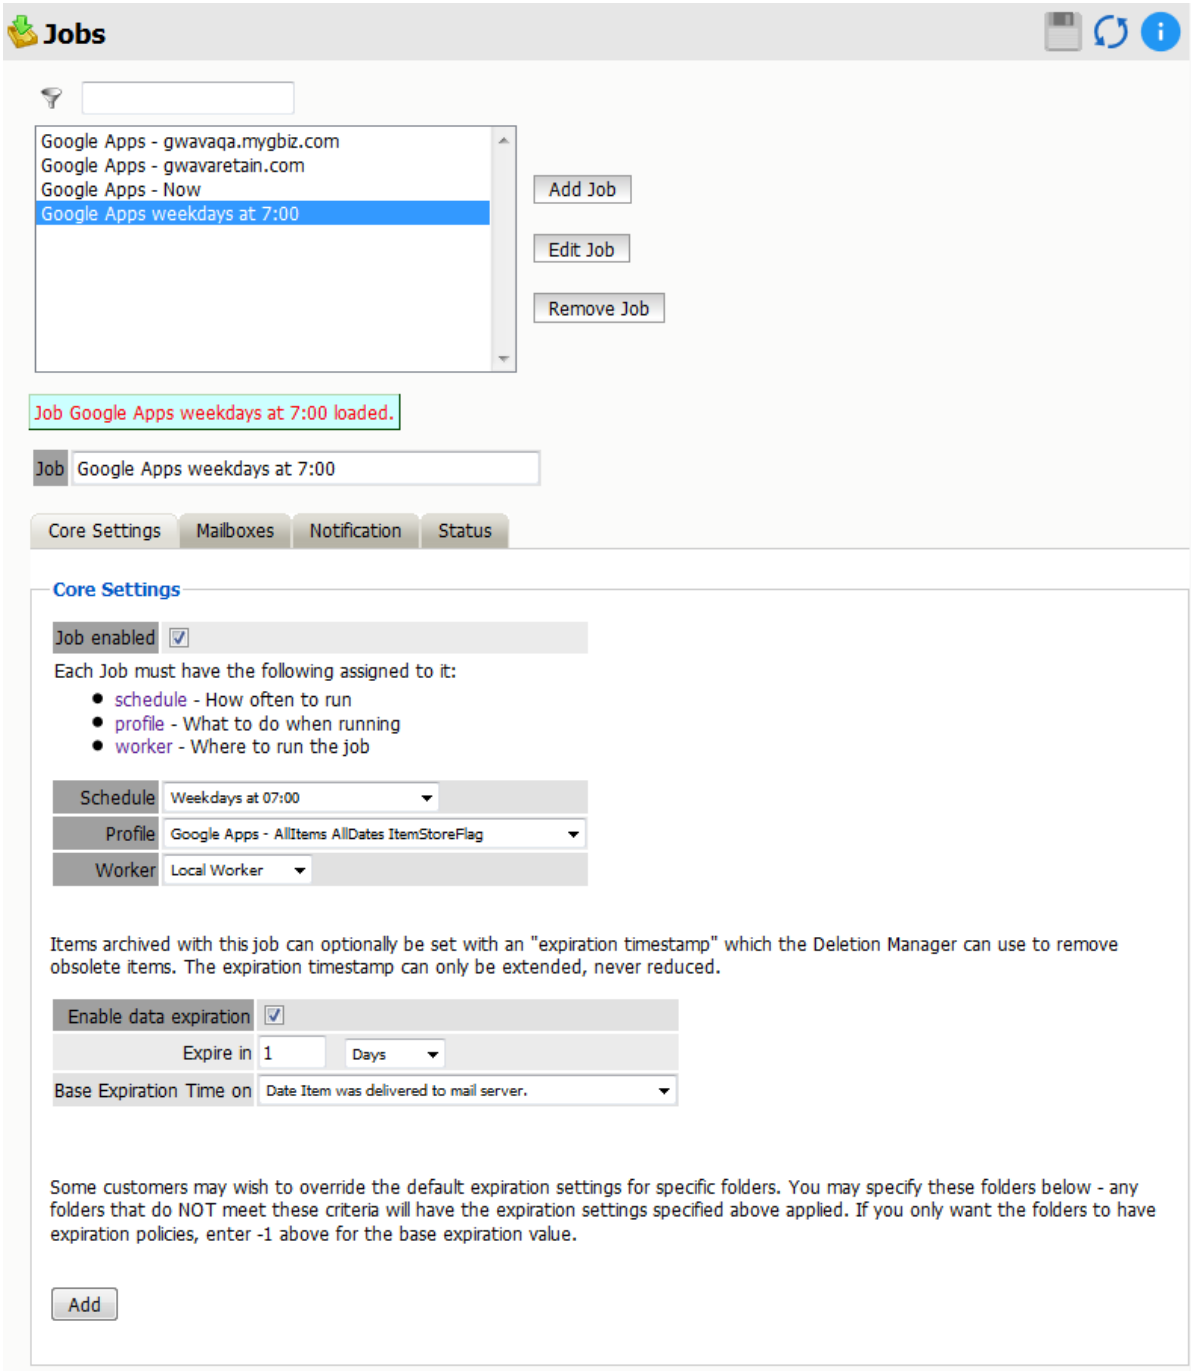
4.14 Google Apps Jobs

All data collection for Google Apps is configured through the Google Apps job interface. The jobs here combine the Profile, the Schedule, and a Worker together to archive a specified mailbox, distribution list, or domain. A Job must have a Profile, Schedule, Worker, and a target, mailbox, list, or domain, before it can be saved or run.

The Schedule, Profile, and Worker selected here will determine what is archived, when it is archived, and what worker does the actual archival work. The Job MUST be enabled to run. Only previously configured Schedules, Profiles, and Workers can be selected.

4.14.1 Core Settings tab

Core settings allows you to enable the job and set the Schedule, Profile and Worker needed to run the job.



The Data Expiration is an option to place a timestamp on the mail in the Retain database, which allows for ease of automation for the deletion manager. In addition, devices such as NetApp, Centera, and Hitachi HCAP may use this number to enforce hardware level protection of the stored item so that no one (including Retain) may delete the item before its expiration date.

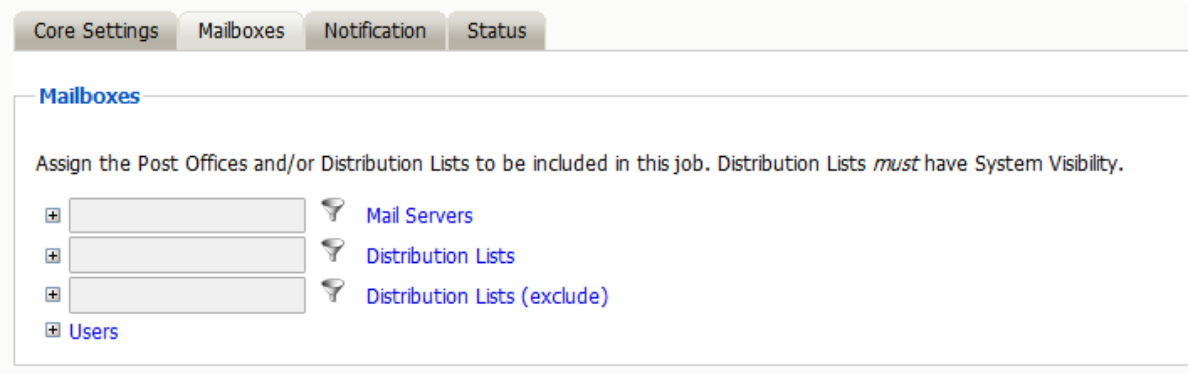
Job Expiration is not retroactive for mail in the database, and only applies to mail archived by the job that it is active for. The base folders and criteria are specified under the custom expiration dates accessed through the ‘Add’ button at the bottom. All messages included in any specified folders will have a differ-ent date or be exempted from the standard expiration date. To have messages with custom job or folder expiration dates properly expire, the deletion management date scope must be set to delete messages with an Expiration Date older than 1 day.

4.14.2 Mailboxes tab

The mailboxes tab is where the administrator specifies which entities (mail server(s) and/or Distribution List(s)) are to be scanned.

Expand the Post Office and/or Distribution List trees, and check off the items you want to be dredged.

NOTE: If you desire to have a job backup a single user, or selected group of users, select the Users menu and assign the user(s) desired.



The users section allows you to select individual users to include, or exclude them from an archive job. For example: you can select an entire Mail Server to be archived, and then expand the users section to include or exclude users to the job.

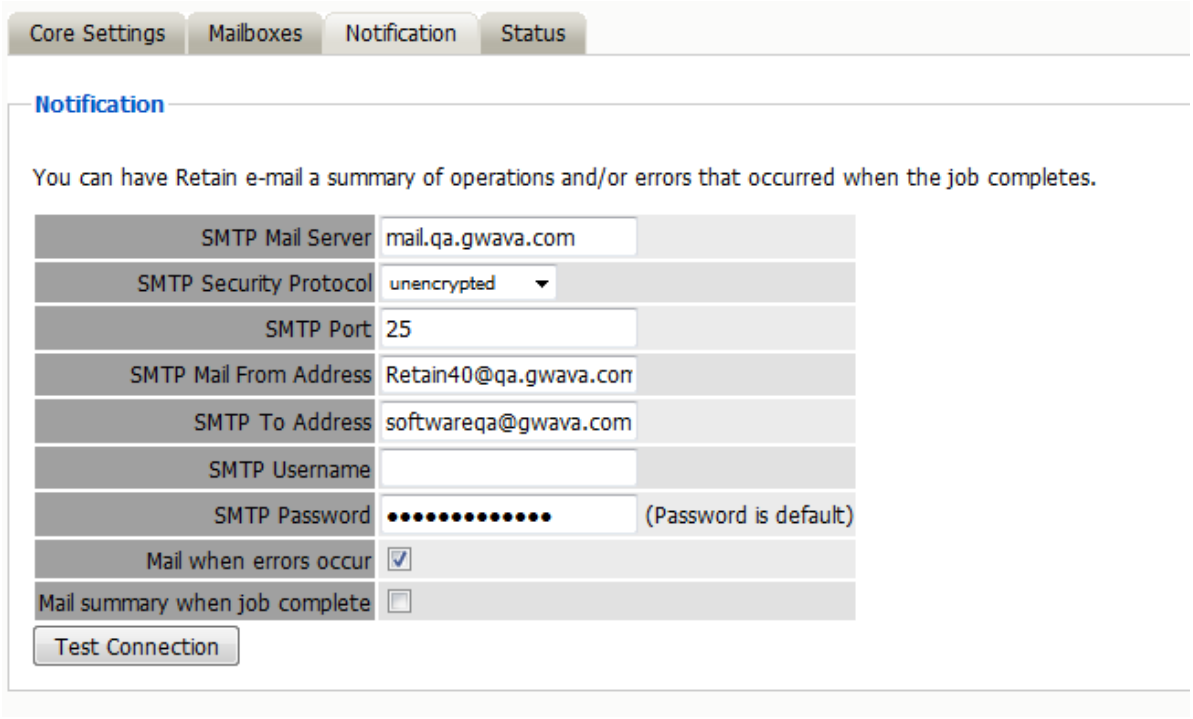
This can also be used to select only certain users in the system for an archive job.

To add a user to the Include or Exclude list, select the respective ‘Add user’ button and search for the user. It can be helpful to unselect the ‘only show recently cached items’ option.

Add the selected users to the list in the search window, then select ‘Ok’ to add them to the include or exclude list.

4.14.3 Notification tab

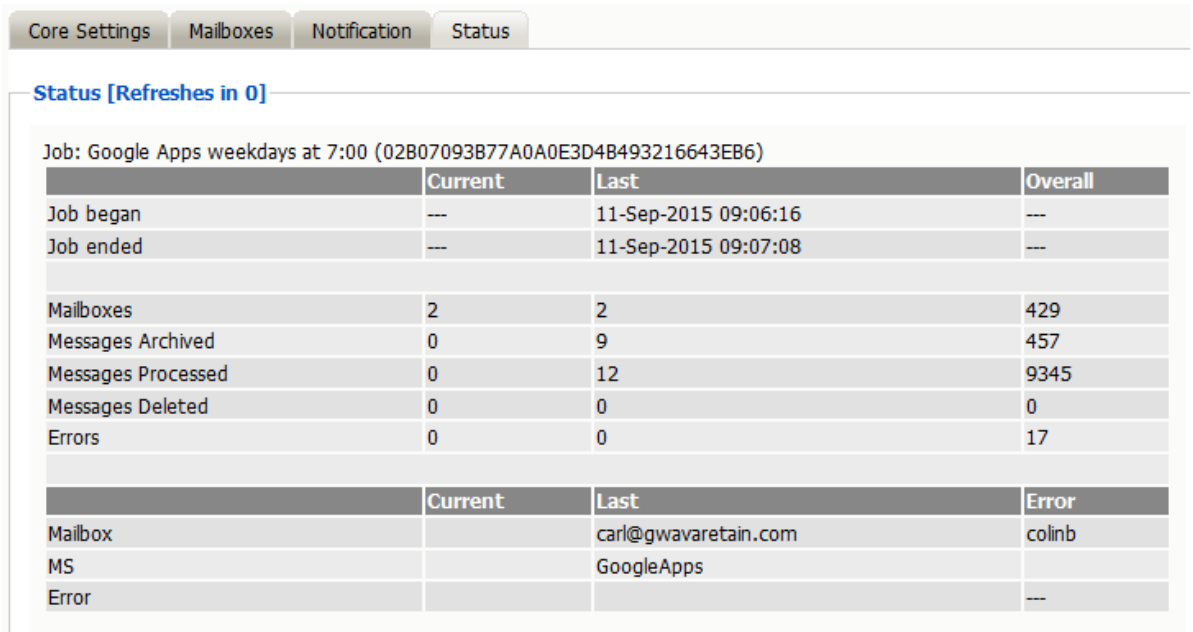
When a job is run, the notification option allows the administrator to be emailed a summary of each run-ning job if desired.



For notification to function correctly, the SMTP information for the desired SMTP server must be fully filled-out. How much information is required varied depending on the mail system used.

4.14.4 Status tab

The Status tab displays the status of any currently running jobs, as well as the stats of the last completed job.



This tab is informational only unless a job is currently running. If the selected job is running, an option to abort the job is displayed.

This window will refresh every ten seconds to keep you up to date as to the status of the selected job. The completed line is displayed during an active job. The completed status is a display of how many mail-boxes have been completed, the job mailbox total, and gives an incrementing percentage for the amount completed. This amount is based entirely on the number of mailboxes, not the amount of mail. Because the last mailbox could be larger than the rest of the system, this percentage may not be accurate accord-ing to time.

You may now configure Schedules, Workers and Jobs.

4.15 Bloomberg Module

The Bloomberg module allows the collection of Bloomberg IM and Mail messages from the daily FTP downloads allowed from the Bloomberg System.

As electronic business communications, Bloomberg Mail is subject to the same compliance regulations as other electronically stored information. Potential litigation makes it critical for companies to archive this data and make it available for discovery. The U.S. Securities and Exchange Commission (SEC) also mandates retention.

4.15.1 Prerequisites

On the Bloomberg server create a user that Retain can use to log into the Bloomberg SFTP server.

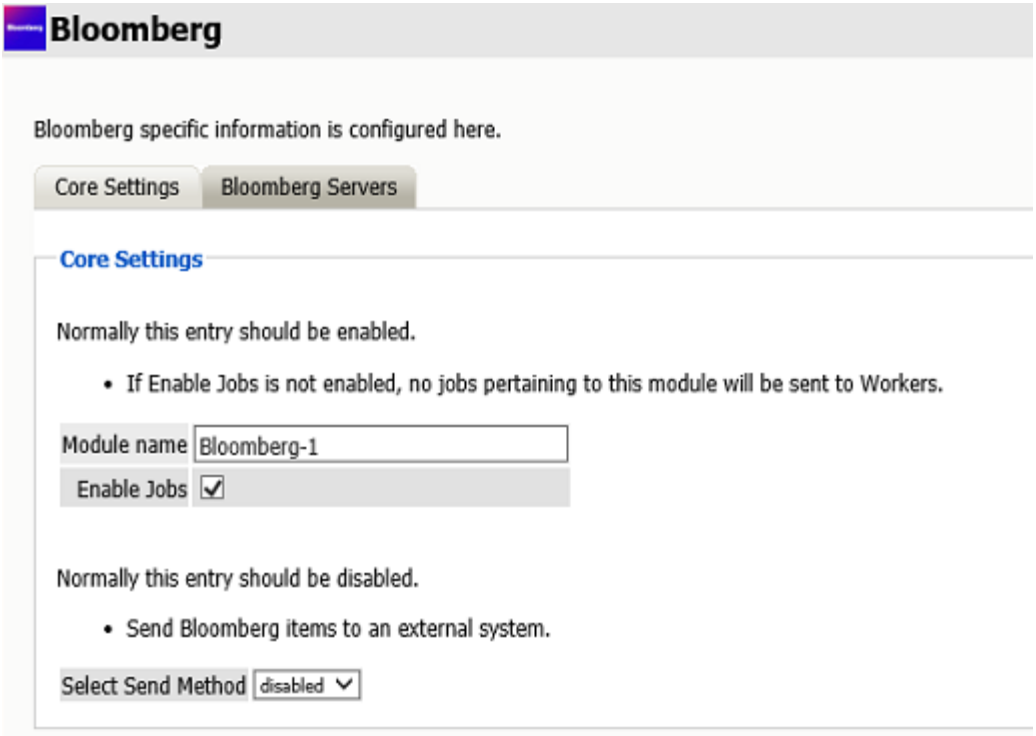
4.16 Module Configuration

4.16.1 Core Settings tab

The module needs to be enabled on this page to make it active in the Retain system.

The module can be given a name.

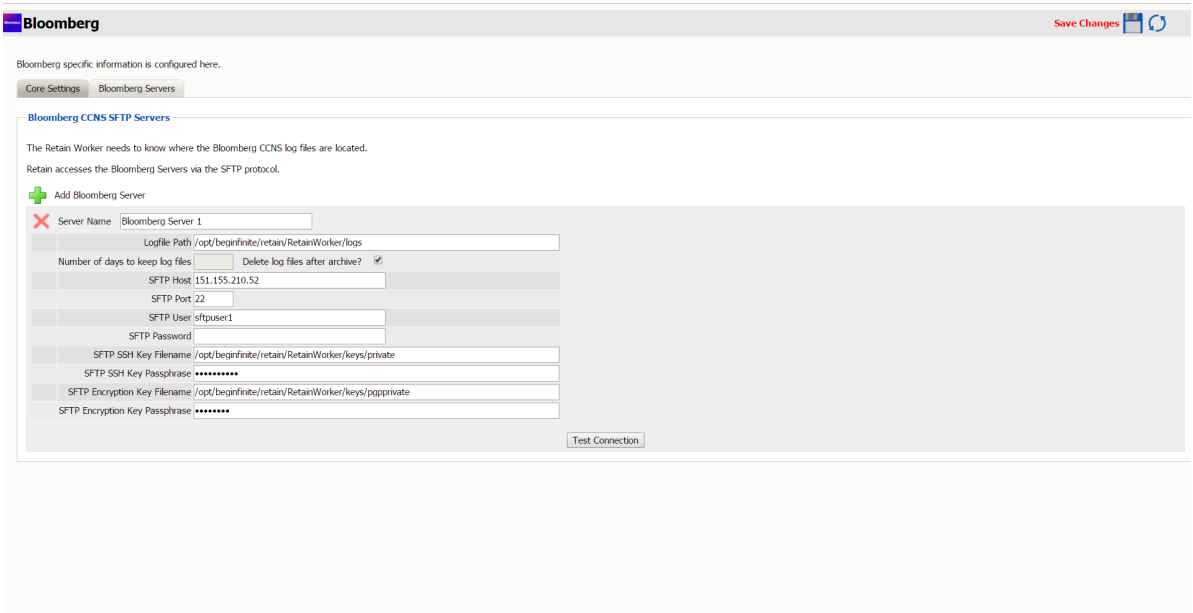
The Send Method option enables either the SMTP Forwarding or FTP features. For either feature to appear and function, the Module Forwarding tab must be configured on the Server Configuration page. See that section for more information.



Select Send Method: Send Bloomberg items to an external system via FTP or SMTP. Normally this entry should be disabled.

4.16.2 Bloomberg Servers tab

Add Bloomberg Server: Click on the green plus sign to add a Bloomberg server to download from. Multiple servers can be added.



- **Server Name:** The name Retain will use to reference Bloomberg messages. You can use any name you wish.

- **Logfile Path:** Where the worker will transfer the logfiles to locally before processing. The log file path must be unique for each server.
- **Delete log files after archive?: Numbers of days to keep Logfiles:** By default, Retain processes and archives the log files, keeping the log files in the Logfile Path directory indefinitely. Disabling this option allows you to set the *Numbers of days to keep Logfiles* option.
- **Number of days to keep Logfiles:** How many days to keep the files downloaded from Bloomberg before Retain deletes the logfiles from the log file directory. 0 means never delete the files.
- **SFTP Host:** The IP address or hostname of the Bloomberg server.
- **SFTP Port:** Port 22, by default, but can be changed by the admin.
- **SFTP User:** This user is created on the Bloomberg server by the Bloomberg server administrator. The user can authenticate either through password or SSH using a private key.
- **SFTP Password:** This password is created on the Bloomberg server.
- **SFTP SSH Key Filename:** The local path to the SSH key file downloaded from the Bloomberg Server.
- **SFTP SSH Key Passphrase:** Created on the Bloomberg server.
- **SFTP Encryption Key Filename:** The local path to the Encryption key file downloaded from the Bloomberg Server.
- **SFTP Encryption Key Passphrase:** Created on the Bloomberg server.

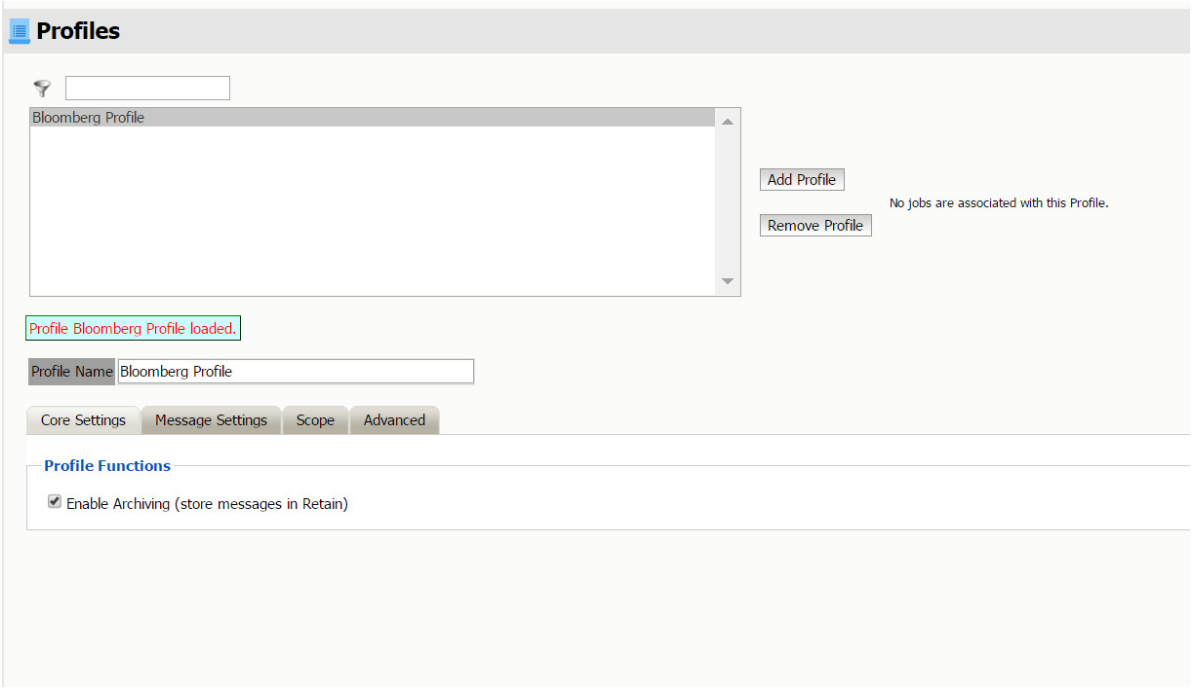
4.16.3 Bloomberg Profile

The Profile allows you to set what is to be dredged from your Bloomberg system.

Click on “Add Profile” and provide a profile name, or select an already existing profile to access the con-figuration tabs. All changes made on this page must be saved by selecting the “save changes”, disk icon, at the top right of the page. Tabs may be changed and navigated through without affecting new settings, but any move to another page will require saving, or abandoning the changes made.

Core Settings tab

Core settings allow you to enable archiving.

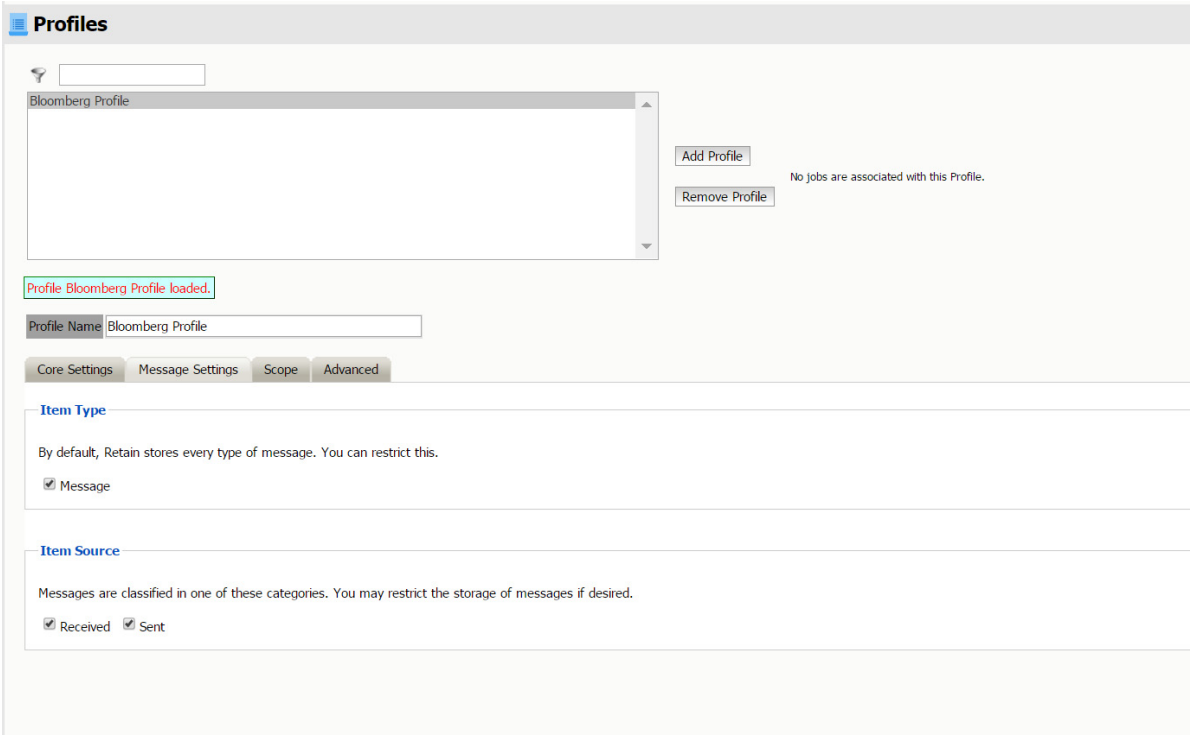


4.16.4 Profile Functions

The Profile Functions tell the Retain Server what to do with the mail it archives from the Bloomberg sys-tem. If Archiving is not enabled, mail will not be archived by Retain.

Message Settings tab

The message settings tab allows you to restrict the types of messages to be archived.

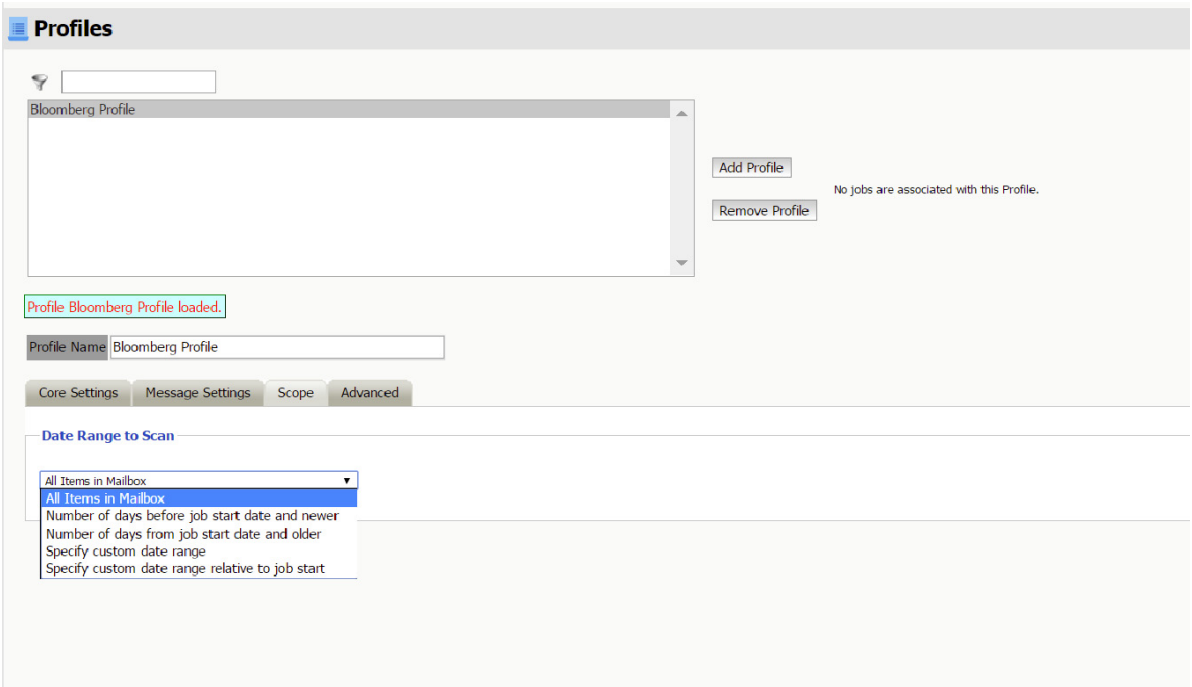


There is only one type of message, disabling this will prevent archiving. You can restrict the item source to received or sent items only or both.

By default, all items are selected for collection.

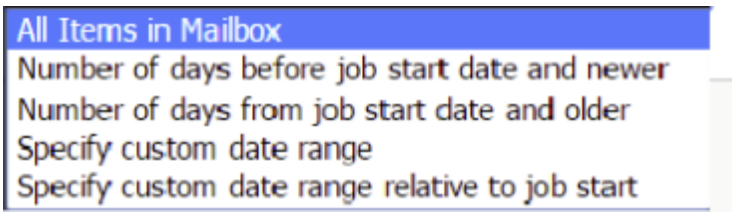
4.16.5 Scope tab

This is the most critical tab to fill out as it sets the limits on how much to archive.



4.16.6 Date Range to Scan

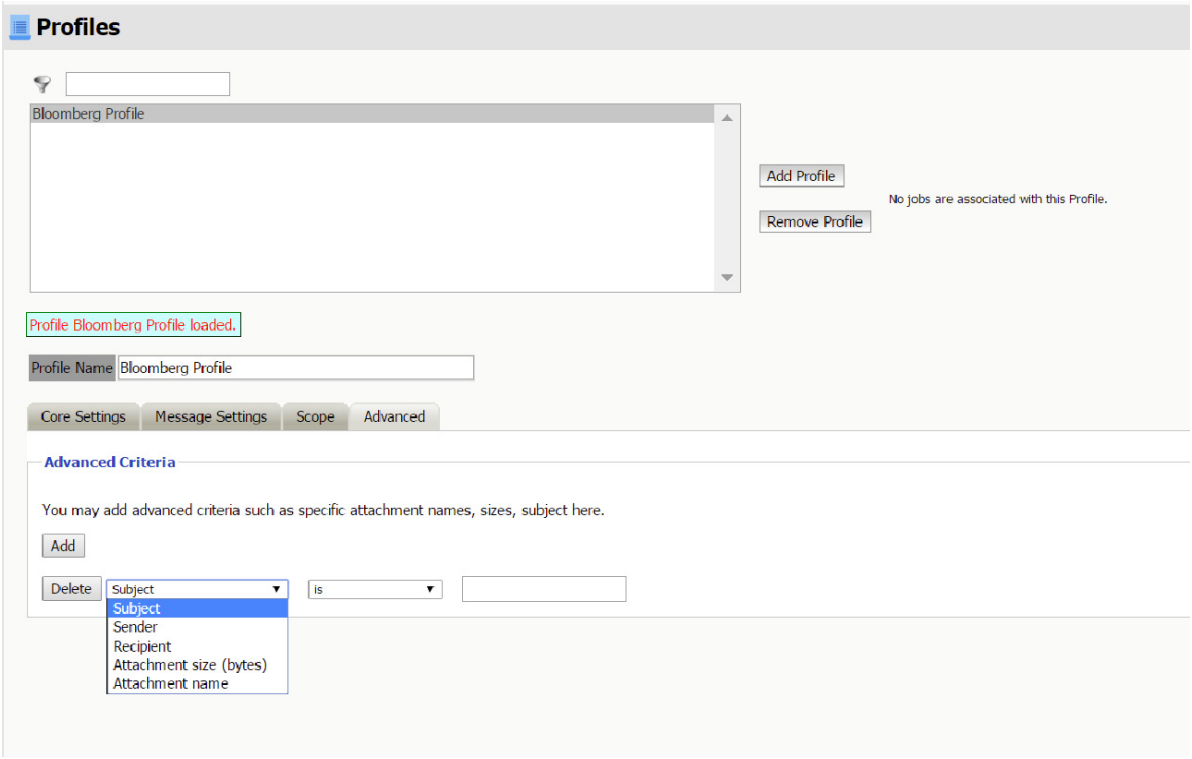
The Date Range determines which message items are collected, depending on the date of the message.



- **All Items in Mailbox:** All items in the mailbox starting from 1/1/1970, duplicates will be processed but not stored if they already exist in the Retain archive.
- **Number of day before job start date and newer:** Only items from the relative number of days from the time the job began will be archived. E.g. messages that came into the email system 7 days ago or less.
- **Number of days from job start date and older:** Only items previous to the relative number of days from the time the job began will be archived. E.g. messages that came into the email system 7 days ago or more.
- **Specify custom date range:** Only items between two absolute dates will be dredged.
- **Specify custom date range relative to job start:** Only items between two relative dates will be dredged. E.g. messages that came into the email system between 7 and 5 days ago.

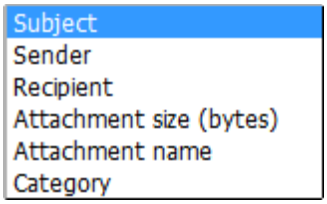
4.16.7 Advanced

The Advanced tab allows you to limit what is stored by Retain. This must be used with caution as this opens holes for data to be lost through. It is recommended to store everything since storage space is inexpensive.



Advanced Criteria

If you want to be more specific as to what to dredge or not to dredge, add the criteria here. Each line will be logically AND-ed together. Think “Dredge all items where the following is true:” Criteria A AND Criteria B AND Criteria C AND etc.



You may select based on:

- Subject
- Sender
- Recipient
- Attachment Size (in bytes)
- Attachment Name
- Category

And whether they are equal to, not equal to, contain or do not contain the item you specify.

This gives you great flexibility and granularity. It allows you to customize dredges and retention for many different groups, or even individuals.

You may now configure Schedules, Workers and Jobs.

4.17 Mobile Module

Mobile Module

The Retain for Mobile module allows the capture and archiving of MMS, SMS, and phone call information from Android smart phones. The mobile application, Retain Service, supports Android 4.4+ and utilizes the REST API in Retain. The Retain Service APK can be accessed on the 'tools' page. The Tools page is accessed by selecting 'Tools' from the top right of the Retain Administration interface.

To enable mobile data archiving the following tasks must be completed, in order:

For the newer application supporting Android 4.4+

1. The Router must be installed and be available outside the host network. (in DMZ) Server or Router must have a 3rd party signed TLS certificate.
2. The Mobile Module must be configured.
3. The Retain Android application (4.4+), must be installed to users' phones.
4. Devices must be registered under 'Device Management'.
5. Users must use the registration email to verify device registration with the Retain server.
6. The Storage path folder(s) must be manually created.

NOTE: Jobs and Workers are not required with Retain Service, they are only present here to provide support for the legacy Android application. Retain Service utilizes the REST API, and will function as soon as the above steps have been completed. No additional work in necessary.

The Retain Router should be installed in the DMZ area of the host network. The Mobile Router or the Retain Server, whichever is to be used, requires a TLS certificate signed by a 3rd party for trusted communications. The Router must be available to the Internet and the Retain Server. The Router and Worker may be installed on the same machine, or apart.

It is recom-mended to use a Mobile Device Management (MDM) system to distribute Retain Service to devices. When using an MDM to distribute the Retain app, the system can be set to not allow app removal. The MDM will need the package name for Retain Service, which should be set to *com.gwava.retain.mobile*. Consult the MDM's documentation for instructions.

MDMs that have been tested with Retain include:

- ACE
- Airwatch
- BES
- Citrix XenMobile
- Good Dynamics
- Google XenMobile
- MobileIron

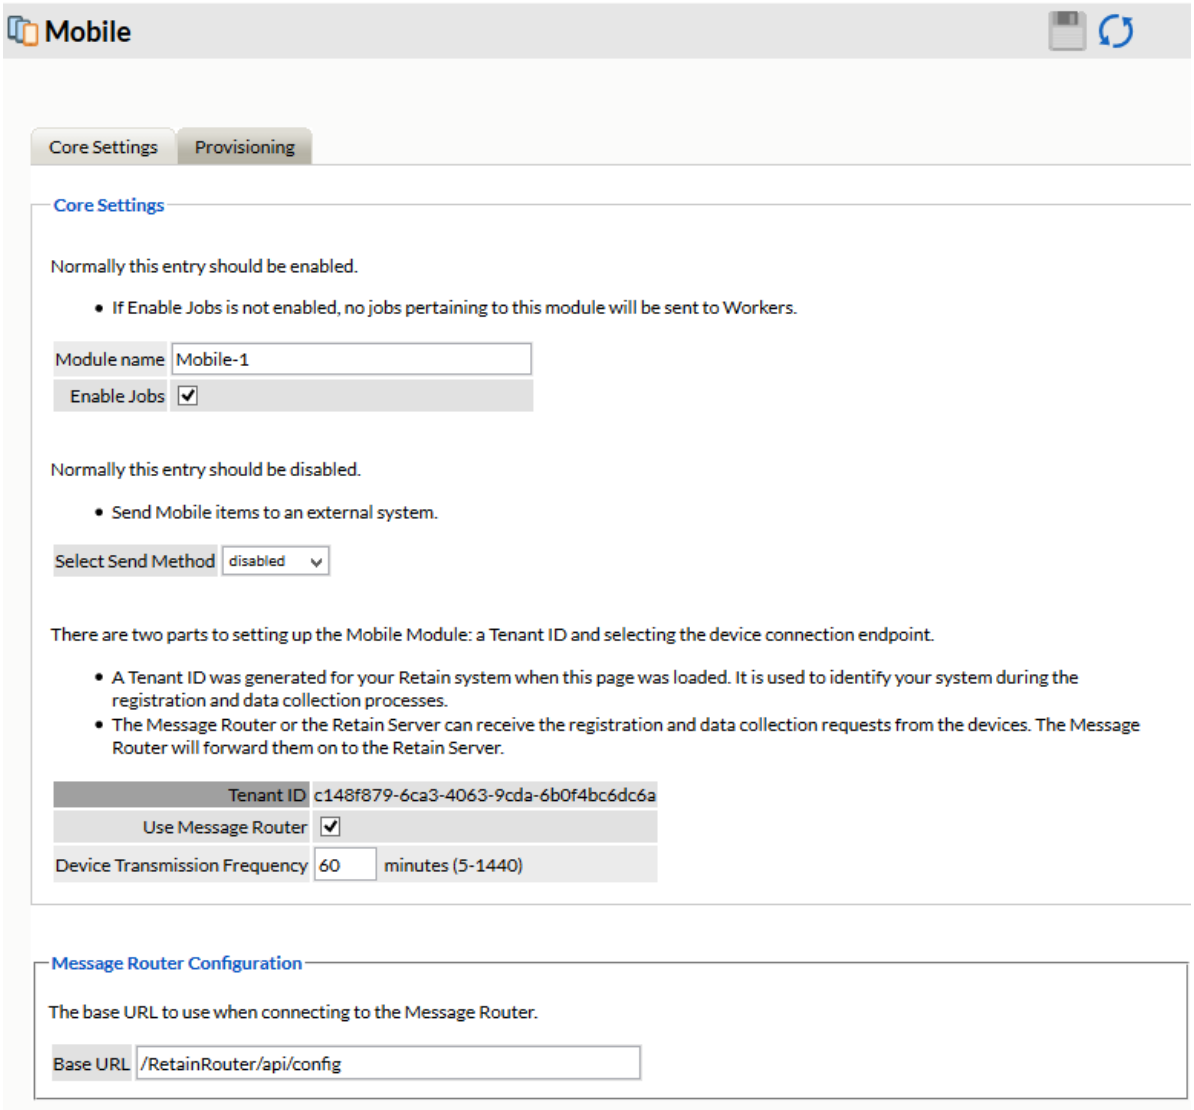
The application sends collected information, approximately every hour, to the Retain Router. No data will be archived until the application is registered and configured.

Core Settings

The module needs to be enabled on this page to make it active in the Retain system.

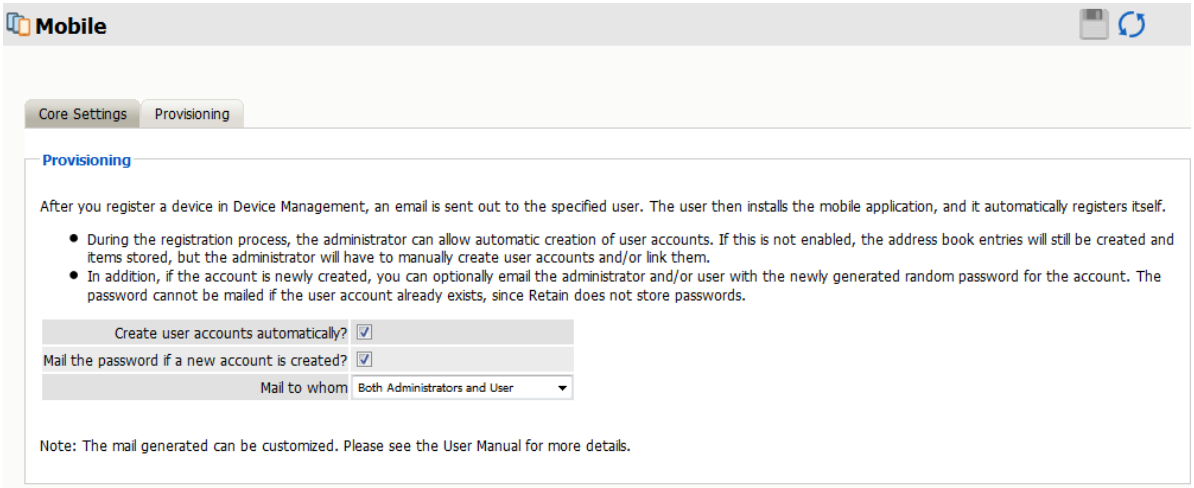
The module can be given a name.

The Send Method option enables either the SMTP Forwarding or FTP features. For either feature to appear and function, the Module Forwarding tab must be configured on the Server Configuration page. See that section for more information.



Provisioning

Provisioning options allow for automatic user accounts, and associated passwords mailed to the desig-nated address.



After the module configuration has been saved, the Mobile, Profile, and Device Management menu pages will become be available.

Worker

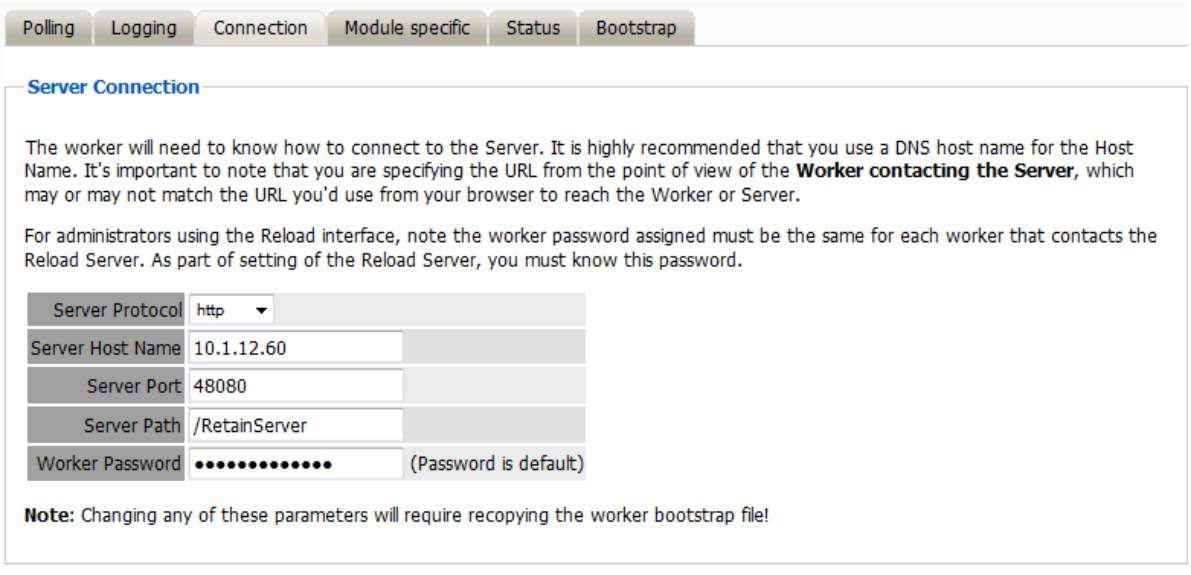
The Worker is only required for legacy Android application support and is not necessary with the current application. The following information is only necessary for supporting older application systems.

The Mobile worker is a special worker configuration which contacts the Retain Router with all the information it needs to function. This includes the configuration settings for the data path and connection information. In general, the Mobile Worker is configured the exact same as most workers in the Retain system, however there is one setting for the Mobile Worker which is specific to the Mobile Module; the data path. The data path is essential for the worker to function with the Mobile Module. An existing worker may be used with the Mobile Module, however, it must have the bootstrap file re-uploaded once the data path has been configured.

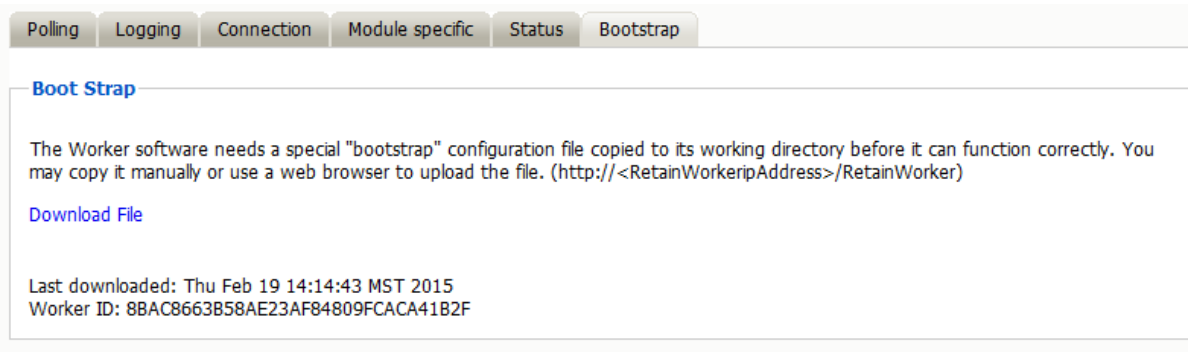
The mobile data is delivered to the Retain Server through a router service, which must be informed how to contact the Retain Server. While the Router is gathering information between jobs, the mobile data is stored. The stored location is specified by the data path. This is why the Router and the worker must have direct and constant access to the data path; the data path stores the message data collected by the Router and retrieved by the Worker.

To ensure connection to both the data path and the Retain Server are correct, the Connection and the Module Specific tabs must be reviewed and configured. The Connection tab holds the connection address which the Retain Worker will use to contact the Retain Server. Depending on where the Server and the Worker are located in the network, and because the Worker may be installed alongside the Router in the DMZ, the connection information must be accurate to allow a clean connection to the Retain Server. An IP Address will work, but the DNS hostname of the Retain Server is also recognized and supported.

The port and address must be open or forwarded through security and firewalls from the Worker to the Server. In addition, the Retain Server will not accept connections without the password set here. The password is randomized and doesn’t need to be changed.

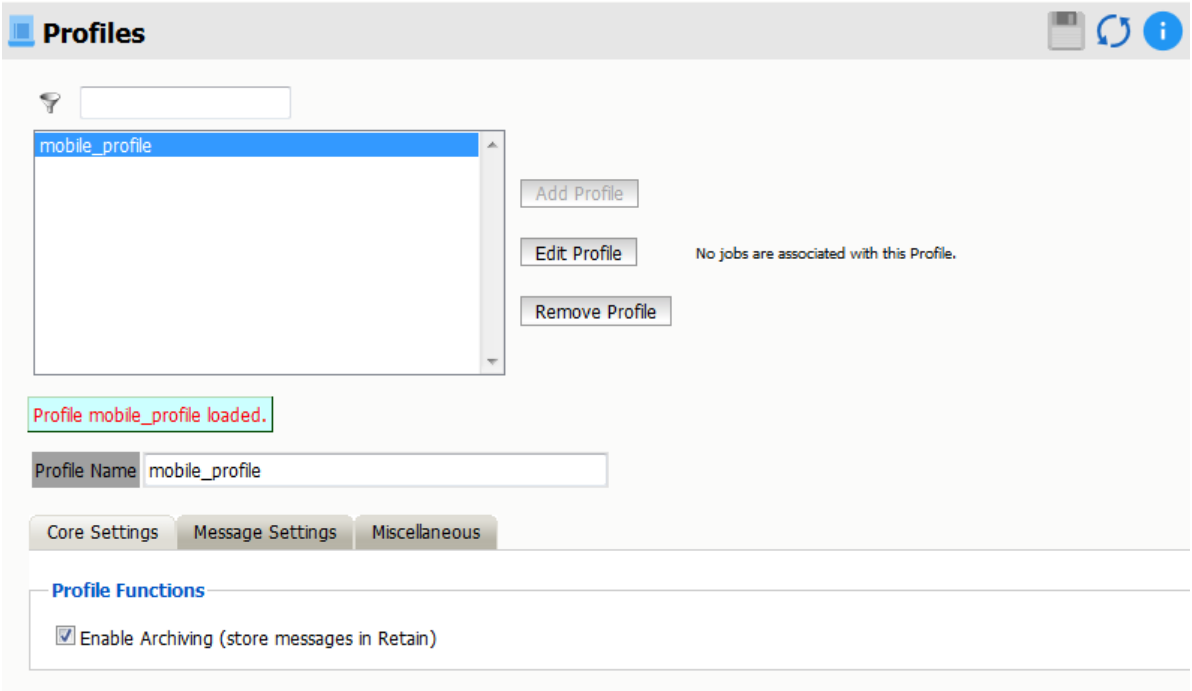


Once the Mobile Worker has been configured, save the changes. Once the settings have been saved, the worker will be created, and the configuration saved to the bootstrap file. The bootstrap file must be uploaded to the Retain Worker before the Worker will function. Select the 'Download File' link to save the configuration file. Browse to the worker's page, and upload the configuration file. Once uploaded, the worker automatically reads the configuration and checks with the Retain Server for jobs.



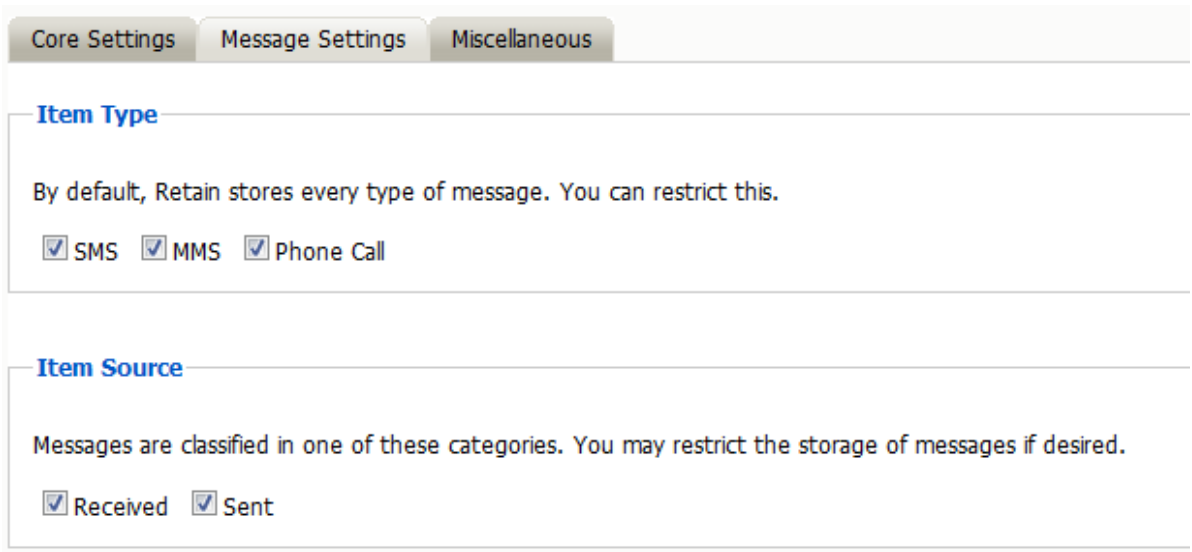
Profile

The Mobile Profile allows the administrator to dictate what types of message data is collected from registered mobile devices. The profile is universal, and once configured works for the entire mobile system.

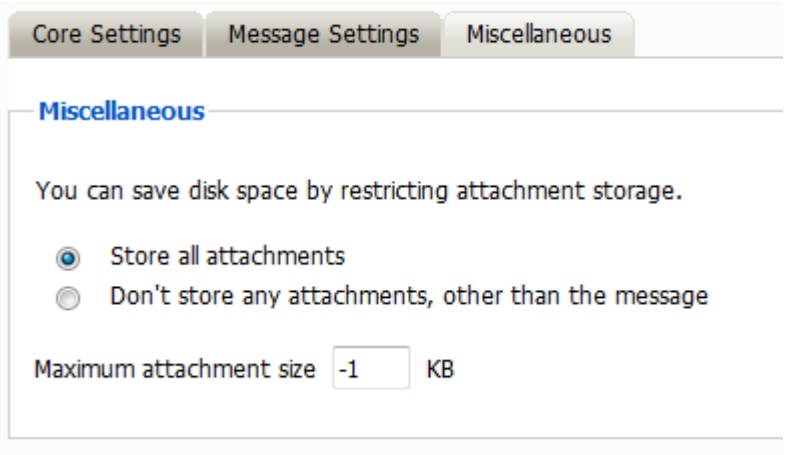


The Mobile Profile is fairly straight-forward. To enable message data collection for the mobile system, the profile must be activated.

Under Message Settings, the different types of messages which can be archived from Mobile devices are listed and configured. The message source and type are available for selection or exclusion.



Finally, the Miscellaneous tab allows for configuring attachments. Attachment size and general attachment settings can be configured here. If there is a maximum limit, set it in the field below.



The attachment size limit is listed in KB. If the size limit is configured to -1, then there is no limit and all attachments will be archived no matter the size.

There can only be one profile for the Mobile module. Once a profile has been created, the option to create a new profile will be disabled and grayed-out. If a new profile is desired, the existing profile must first be deleted. If there is no profile active in the system, the option to create a new profile will be active.

Once the module, worker, and profile have been configured, devices must be added through Device Management.

Device Management

Devices are added to and deleted from Retain through the Device Management interface. Device Management is split into two sections: Mobile and Blackberry. In both sections, multiple devices may be selected and acted on. In addition, all selected devices remain selected until either an action is taken or the administrator has navigated away from the page.

Mobile

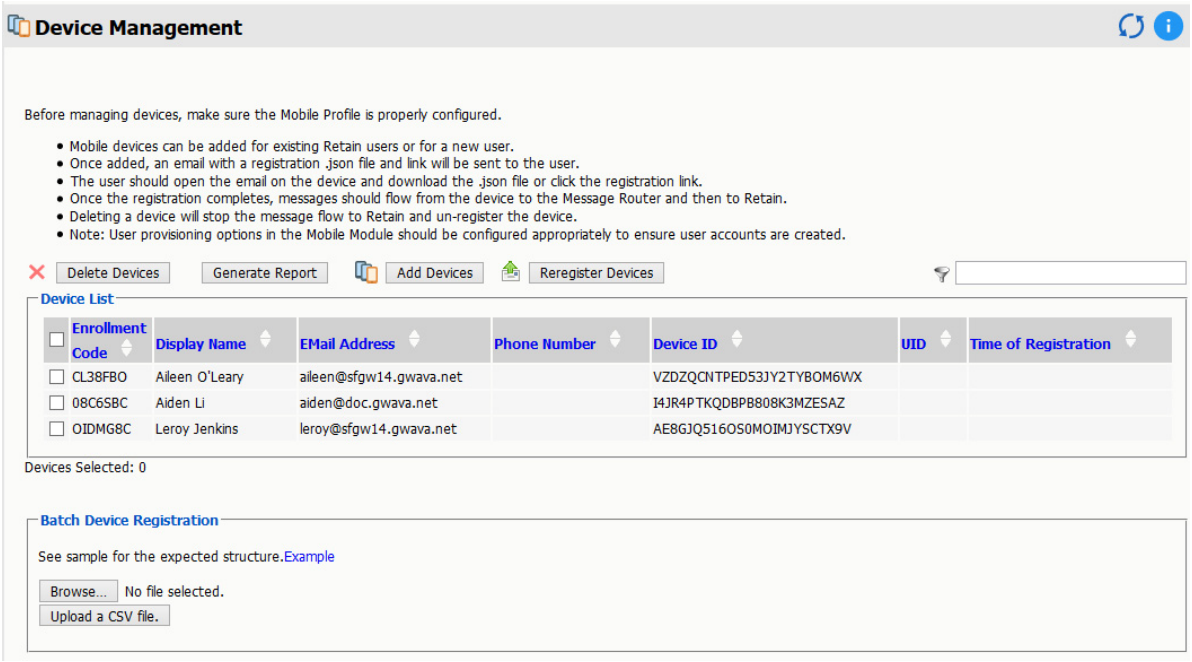
Once a device is added to the interface, Retain creates an email which contains an activating link. Users need to open this email and follow the link on their mobile device. The link should be activated through the Retain application on the Android device. The activation email can be customized. The template this email is generated from is located in the file:

…\RetainServer\WEB-INF\classes\config\mobile\registeruser.html

The email sent to users may be customized, including any specific instructions for the organization, can be added or modified here, and all emails sent to users will contain information placed in this template. Information on the email variables is listed at the end of this page.

When a device is added, an email is sent to the user with a unique enrollment code. Devices which have been activated will show a Device ID and a UID. Once registered, the Retain app will start collecting data setup in the profile section, (SMS, MMS, and phone call), and will send it to the Mobile Router to be archived in the Retain System. NOTE: While it is possible to register devices without utilizing a registration email, the process requires the manual creation of a registration file for each device. If required, call support for details.

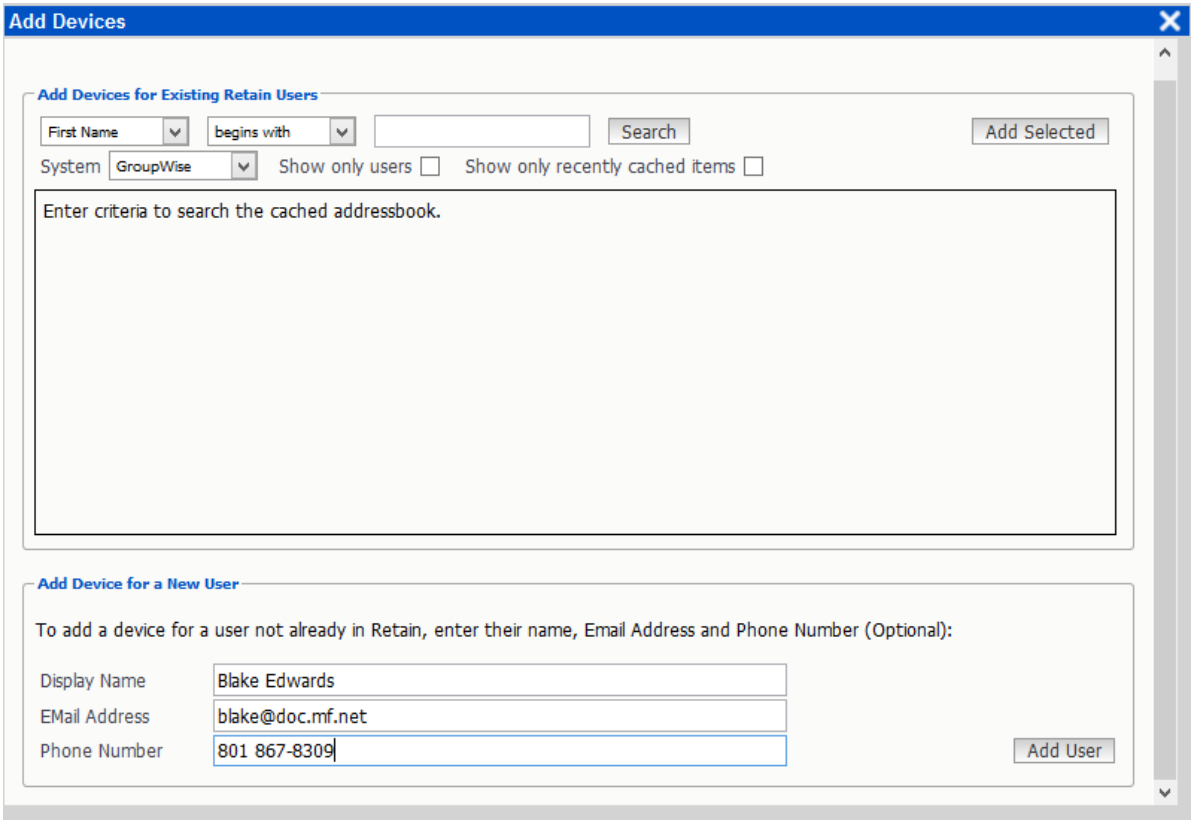
Each column is sortable press the arrows to sort by that column. The search box allows you to filter the Device List.



To add a device to the Retain system, open the Device Management window and select the *Add Devices* button.

The Add Devices window opens, which allows administrators to select and add multiple users. The system may also add users which are not cached in the Retain system, but they must be added one at a time. Manually input the name, associated email address, and phone number at the bottom and select the 'add user' button.

NOTE: Phone numbers can only be added with "Add Device for A New User" or adding the phone number to the CSV file for Batch Device Registration. There is *no option* to edit the user to add the phone number when adding the user by searching "Add Devices for Existing Retain Users"



Devices are only enrolled once. If a user acquires a new phone, or if they change phones, a new registration needs to be completed to maintain the correct codes.

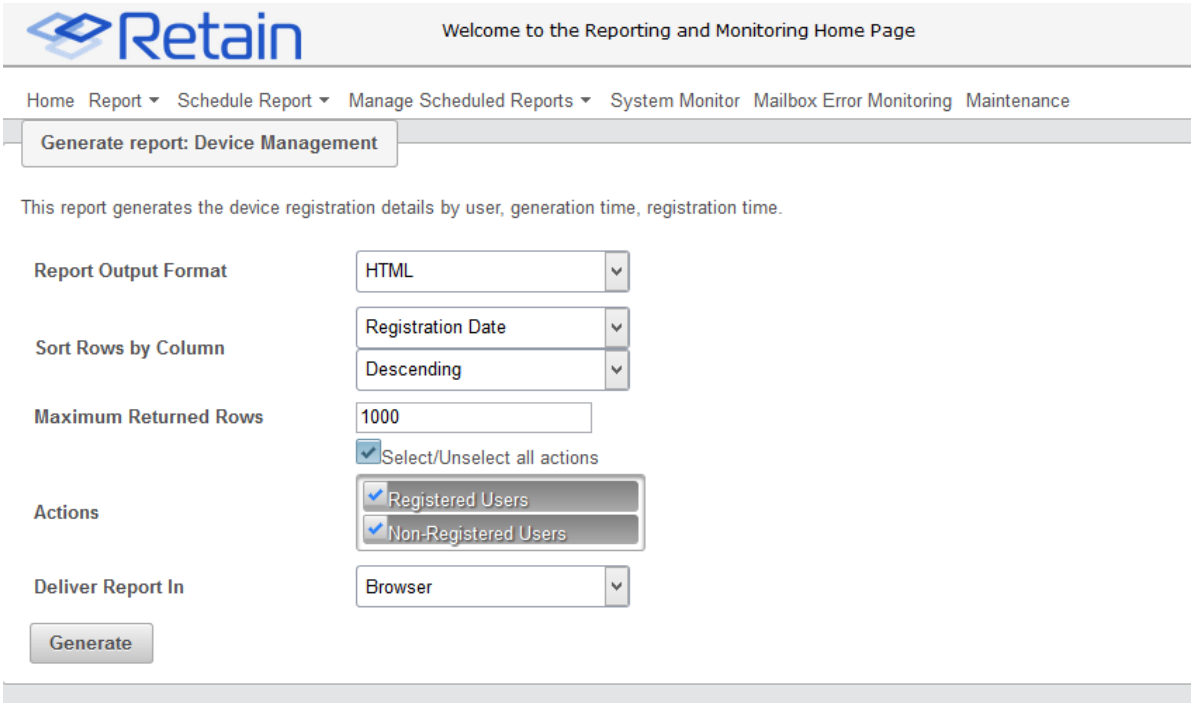
To re-register a device, select the desired device and press the *Reregister Devices* button. Once pressed, Retain will recreate the enrollment, remove the old Device ID, and resend the registration email. The end user will need to repeat the enrollment process stated in the email to register the new device.

A device without a Time of Registration is not registered. A Not Registered device has had the registration process begun but the user has not completed the registration process. No archiving is being done until the process is complete.

Once devices have been added and registered, they are accessible to be modified, removed from, and archived by the Retain system.

The *Delete Device* button will remove the device from device management but will not remove the data.

The *Generate Report* button will open a new Reporting and Monitoring: Generate report window, where a report can be created.



Batch Device Registration

A CSV file can be created with a list of devices to allow for the batch creation of many devices.

The format of the CSV file can be found by downloading the sample file. Generally of the form:

Jane Doe,Jane@Gmail.com

Bob Wood,Bobby@Microfocus.com

Joe Cole,Joe@gmail.com

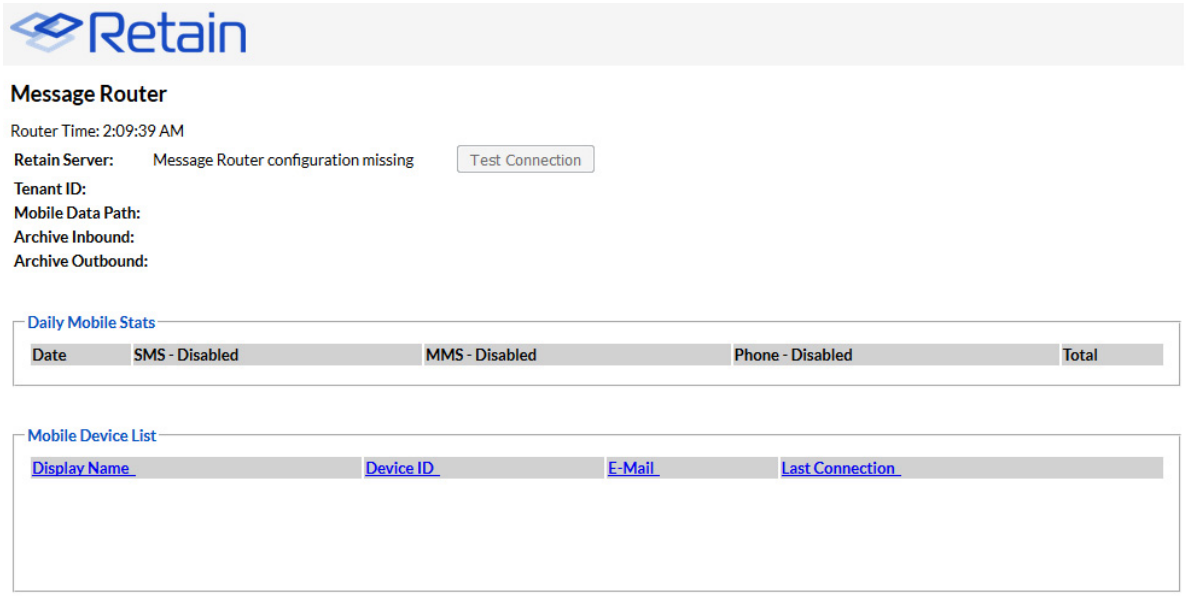
Mobile Router Page

The Retain Router has its own web page. This page can be used to test the connection to the Retain Server, as well as verify that communication is open and that devices registered in the Retain Server are listed in the Retain Router.

To access the Router page, open a browser and enter the connection URL.

http://Retain_Router_IP_or_Domain/RetainRouter

For example: http://192.168.1.21/RetainRouter



Devices which are active are listed along with daily statistics. The daily statistics are reset every night.

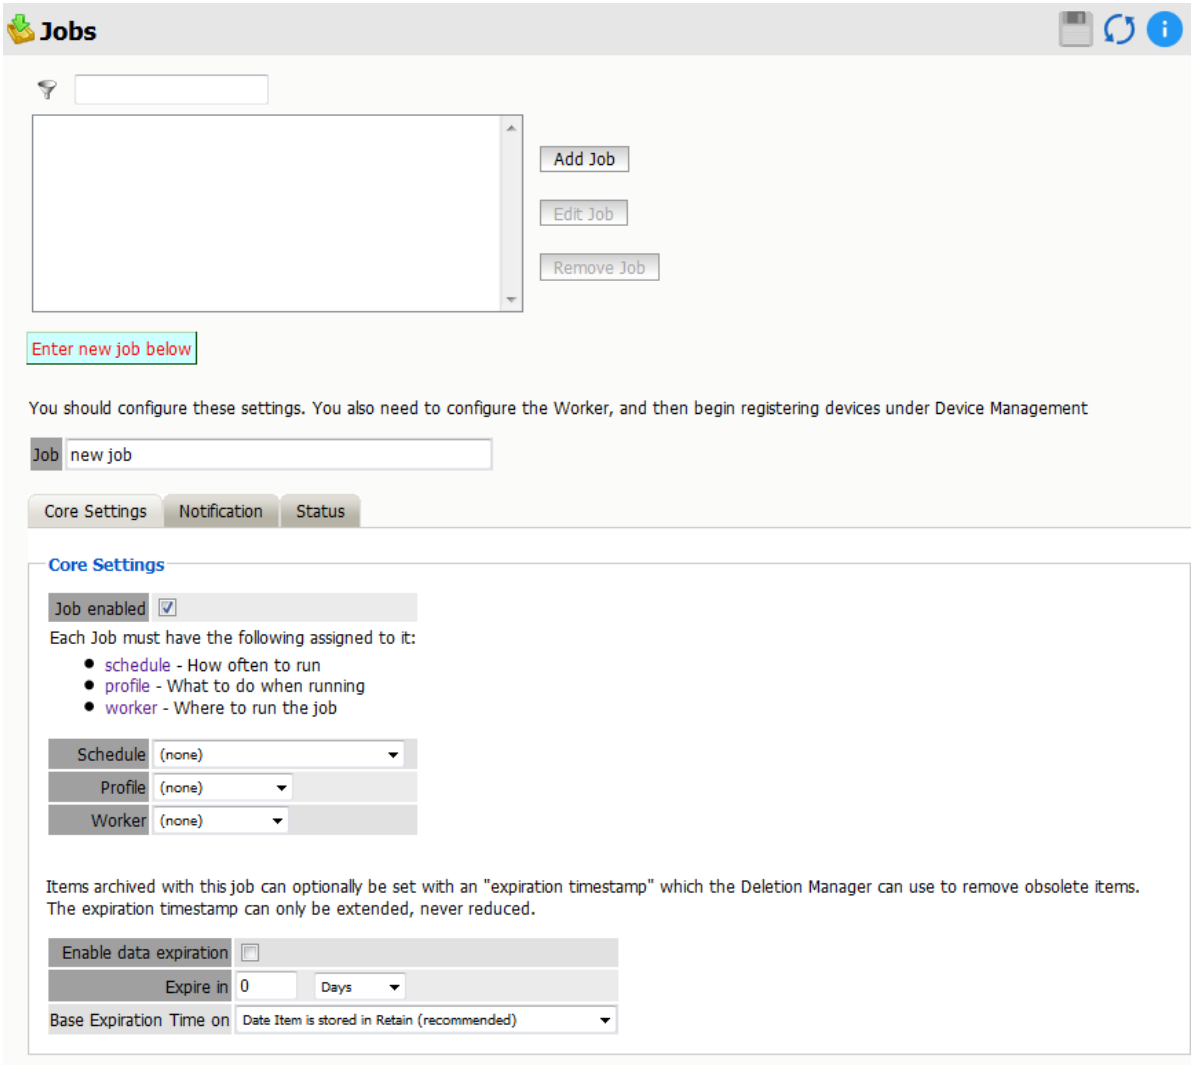
This page is mainly an informational page.

Jobs

Job configuration is only required for legacy Android application support and is not necessary with the current application. The following information is only necessary for supporting older application systems.

Core Settings

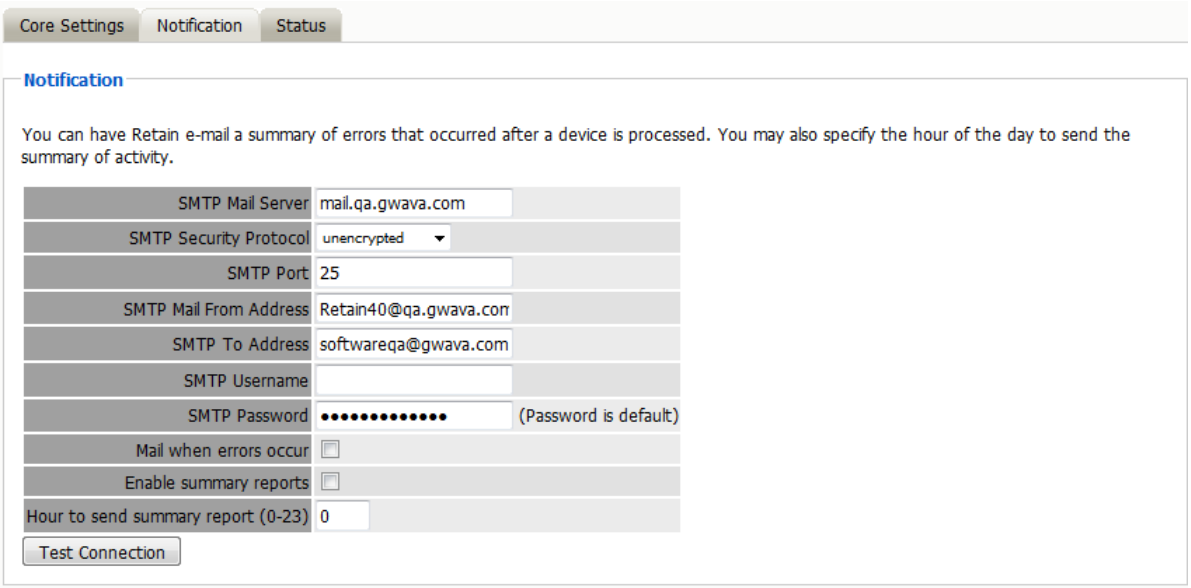
The core settings tab contains information on the schedule, profile, and worker utilized for the mobile module, as well as the data expiration date.



The Data Expiration is an option to place a time stamp on data in the Retain database, which allows for ease of automation for the deletion manager. In addition, devices such as NetApp, Centera, and Hitachi HCAP may use this number to enforce hardware level protection of the stored item so that no one, including Retain, may delete the item before its expiration date. Job Expiration is not retroactive for messages in the database, and only applies to messages archived by the job that it is active for. In order to have messages with custom job or folder expiration dates properly expire, the deletion management date scope must be set to delete messages with an Expiration Date older than 1 day.

Notification

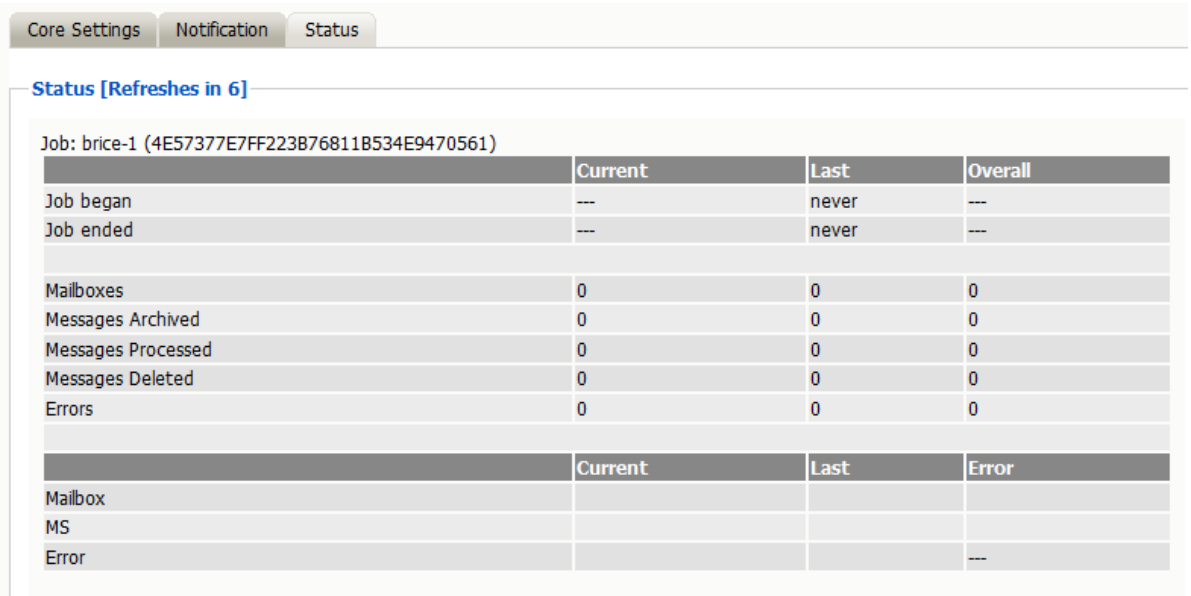
When a job is run, the notification option allows the administrator to be emailed a summary of each running job if desired.



For notification to function correctly, the SMTP information for the desired SMTP server must be fully filled-out. How much information is required varied depending on the mail system used.

Status

The Status tab displays the status of any currently running jobs, as well as the stats of the last completed job.



This tab is informational only for the Mobile module.

Mobile Registration Email variables

The Retain registration email template (…\RetainServer\WEB-INF\classes\config\mobile\registeruser.html) has many variables which allow for a customized registration email. There are a few variables which must remain unchanged, but the rest may be removed or modified, as desired, to adapt the email to any organization.

This is a quick description of the different variables, and which can be modified. They are distinguished by double brackets “[[“and “]]”. The Retain Server code that sends out the email sets these variables with the product specific and user specific information.

Typically, only #1 will want to be changed – for use in mobile device management software.

If an admin wishes to make the registration email to look like it is coming from their company, they should modify #’s 1-2, 6-9. However, variables 6-9 may be simply removed if the admin doesn’t want users to know what they are using.

1. Your install location – This is typically the main variable to change. This variable is not included by default in the email to allow admins to utilize mobile device management software to push the Retain app to devices. It may be included in the email and filled in with the app's location so it can be manually downloaded to devices.
2. Product – This is automatically filled out with the Retain product name. Admins may set this to be the company name. This variable is used in the subject and the body of the email message.
3. Retainheader – This is automatically filled with the Retain product logo. Admins may set this to any branding or company image file. This must be an image file.
4. Name – Each user’s individual username will be input here. This should not be changed.
5. GwavaURL – Automatically filled-in with a link to the Micro Focus website.
6. GwavaURLLabel – This determines text displayed for the GwavaURL link.
7. ProductNameURL – This is automatically filled-in with a link to the product website.
8. ProductNameLabel – This determines text displayed for the productNameURL link.

The registration code and information in that variable must not change. This information is generated by the device management page when the device is added.

The following variables must not change.

Do not modify the following:

- enrollmentCode: This is generated when the device is created on the device manage-ment page.
- tenantid: This is static number for retain. Displayed on the mobile page from Module Configuration.
- email: This is from the Device Management page.
- name: This is from the Device Management page.
- endpoint: This is the address of the Retain Server
- routerURL: This is pulled from the mobile page of Module Configuration.
- activateHREF: This is an identification string, a link specific to the user and installation.

4.18 CellTrust SecureLine

There are two generations of CellTrust Section 4.19, “Generation 1,” on page 236 that uses the module and Section 4.20, “ Generation 2,” on page 240 which uses the REST API.

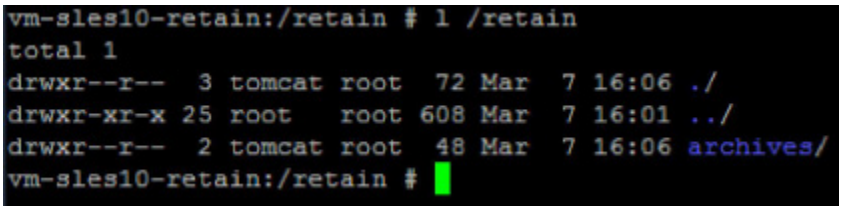
4.19 Generation 1

The CellTrust SecureLine module allows the collection of data from the CellTrust SecureLine server into the Retain system. CellTrust SecureLine supports the archival of SMS from both Android and iOS phones. All devices which are supported by the SecureLine server are covered by Retain.

To archive CellTrust SecureLine information, the SecureLine server must be configured to log SMS data. The server log files are what Retain uses to collect the information about SMS activ-ity. If the SecureLine server has not been configured to log the data, Retain will be unable to archive the information. See below for SecureLine server configuration.

Configuring the CellTrust SecureLine Module is incredibly simple. Essentially, the module needs to be enabled and configured with the log file location. Afterwards, the data is treated the same as any other data entering the Retain archive. Data is archived in the database under the user/phone it originated from and can be searched and viewed through the Search Messages interface.

To enable the CellTrust module, select the ‘Configure’ link under the ‘CellTrust’ icon from the module page.

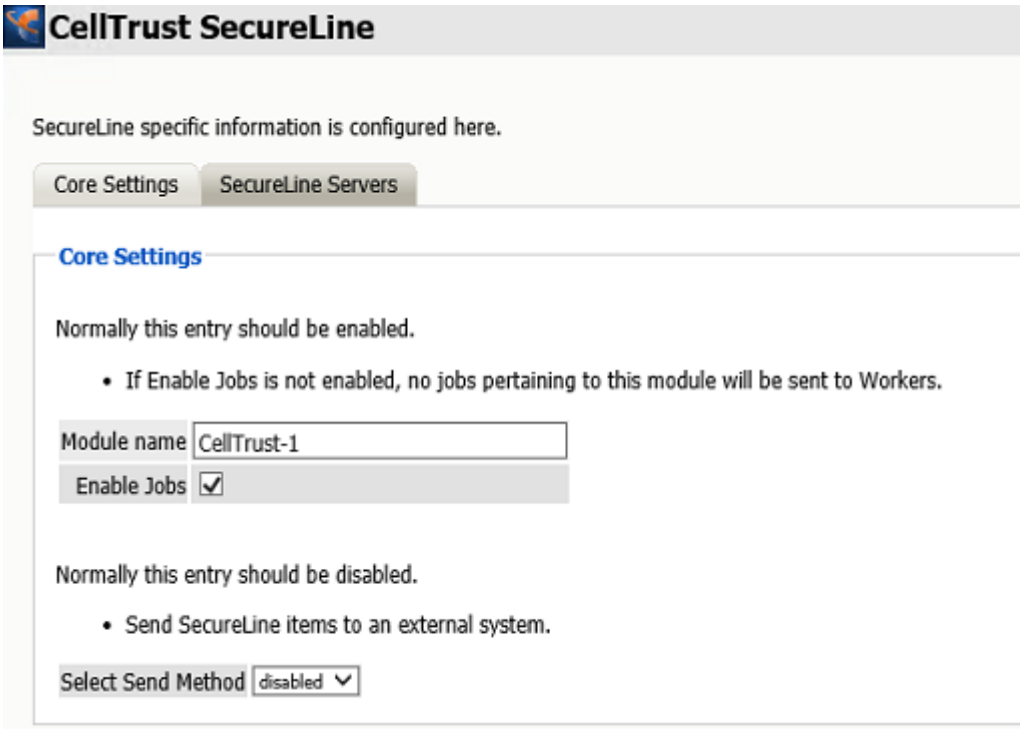


4.19.1 Core Settings

The module needs to be enabled on this page to make it active in the Retain system.

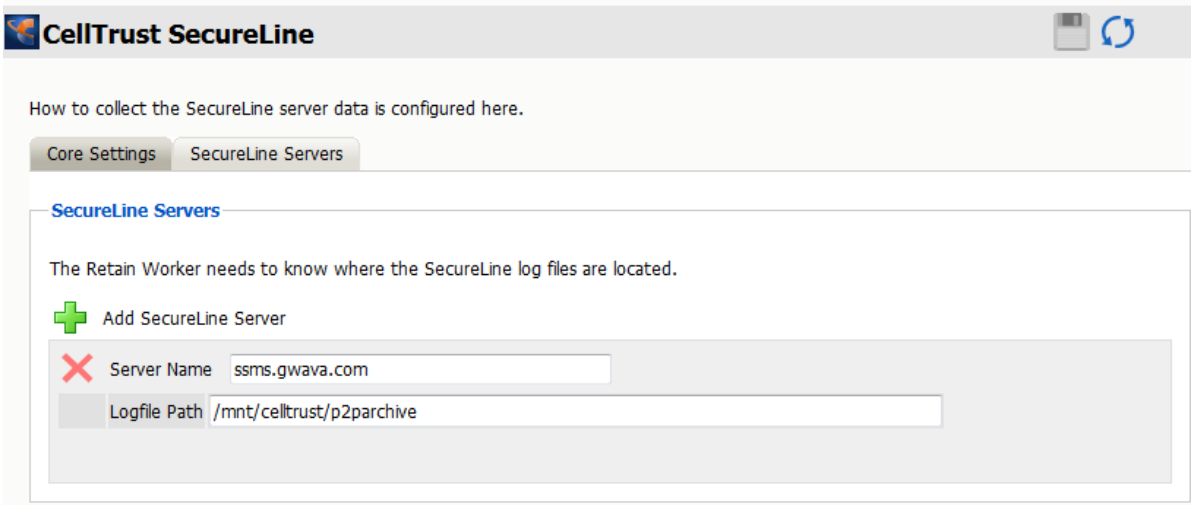
The module can be given a name.

The Send Method option enables either the SMTP Forwarding or FTP features. For either feature to appear and function, the Module Forwarding tab must be configured on the Server Configuration page. See that section for more information.



4.19.2 SecureLine Servers

The log file path for each server in the system is required to archive the SMS data. To add a server to the Retain process, select the Green Plus Sign(+). Give the server a name (any name will do), and provide the log file path.



The log files must be directly accessible to the worker tasked with archiving the CellTrust SecureLine data, whether through direct access or through a mapped drive. The logfile path is the direct file access path to the log files of the CellTrust SecureLine server, (/opt/PM/scheduler/external by default), and should be specified as the worker will access the log files.

4.19.3 Archiving/Data Collection – About Workers, Profiles, Schedules, and Jobs

Getting data into the archive is the core functions of Retain. As such, sub-functions have been broken out to give you as much flexibility as possible in deciding your archive policy.

Your archiving policy will cover several points:

- What to archive (jobs)
- Includes what message server(s) to collect data from.
- What worker
- When will the data be collected (schedule)
- How to collect the data will cover (profile)
- Types of items or users to archive.
- Date range of items to be archived.
- Avoiding redundant data collection.
- Utilizing storage flags to achieve 100% retention.
- Whether to store attachments.
- Advanced criteria, such as attachment names, or folder locations.

4.19.4 Retain Job components

Schedule

The Schedule is the time framework that is set for the job to run on and provides the choice of how often, and when, the job is run. Single time, (job runs once), or a recurring time frame are available. Schedules are universal and are not module or job specific. They may be used for one or all jobs.

Profile

The Profile dictates which kinds of files to retrieve from the message system. All message system items are selectable to be archived, but you may also specify time periods, exclude folders or subjects, and how the data is stored. Profiles are module specific, and the profile section contains sections specified by each module and information only applicable to that module.

The Profile also has a very important function, duplicate checking. This keeps the Retain system from storing multiple instances of the same message. Make sure you check and read the section on Purge (GroupWise), Retention (GroupWise), and Item Store (Exchange/GroupWise) flags.

Worker

The worker is the part of Retain that actually pulls the data from the message system servers. The Worker contacts the Retain Server and gathers information on which jobs are associated with it, when to start them, and which items to get. Workers can only run one job at a time. Workers are not module specific, but may require module specific settings. These settings, where applicable, are found under the module specific tab.

Jobs

The Jobs section is where the Worker, Profile, and Schedule are all brought together and associated into a comprehensive configuration by allowing the user to tie their configuration to a selected SecureLine server and selected user list. The desired SeucureLine server must be selected. The Worker will not start archiving anything until a created job is bound to a Worker, Schedule, and Profile.

4.19.5 Secureline Server Configuration

Retain requires the CellTrust SecureLine server to create log files for the worker to access, digest, and send to the Retain Server for storage. To configure the SecureLine server to create these files requires the scheduled run of a Gwava Report script, a tool provided by CellTrust.

The script may be scheduled through the quartz scheduler.

To configure the quartz scheduler to run the GwavaReport script, navigate to the /opt/PM/bin directory and run the start_sched_client.sh script, then follow the desired prompts.

For example:

```
[root@dev04120lv bin]# ./start_sched_client.sh
Hi, welcome to PrimeMessage Scheduler configuration utility
You will be prompted to enter different values. Every time you finish to enter
something - just hit "Enter". You can quit any time - just enter single 'q'
character in the beginning of the line.
GOOOOD LUCK!
Select operation ( 1 - add job; 2 - remove job; 3 - pause job; 4 - resume job; 5 -
current jobs):
1
Enter job name:
GwavaReport
Enter job class name ( omit ".class" part ):
GwavaReportJob
Please select Trigger type( 1 - Simple, 2 - Cron):
```

At this point, select whether you want to use a Cron configuration for the triggers or a simple type – this is an example using the simple type:

```
Please select Trigger type( 1 - Simple, 2 - Cron):
1
Enter delay ( in seconds ) before first execution:
30
Enter number of executions ( 0 for indefinite ):
2
Enter interval between executions ( in seconds ):
30
Enter job additional params ( String or int only)
When finished, enter empty line for attribute name
Enter attribute name:
```

The GwavaReport has no quartz attributes, so hit enter, then ctrl-c to exit the script.

Schedule the report to run as often as desired according to traffic and desired archival.

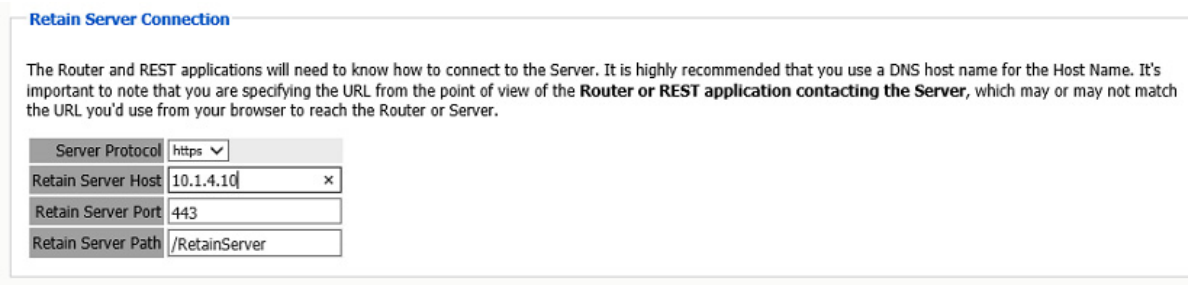
Once the report tool has run, it will place the necessary log file in the /opt/PM/scheduler/external directory. The files here are required by the Retain Worker. The Retain Worker required direct file access to these files. Whether by export or by local access, configure the CellTrust SecureLine module with the location of these files, the file location must be from the worker's point of view.

4.20 Generation 2

CellTrust generation 2 uses the REST Collector model.

4.20.1 Configure Retain Server Connection

Under *Server Configuration/Communications* you must setup the Retain Server Connection. This is the hostname or the public IP Address of the Retain server the host will be entered into the CellTrust system to connect to Retain. The port will be 48080 for HTTP or 443 for HTTPS protocols.

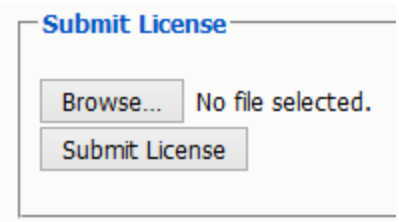


4.20.2 Install REST Licenses

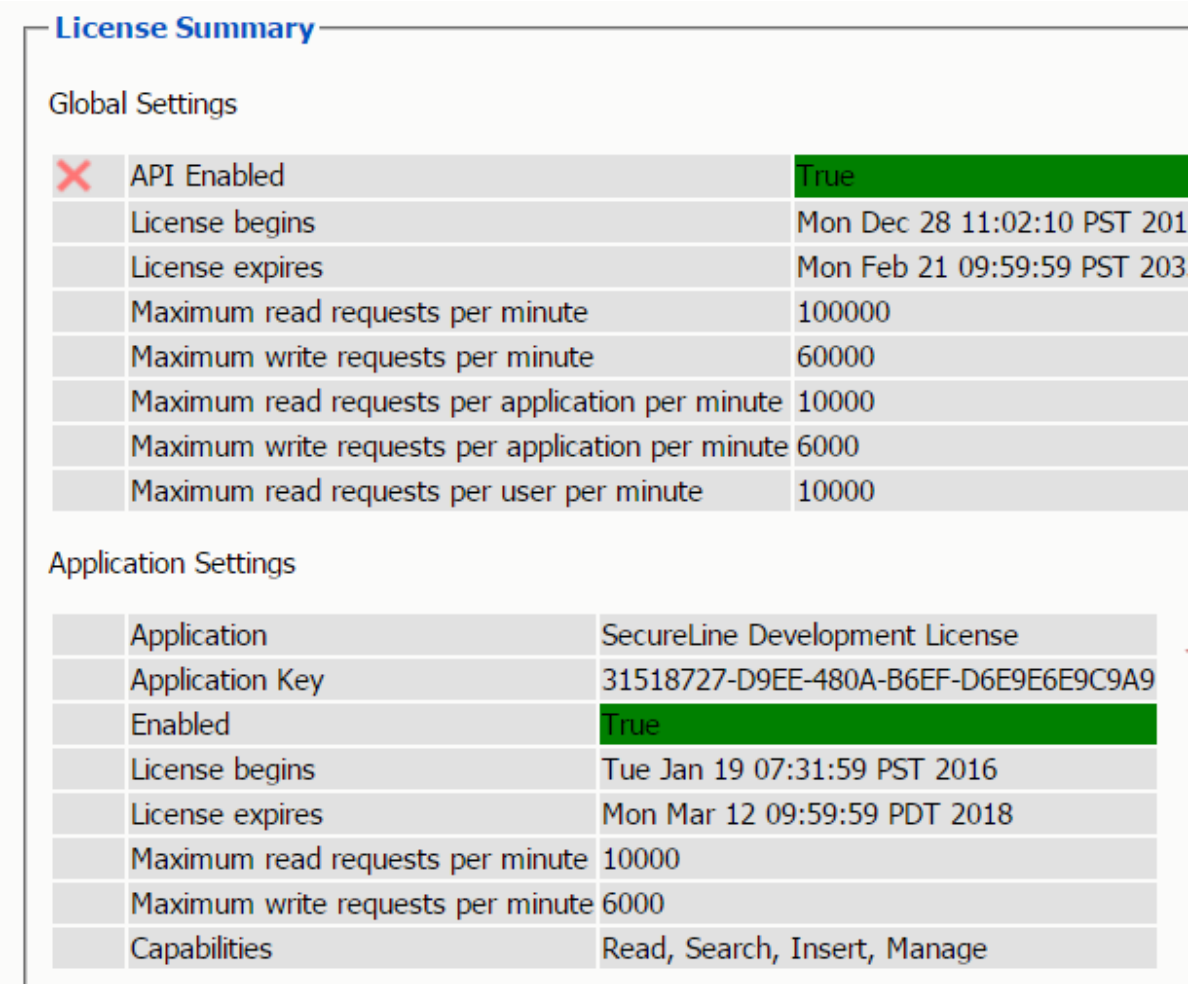
Beyond Installing the module license PEM file to the Retain Server under the License page, you must install the REST license files. These will be downloaded as a ZIP file, there is no need to unzip this file.

Go to *Server Configuration/REST API.*

Scroll down to Submit License, Browse to the ZIP file and press Submit License.



The license will be uploaded to the Retain server and appear in the list under Application Settings.



4.20.3 Create REST Collector

Create a REST Collector by providing:

Collector Name



The Key and Secret will be used on the CellTrust side of the system for authentication.

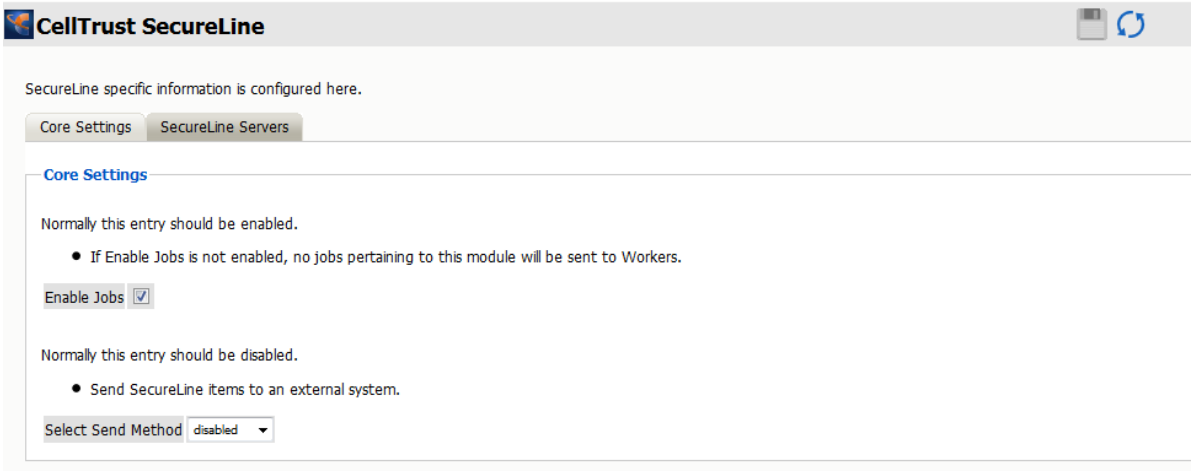
Then press "Add collector" to enable the collector.

You may delete the collector.

Export the collector allows you to import the Key and Secret into the CellTrust system.

4.20.4 Configure the Module

To configure the CellTrust module, all that is required is to enable the Jobs in the module and Save Changes.



The “Enable Jobs” box must be checked to have data archiving functionality.

The Send Method option enables either the SMTP Forwarding or FTP features. For either feature to appear and function, the Module Forwarding tab must be configured on the Server Configuration page. See that section for more information.

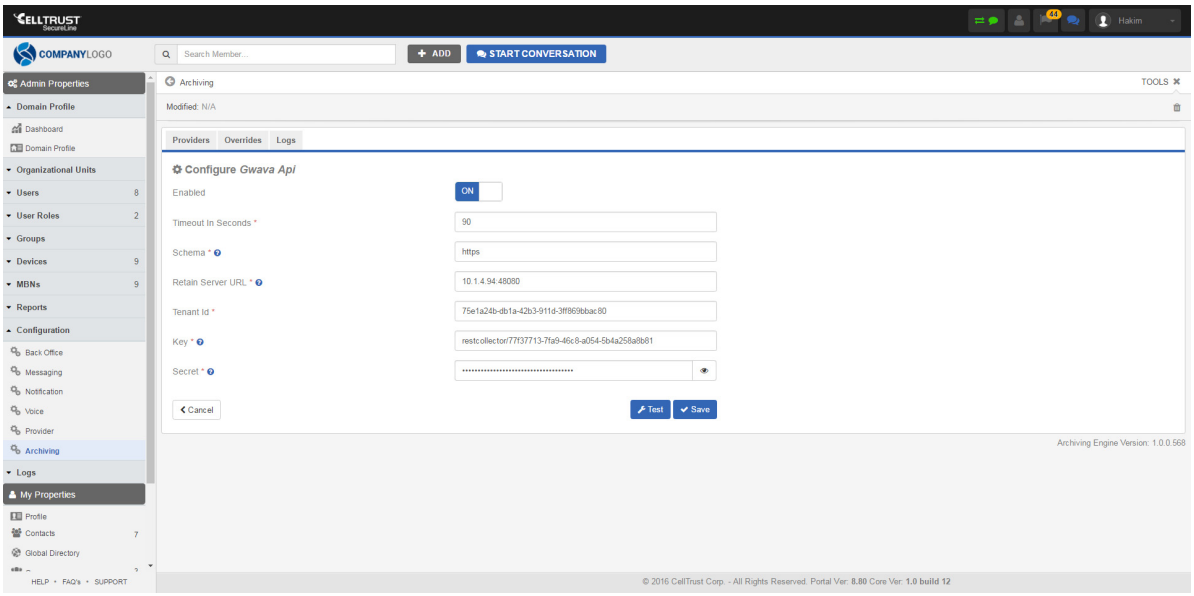
Save changes and the module is configured. You may need to disable and re-enable "Enable Job" before it will allow you to Save Changes.

There is no need to configure the SecureLine Servers.

Configure CellTrust

The REST connection to the collector needs to configured by the Admin on the CellTrust system side.

Provide the Retain Server URL, Tenant ID, Key and Secret.



Once the connection is made data will begin to flow into Retain from CellTrust and will appear in the user's Retain mailbox.

There are no Profiles or Jobs associated with this Module.

4.21 GBS Notes Module

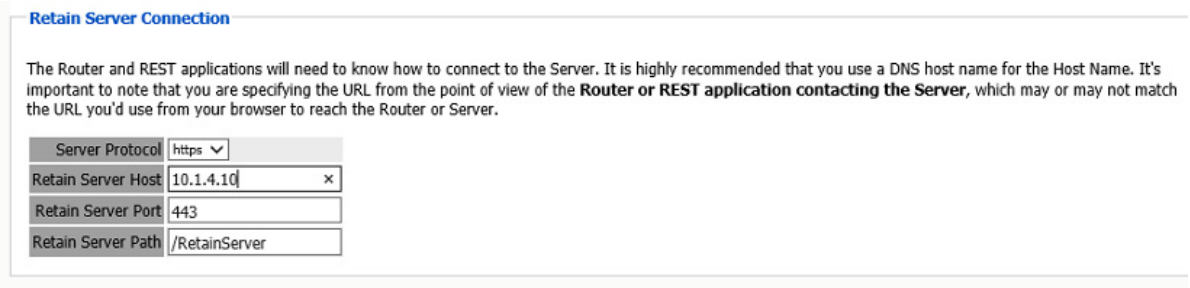
The GBS Notes module allows Retain to archive Lotus Notes/Domino communications via GBS. This is different from most modules in that there are no Profiles or Jobs. A REST API collector is connected between the GBS Notes and Retain systems and the data flows into the Retain archive as it created in the GBS Notes system. Data is archived immediately as it enters or is created in the GBS Notes system.

4.21.1 Install Module License

You will receive an email with your license validation keys. Go to *http://licenses.gwava.com*, enter each key for the server and modules and download the .PEM and .ZIP files. Upload the .PEM files to the Retain Server on the *Licenses* page. The REST API license .ZIP file will be uploaded to a different location discussed below.

4.21.2 Configure Retain Server Connection

Under *Server Configuration/Communications* you must setup the Retain Server Connection. This is the hostname or the public IP Address of the Retain server the host will be entered into the GBS system to connect to Retain. The port will be 48080 for http or 443 for https protocols.

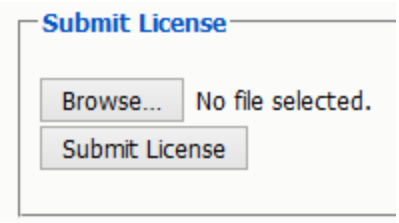


4.21.3 Install REST Licenses

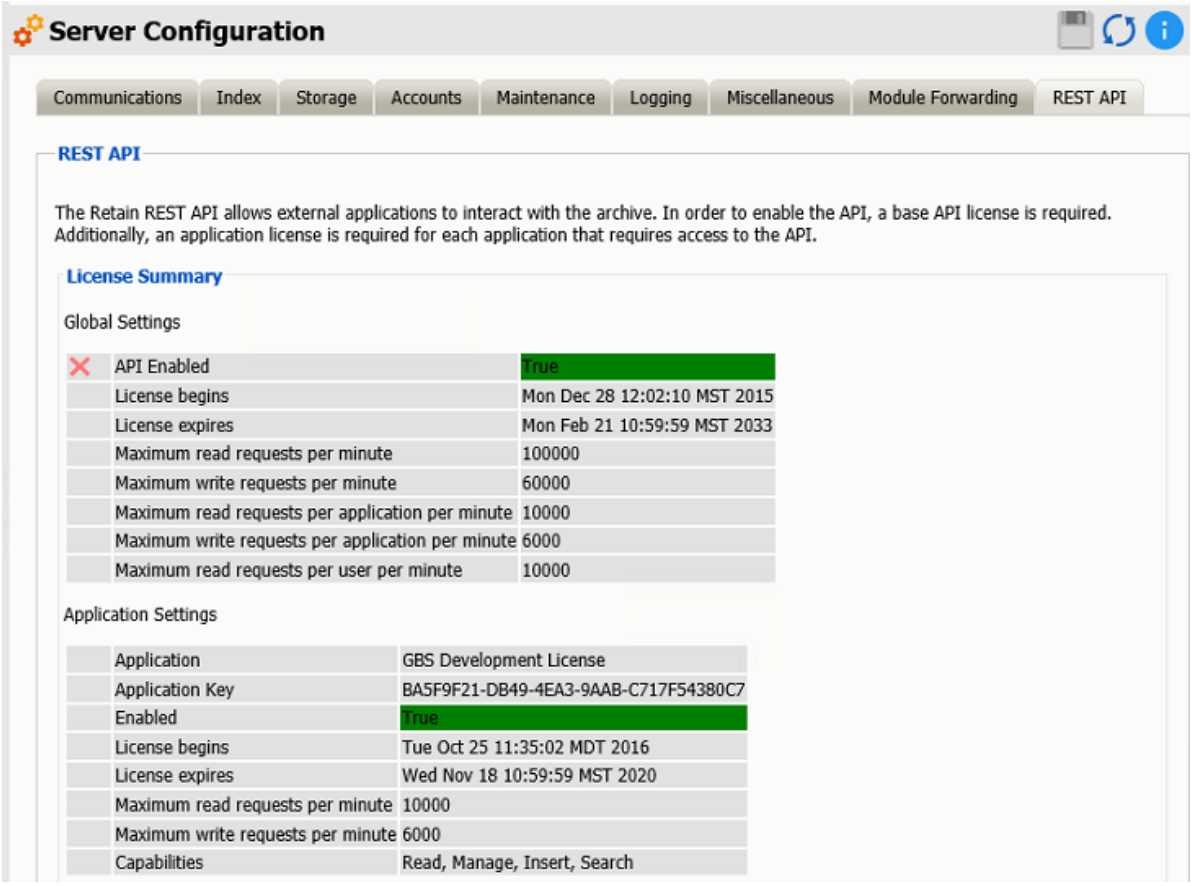
Beyond Installing the module license .PEM file to the Retain Server under the License page, you must install the REST license files. These will be downloaded as a .ZIP file, there is no need to unzip this file.

Go to *Server Configuration/REST API.*

Scroll down to Submit License, Browse to the .ZIP file and press Submit License.



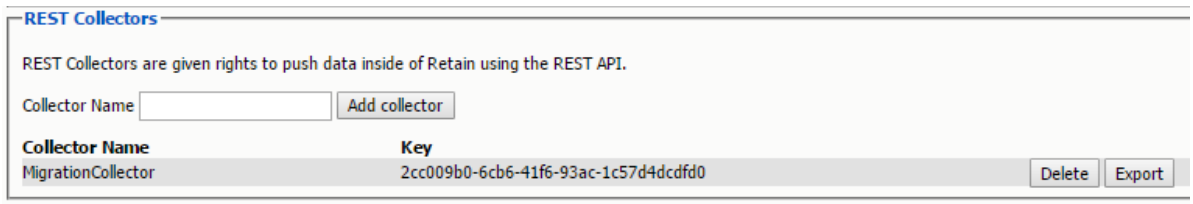
The license will be uploaded to the Retain server and appear in the list under Application Settings.



4.21.4 Create REST Collector

Create a REST collector under Server Configuration

Some modules use the REST API for data collection. You need to create a collector to connect to the message system. You may have multiple REST collectors.



Create a REST Collector by providing:

1. The *Collector Name* can be anything you want, but we recommend something easy to identify with the system it connects with. A Key and Secret will be created for the collector. The Key and Secret will be used on the message system side of the system for authentication.
2. Then press "*Add collector*" to enable the collector.
3. *Export* the collector allows you to export the Key and Secret to disk for use with the message system.

Delete removes the collector.

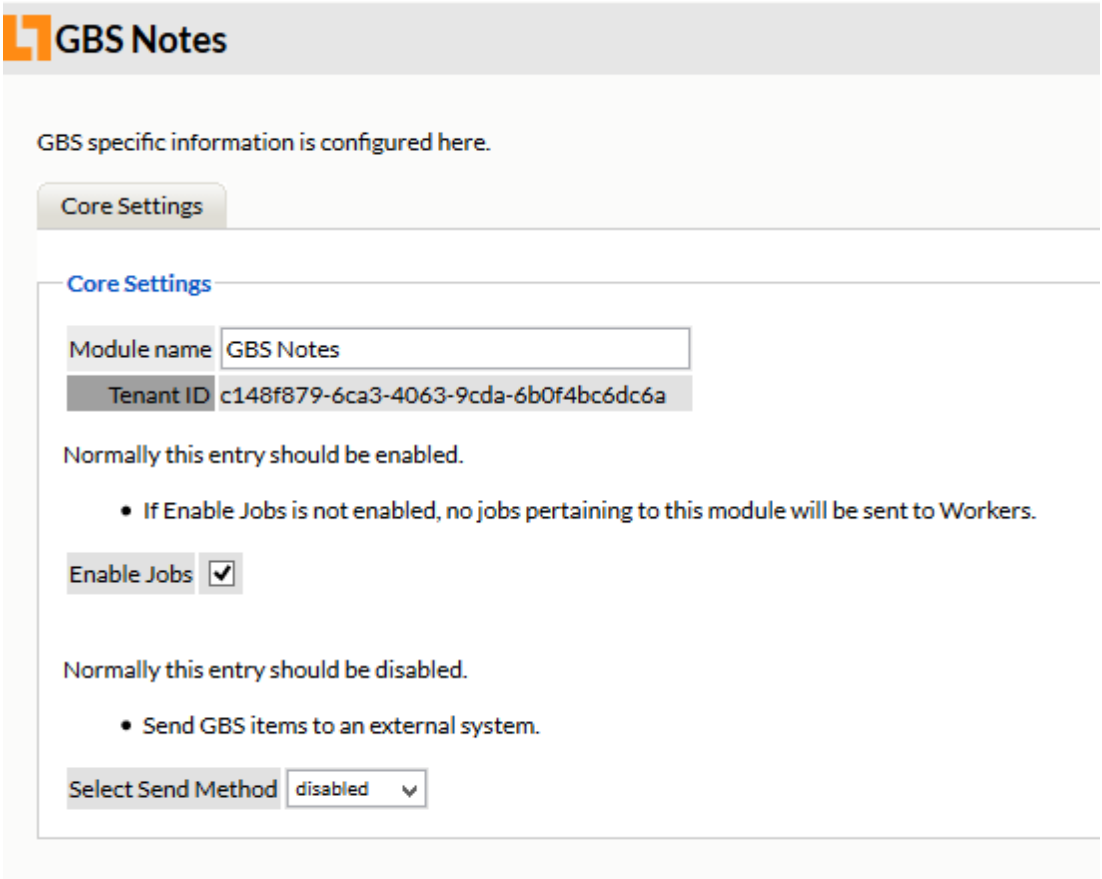
4.21.5 Configure the Module

The module needs to be enabled on this page to make it active in the Retain system.

The module can be given a name.

The Send Method option enables either the SMTP Forwarding or FTP features. For either feature to appear and function, the Module Forwarding tab must be configured on the Server Configuration page. See that section for more information.

To configure the GBS Notes module, all that is required is to enable the Jobs in the module and Save Changes.



The “Enable Jobs” box must be checked to have data archiving functionality.

The Send Method option enables either the SMTP Forwarding or FTP features. For either feature to appear and function, the Module Forwarding tab must be configured on the Server Configuration page. See that section for more information.

Save changes and the module is configured. You may need to disable and re-enable "Enable Job" before it will allow you to Save Changes.

Configure GBS Notes

The REST connection to the collector needs to configured by the Admin on the GBS Notes system side.

Once the connection is made data will begin to flow into Retain from GBS Notes and will appear in the user's Retain mailbox.

There are no Profiles or Jobs associated with this Module.

4.22 Social Messaging Module

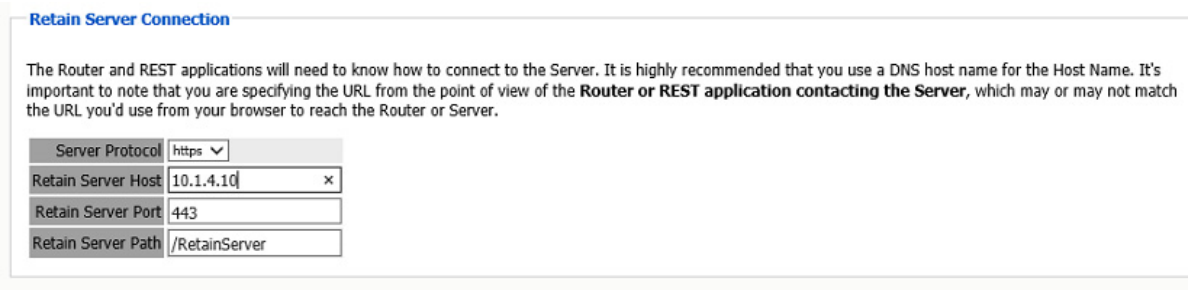
The Social Messaging module allows Retain to archive Social Messaging communications from Twitter, Facebook, Yammer, AOL chats, LinkedIn, Yahoo Messenger, Google Talk, Lync & Skype for Business, and Gmail. The Social Messaging module also archives searches performed on Google, YouTube, Yahoo, Bing, and Wikipedia.

4.22.1 Install Module License

You will receive an email with your license validation keys. Go to *http://licenses.gwava.com*, enter each key for the server and modules and download the PEM and ZIP files. Upload the PEM files to the Retain Server on the *Licenses* page. The REST API license ZIP file will be uploaded to a different location discussed below.

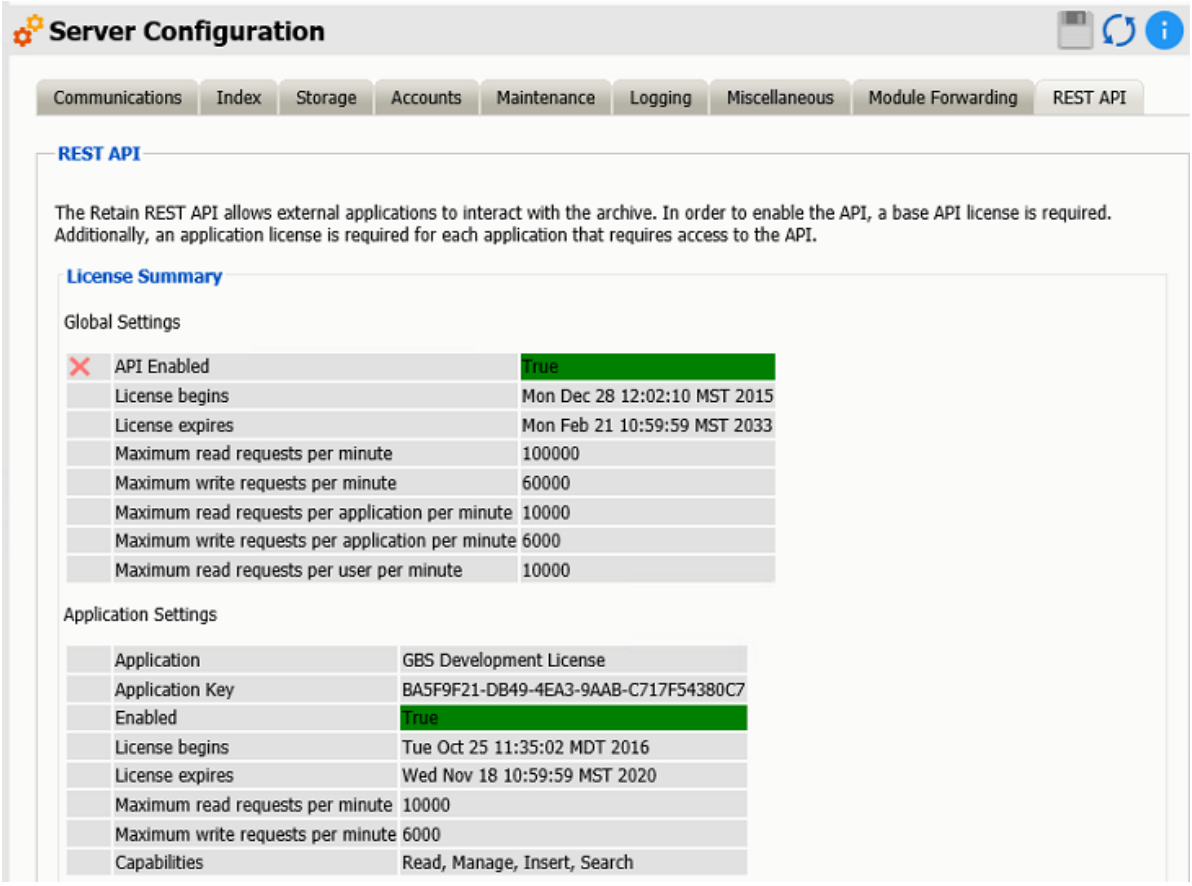
4.22.2 Configure Retain Server Connection

Under *Server Configuration/Communications* you must setup the Retain Server Connection. This is the hostname or the public IP Address of the Retain server the host will be entered into the Social system to connect to Retain. The port will be 48080 for HTTP or 443 for HTTPS protocols.

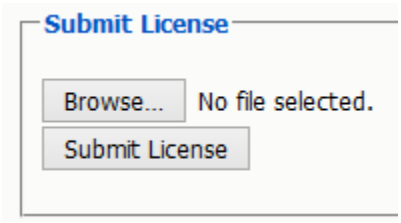


4.22.3 Install REST Licenses

Beyond Installing the module license PEM file to the Retain Server under the License page, you must install the REST license files. These will be downloaded as a ZIP file, there is no need to unzip this file. Go to *Server Configuration/REST API.*



Scroll down to Submit License, Browse to the ZIP file and press Submit License.



The license will be uploaded to the Retain server and appear in the list under Application Settings.

4.22.4 Create REST Collector

Create a REST Collector by providing:

Collector Name



The Key and Secret will be used on the Social side of the system for authentication.

Then press "Add collector" to enable the collector.

You may delete the collector.

Export the collector allows you to import the Key and Secret into the Social system.

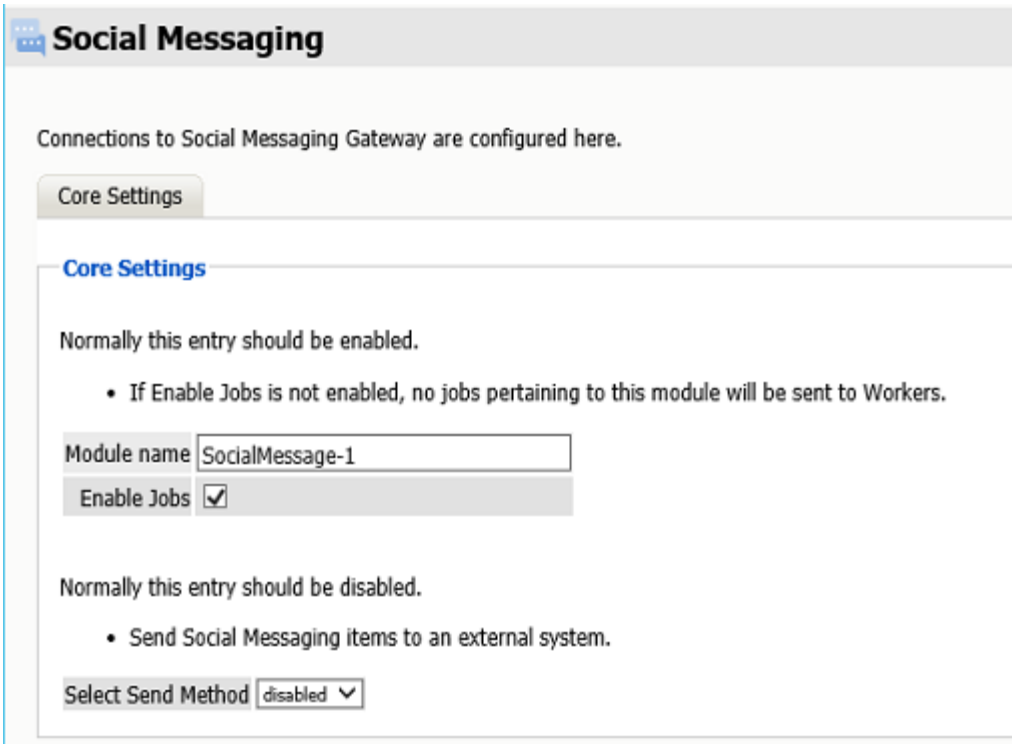
4.22.5 Configure the Module

The module needs to be enabled on this page to make it active in the Retain system.

The module can be given a name.

The Send Method option enables either the SMTP Forwarding or FTP features. For either feature to appear and function, the Module Forwarding tab must be configured on the Server Configuration page. See that section for more information.

To configure the Social Messaging module, all that is required is to Enable Jobs in the module and Save Changes.



The “Enable Jobs” box must be checked to have data archiving functionality.

The Send Method option enables either the SMTP Forwarding or FTP features. For either feature to appear and function, the Module Forwarding tab must be configured on the Server Configuration page. See that section for more information.

Save changes and the module is configured. You may need to disable and re-enable "Enable Job" before it will allow you to Save Changes.

Configure Social

The REST connection to the collector needs to configured by the Admin on the Social system side.

You will need the Tenant ID, Key, and Secret from the Router Configuration Page.

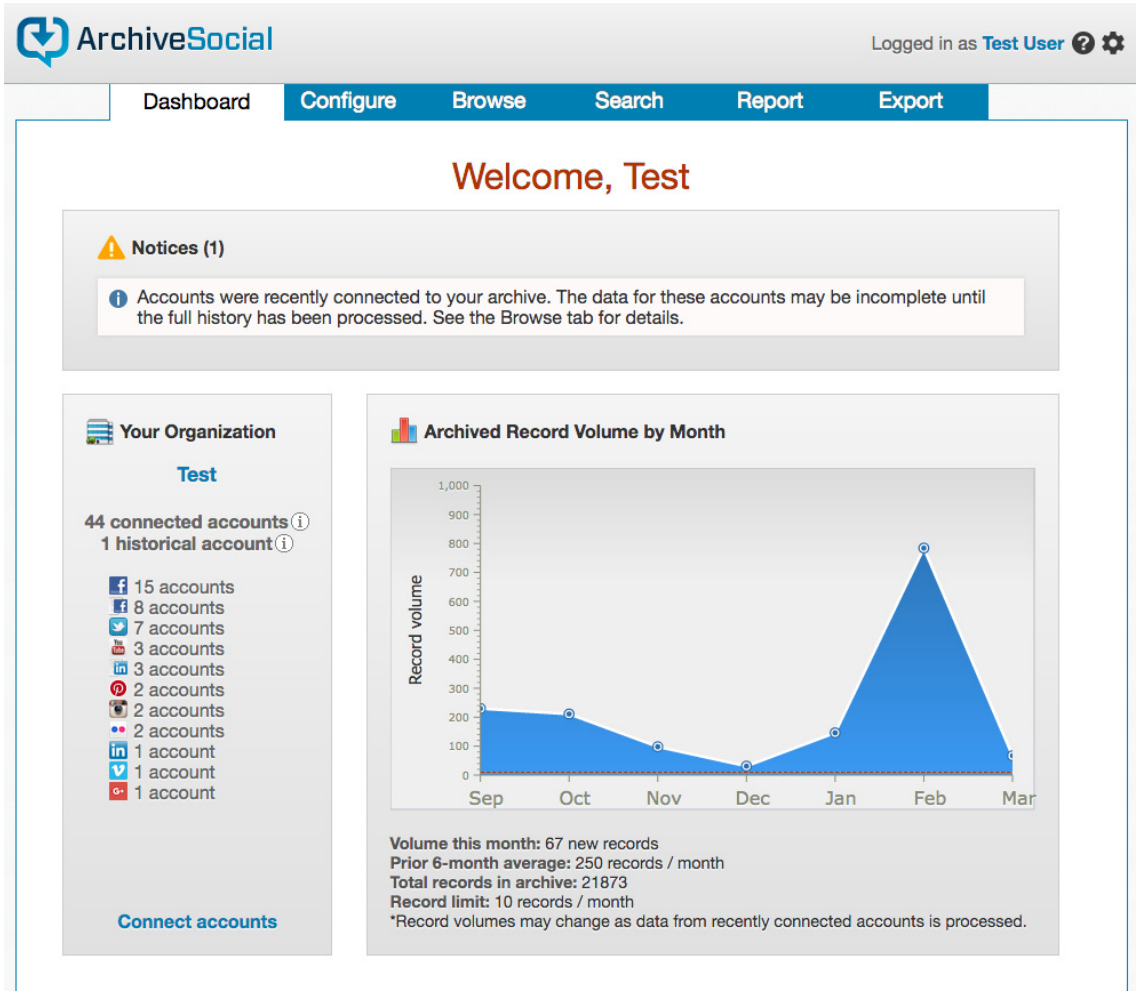
That information must be given to ArchiveSocial.

Once the connection is made data will begin to flow into Retain from Social and will appear in the user's Retain mailbox.

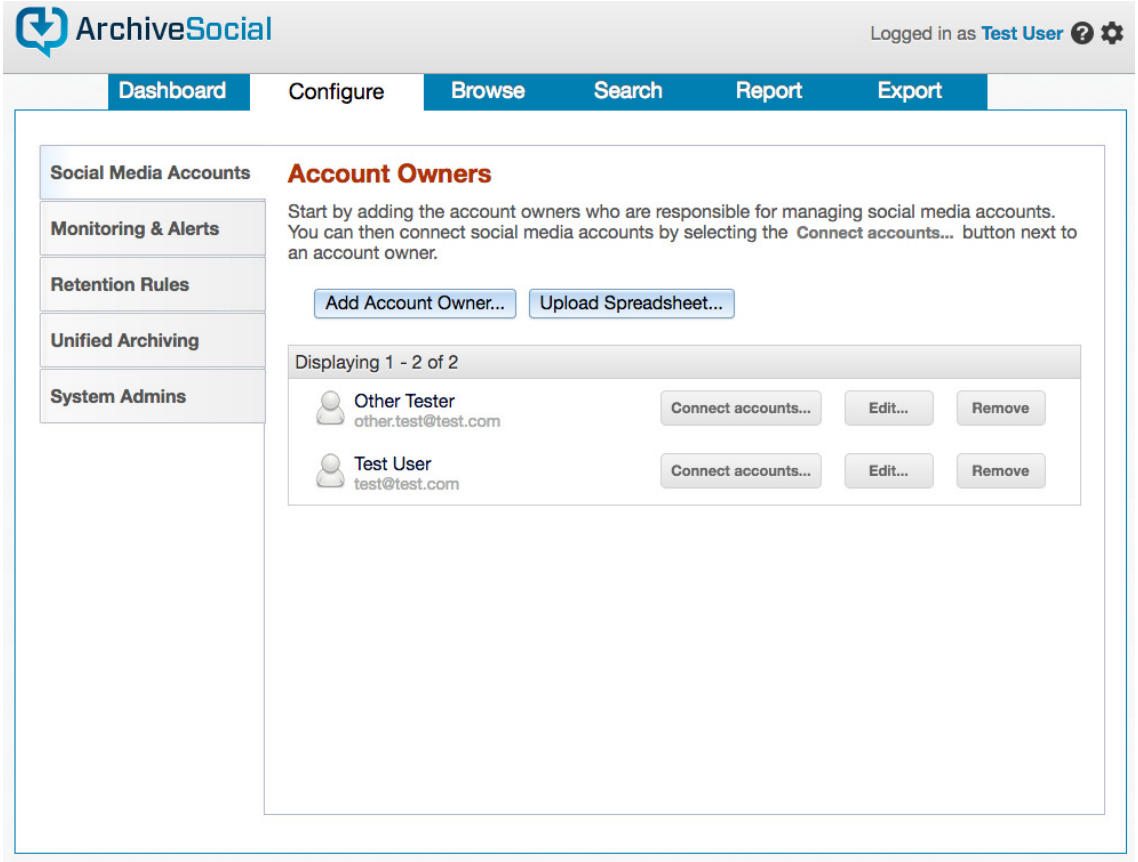
There are no Profiles or Jobs associated with this Module.

In ArchiveSocial

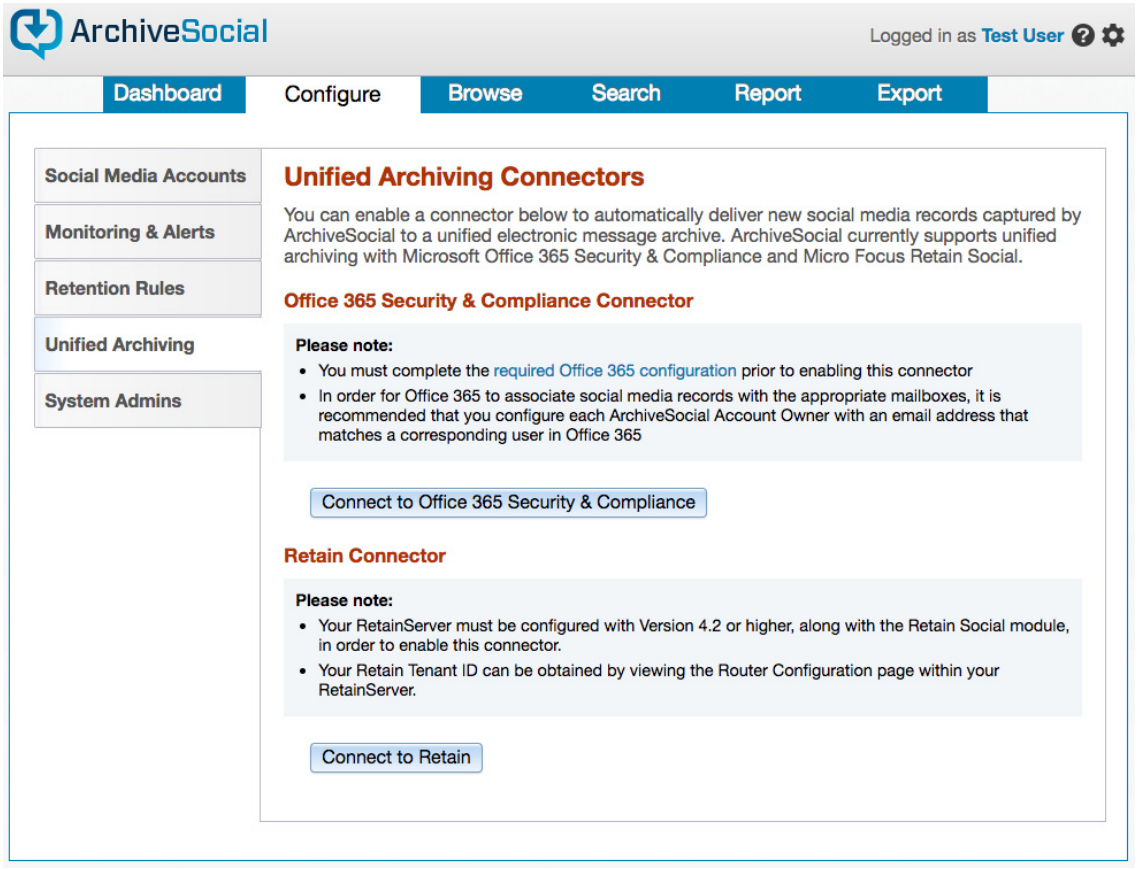
1. Login to your ArchiveSocial account. After logging in you should see the dashboard which looks similar to the screenshot below:



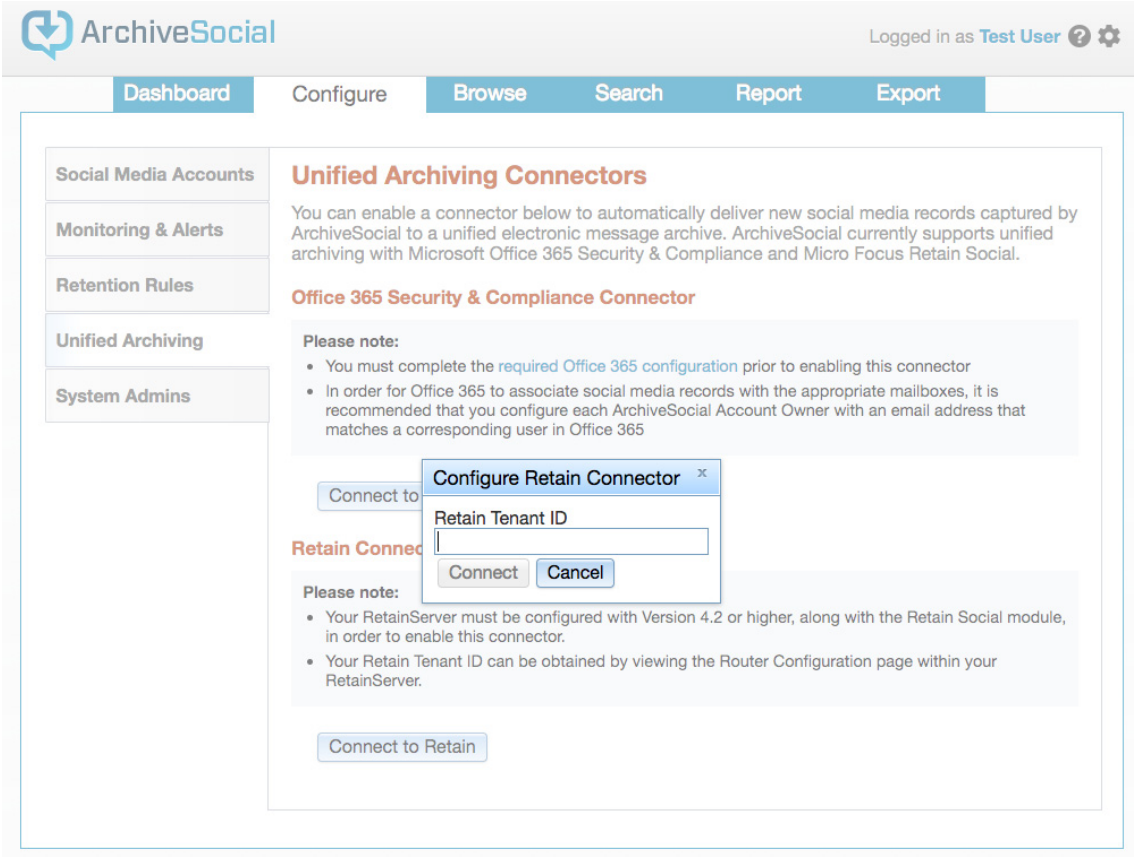
2. Click on the configure tab near the top of the page. This should take you to the configure screen:



3. On the left sidebar there should be a Unified Archiving tab click there to open up page for connecting integrations:


ArchiveSocial
Logged in as Test User
Dashboard
Configure
Browse
Search
Report
Export
Social Media Accounts
Monitoring & Alerts
Retention Rules
Unified Archiving
System Admins
Unified Archiving Connectors
You can enable a connector below to automatically deliver new social media records captured by ArchiveSocial to a unified electronic message archive. ArchiveSocial currently supports unified archiving with Microsoft Office 365 Security & Compliance and Micro Focus Retain Social.
Office 365 Security & Compliance Connector
Please note:
You must complete the required Office 365 configuration prior to enabling this connector
In order for Office 365 to associate social media records with the appropriate mailboxes, it is recommended that you configure each ArchiveSocial Account Owner with an email address that matches a corresponding user in Office 365
Connect to Office 365 Security & Compliance
Retain Connector
Please note:
Your RetainServer must be configured with Version 4.2 or higher, along with the Retain Social module, in order to enable this connector.
Your Retain Tenant ID can be obtained by viewing the Router Configuration page within your RetainServer.
Connect to Retain

4. Click the Connect to Retain button and input whatever tenant id you want to send content to and hit connect and then you're all set. At this point new content captured will be sent to the specified tenant.



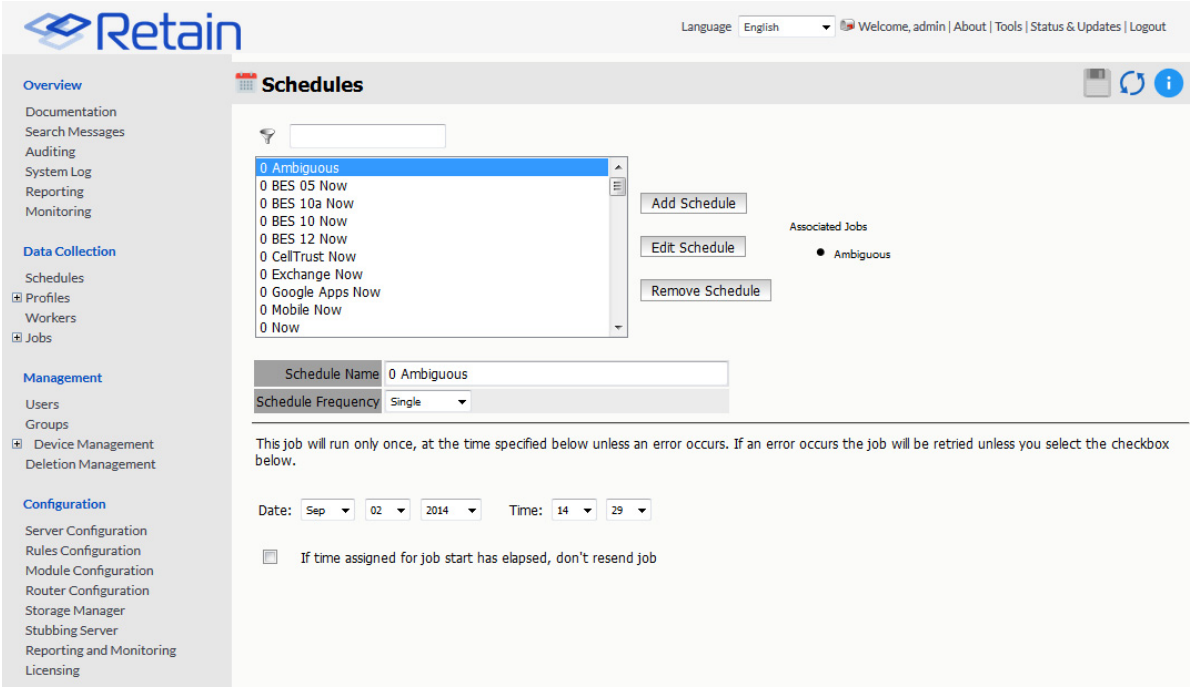
4.23 Schedules

Schedules define when to perform the data collection as defined in a job. a job defines what schedule, profile and worker to use.

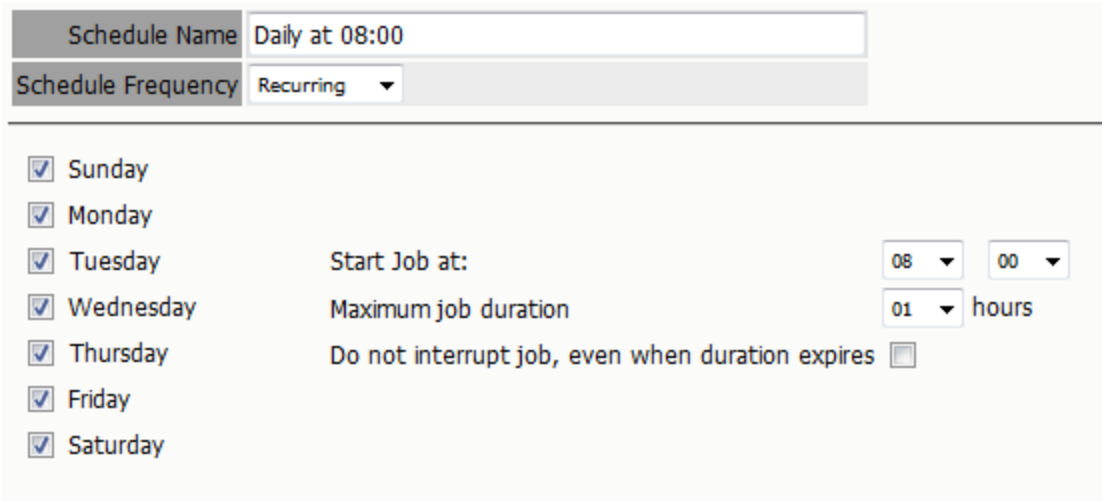
- You can create as many schedules as you need.
- Schedules have no functionality until they are associated with one or more jobs.
- They can be shared across multiple jobs.
- You cannot delete a schedule that is currently associated with a job.
- Schedules can be recurring or one-time.

Single schedules are configured with one start time, and one date.

To keep jobs from queuing up to the worker, enable (check) "If time assigned for job start has elapsed, don't resend job". Otherwise, a job will queue up on the worker if a previous job has not yet completed, causing the new job to start as soon as the current job is complete or terminated.



Recurring schedules are configured using a start and interrupt time, and which days to start on.



Choose when on which days you want this schedule to trigger a job.

You may specify maximum job duration. If the job runs longer than the time limit you set, it will be inter-rupted. For cases where you do not want the job to be interrupted. In an example with a job you know will last longer than 22 hours, you may prevent the schedule from interrupting the job.

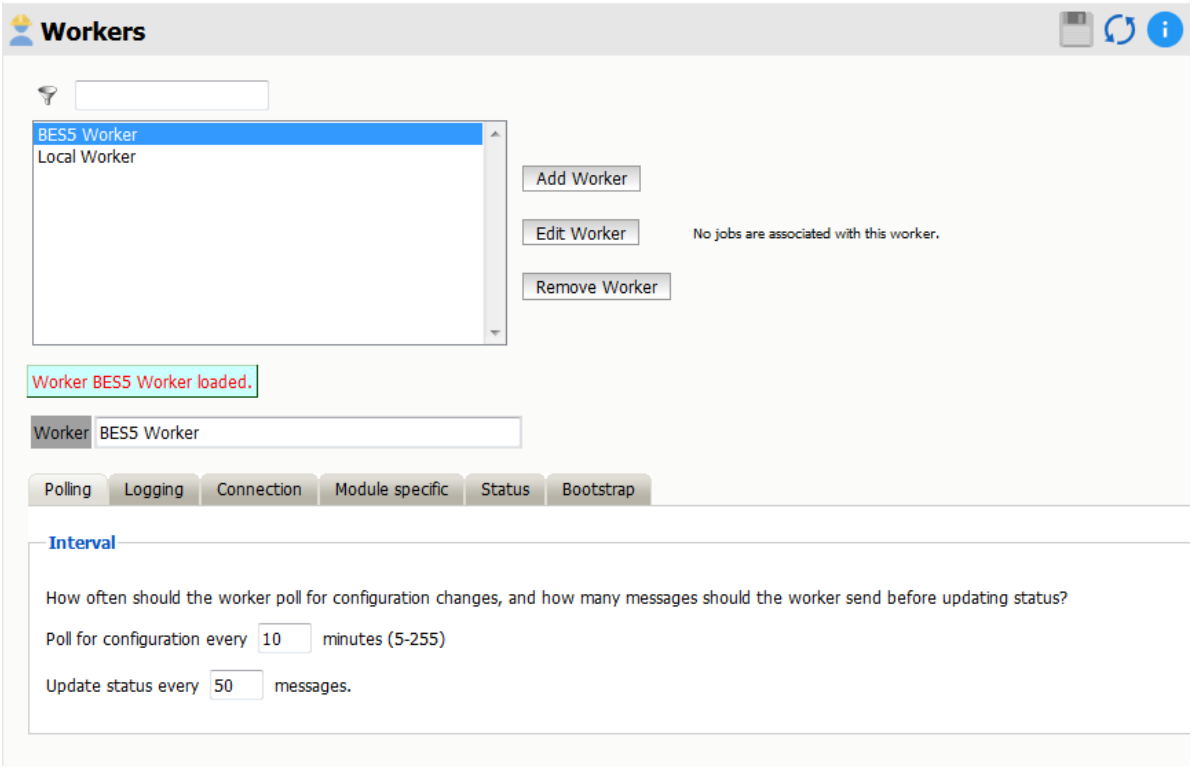
Don't forget to "Save Changes".

4.24 Workers

You will need to set up at least one Retain Worker. The Retain Worker does the actual work of collecting the data from a message system and delivering it to the Retain Server. The Retain Worker component is normally installed local to each GroupWise mail server, or any Exchange server from

which you wish to collect data. Although a Worker may be installed on a machine running Retain Server, this is normally not recommended or sufficient, except for very small (100 or fewer users) systems or for evaluation pur-poses. Ideally, the worker you choose to dredge mail from a desired server is local to the server itself. Using a worker on Server A to pull data from a mailbox server on Server B is possible, but it will use up network bandwidth and will be slower than simply pulling data locally. For this reason, Retain was designed to allow you as many workers as you need, on Linux or Windows.

4.24.1 Worker Configuration Options



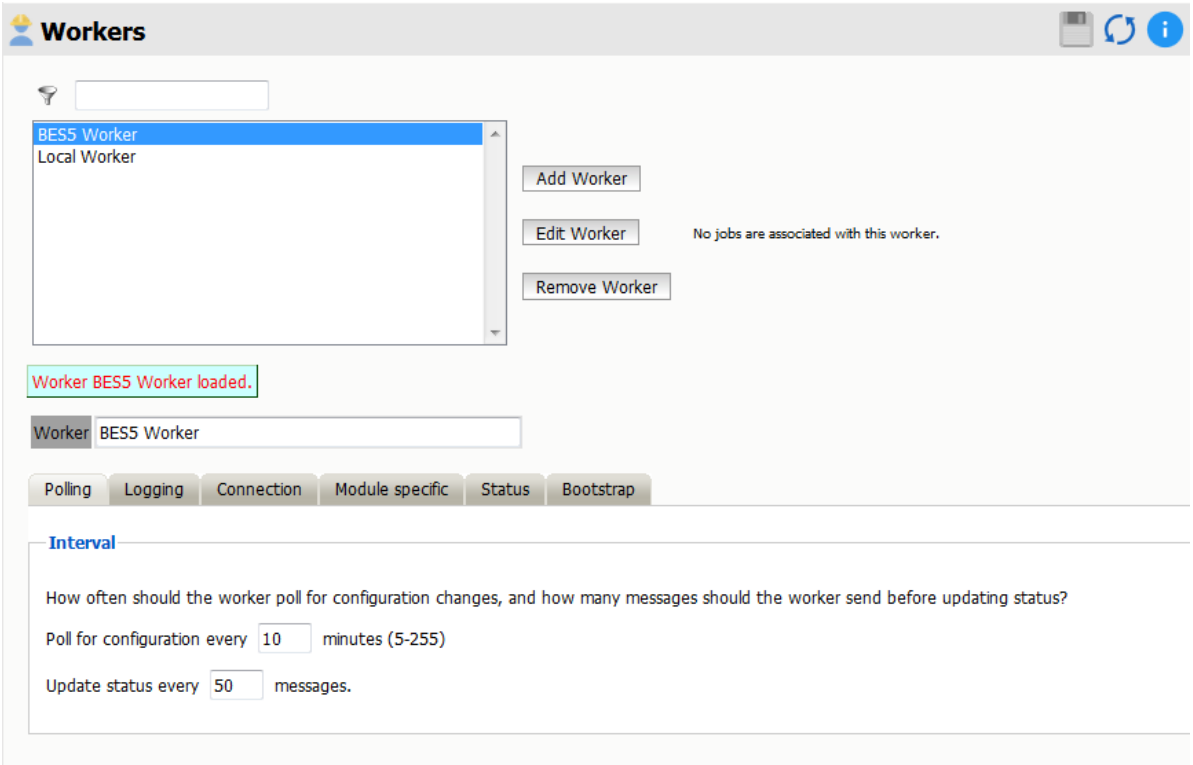
- Polling: Defines how often the worker polls for configuration changes. (Jobs, job changes, etc.)
  - Default is 10 minutes.
  - Default is to update the displayed status every 500 messages while running the job.
- Logging – Controls logging by the Retain Worker.
  - oDefault is normal.
- Connection. Important. This is the information the worker needs to connect to the Server. Double check this setting.
  - Protocol
  - URL/host name
  - password
  - If the connection information is changed, you will need to manually copy the Retain Worker.cfg bootstrap file over to the Worker
  - Correcting or recopying the Bootstrap:
    1. Stop Tomcat.

2. First delete the Retain Worker.cfg file on the Worker machine. For security reasons, there is no way to do this except by manually deleting the file from the server console itself. The path is: Retain Worker/WEB-INF/cfg. Removing the bootstrap resets the worker to initial configuration.
3. Restart and re-upload the bootstrap to the worker.

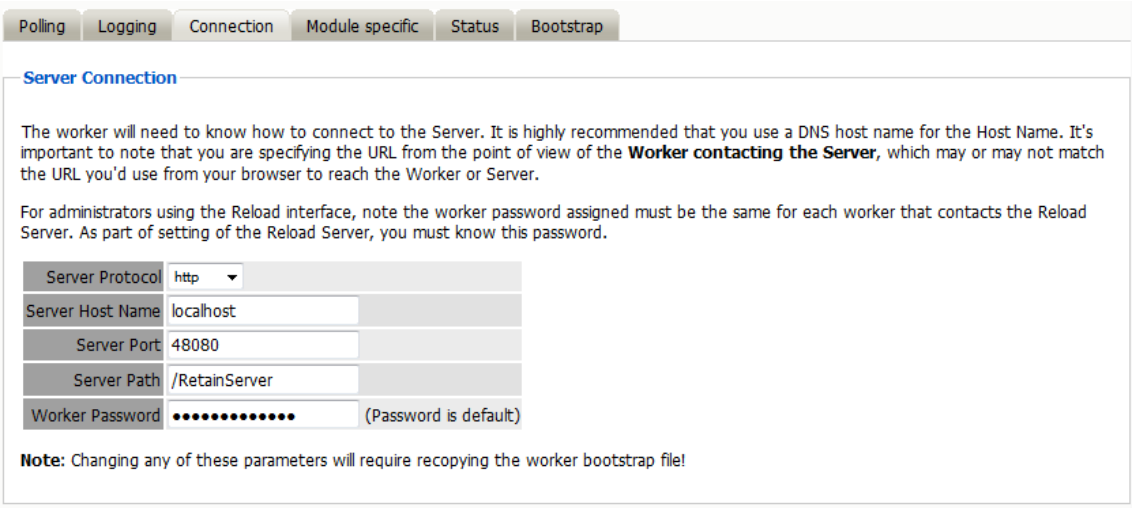
- Module specific – Depends on the module(s) installed and allows the worker to be pointed at a particular post office server.
- Status – Basic status information on the Worker. This status information is updated while the worker is running according to the “Update Status” interval described above.
- Bootstrap – Indicates the download link and when it was last downloaded.

4.24.2 Worker Configuration

From the Retain home page, click Workers. You will see this screen:



1. Click “Add Worker”.
2. Put the name of the worker in the “Worker box”.
3. Example: Linux Worker 1.
4. Check the Connections Tab to be sure the Retain Server IP address settings are correct. This tells the Worker where the Retain Server is located. (Server Host Name or IP address must be correct; this should be populated automatically, but double check for accuracy.) The Server Connection setting must be able to resolve to the Retain Server if placed in a browser address bar.



The Connection tab holds very important contact information for the worker's connection to the Retain Server and the Post Office. Make sure that the Retain Server connection address is correct, or the Retain Worker will not be able to talk to the server.

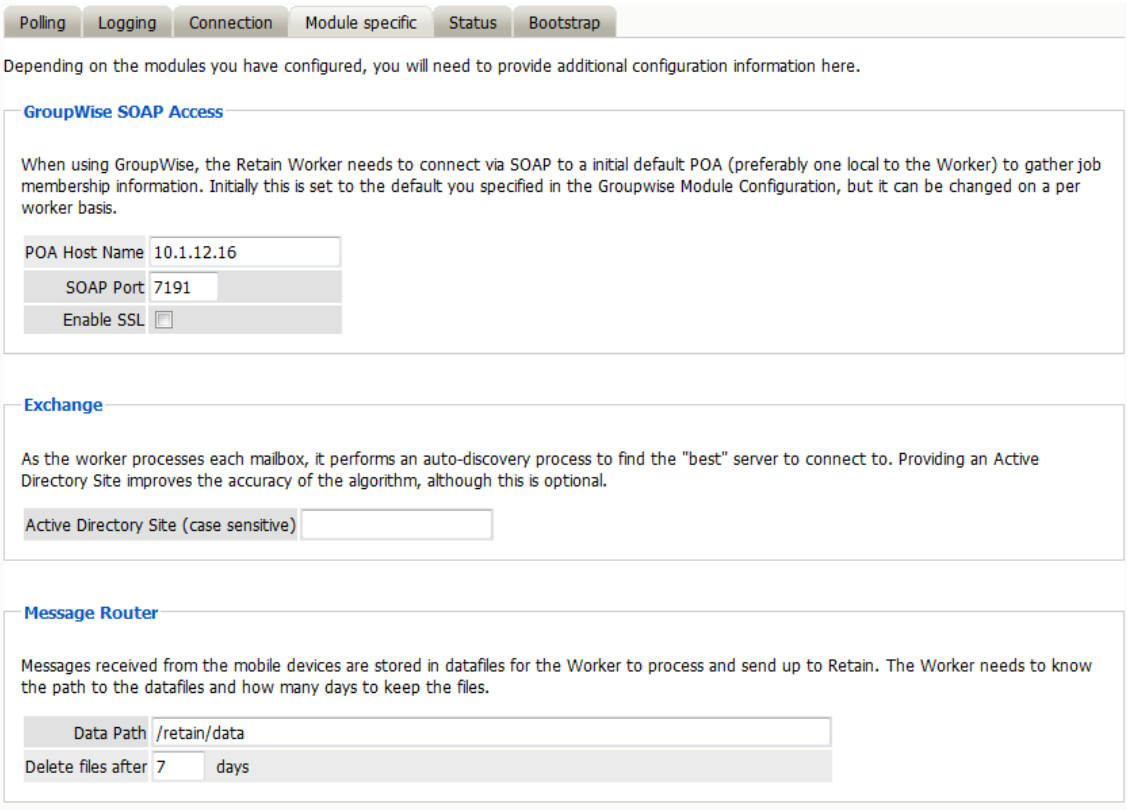
Double check all these settings to ensure that you will have the best performance and connection for your system.

Important Note: It is vitally important that the connection information is correct before you save the worker. On a single server system, it is easy enough – it's the same server. However, in an enterprise set-ting where you might be using different components on different servers, it is extremely important that each component knows how to connect to the Retain Server.

If you do not set this correctly, the Retain Worker will not be able to contact the Retain Server and it will not run jobs.

5. The Module Specific tab contains information specific to individually configured modules. With Exchange, you may optionally specify a specific Active Directory Site the worker is located in. That allows Retain to find the "closest" Exchange Server that is appropriate for the user to be archived. The settings specified here are tied only to the specific worker configured, and does not affect the rest of the system. In general this setting may be ignored unless required. If required, input the appropriate information.
6. Message Router. The message router settings must be set on the message router worker to reflect where the data path is for the Message Router. This path must be specified from the worker's point of view. If the worker is not local to the router, then the path must be fully accessible to the worker; direct file access is required. You may specify how long to keep files in this location. Default is shown.

NOTE: If REST is being utilized, do not configure the Message Router section. The Message Router will default to REST if this is not configured.

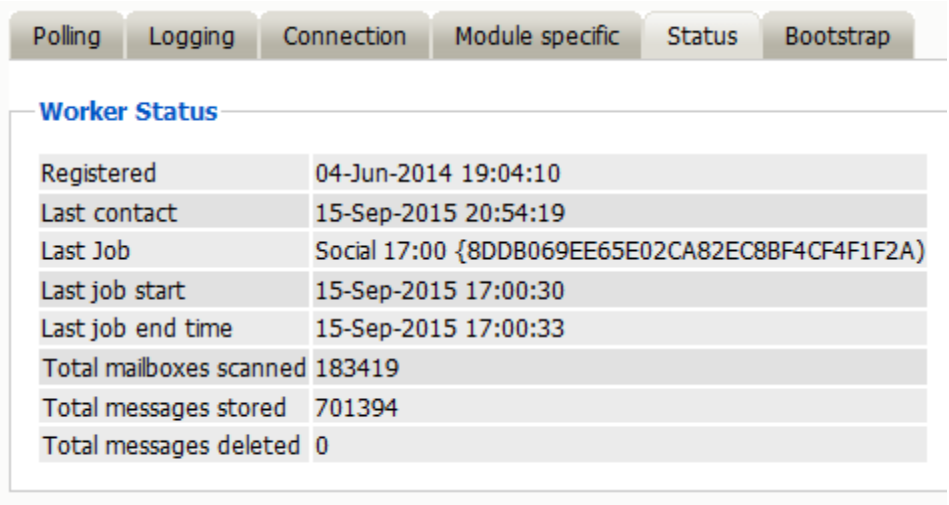


7. After you have checked and set your settings, click “Save Changes” before browsing to another page to finish creating and to save your worker.

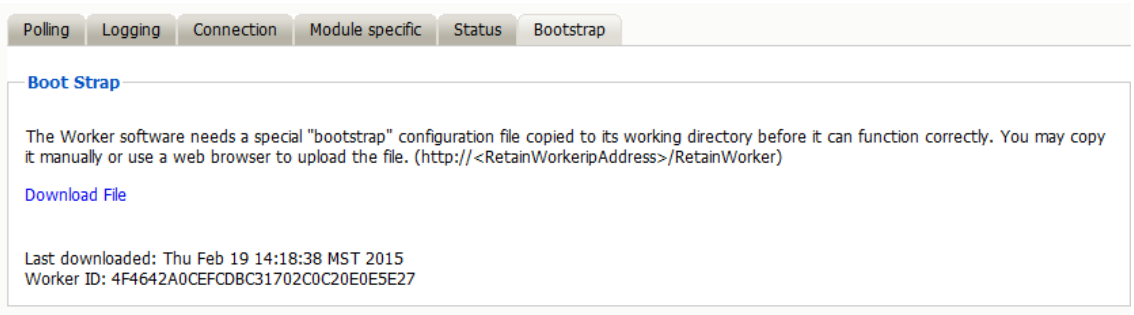
   After the worker has been saved, two additional tabs become visible for settings on the worker. You MUST save changes on a newly created worker before the ‘Bootstrap’ tab is available to be used.

   The status tab for the new worker will not hold any pertinent information until a job has been run or the worker has registered to the server and is in communication. Once a job has been completed, information will be displayed on the general status of the selected Worker.

   Check this page after the worker has completed connecting to the Server for information and jobs. After a job has been run, or is running, more pertinent information, such as total mailboxes scanned and messages stored, will be displayed.



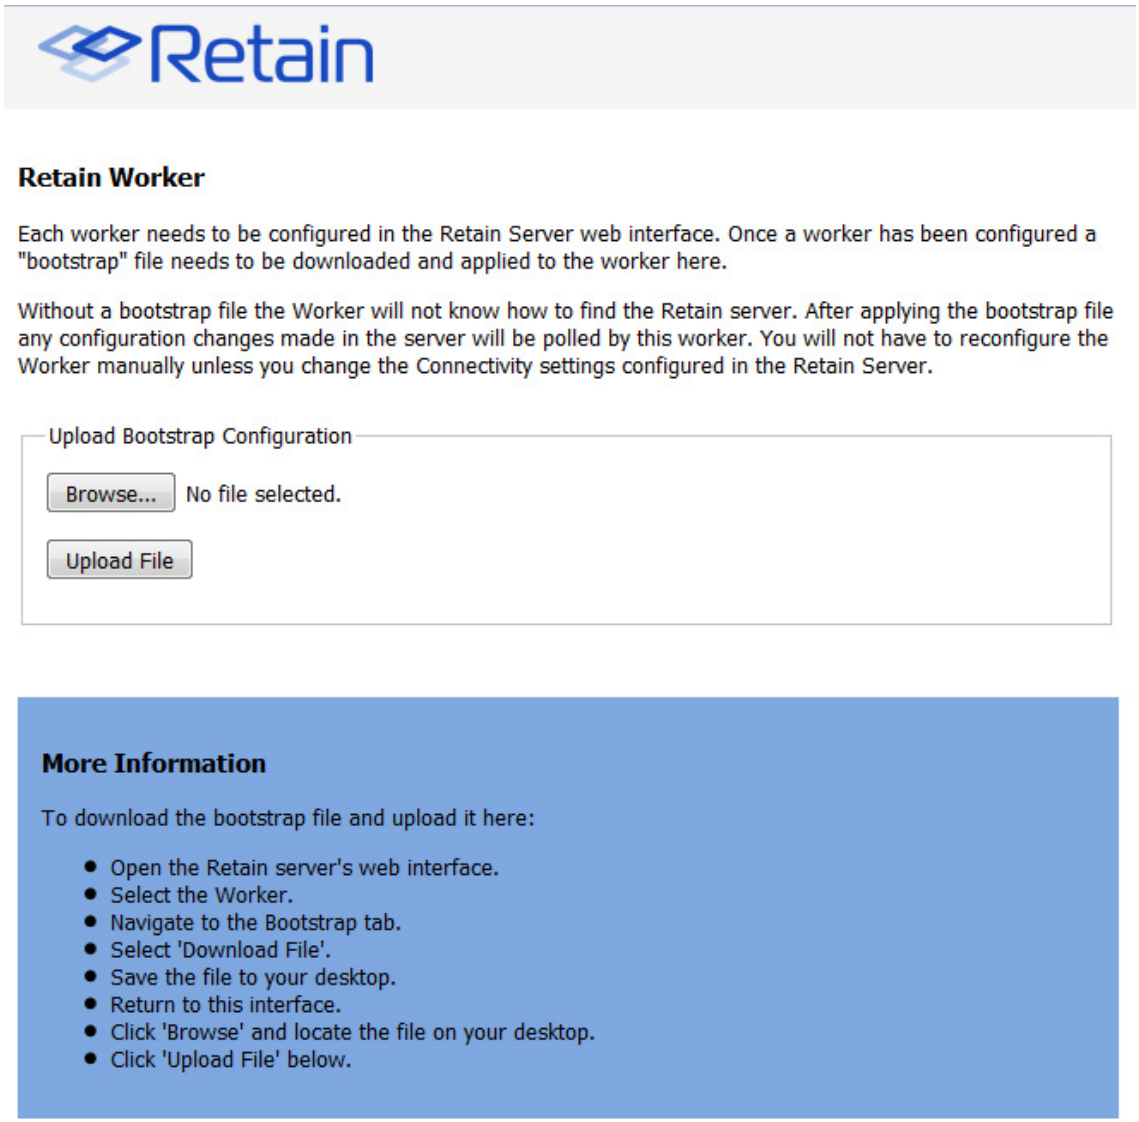
8. After you have saved the changes, the Bootstrap tab appears (see below). Click “Download File” to save the configuration to a local workstation.



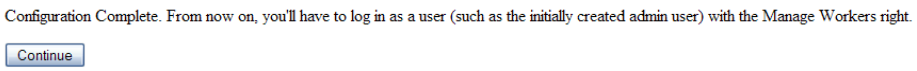
9. Save this file to your hard drive in a place where you can find it. Click “Save”.
10. Open a new browser window and type:

    http://< your-RetainWorker-ip >/RetainWorker

    The upload screen appears when a worker has not been initialized.
11. Upload the bootstrap file.

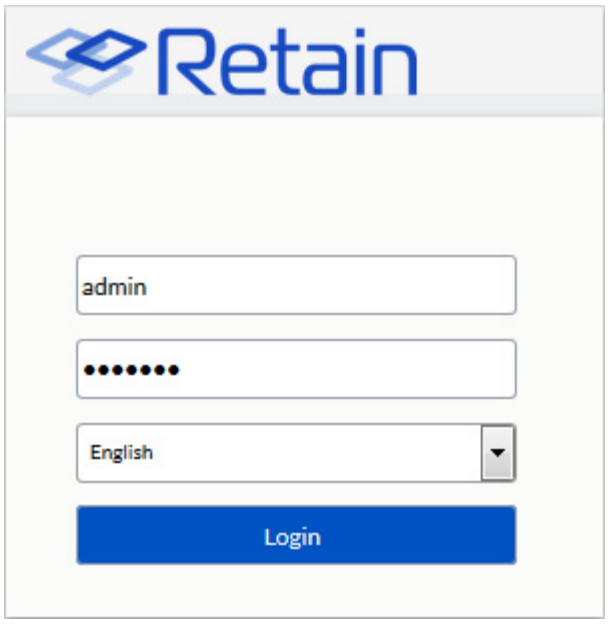


12. Click “Browse” to find the file you just downloaded, and select it.
13. Click “Open”.
14. Click “Upload File”.
15. If the upload went well, you will see this screen:

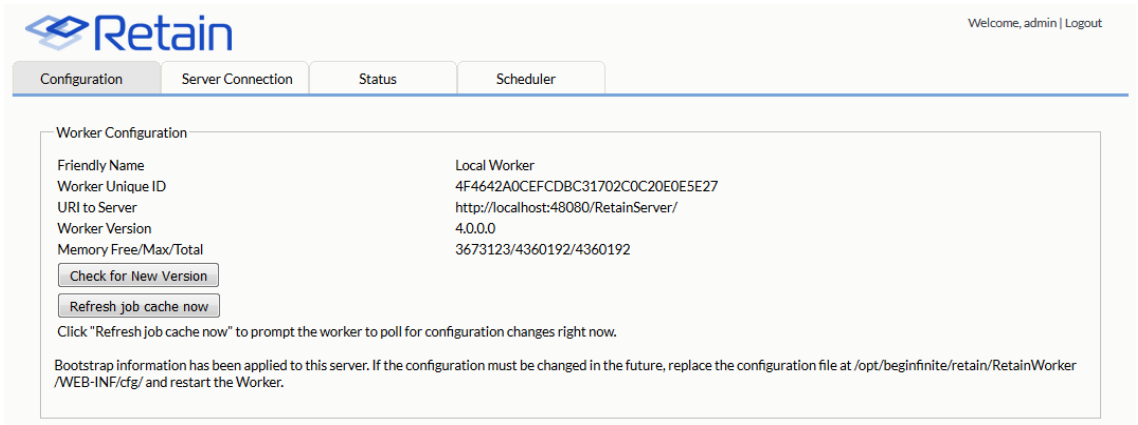


After you click the “Continue” button, you will be returned to the login page.

16. 16.Login as Admin.

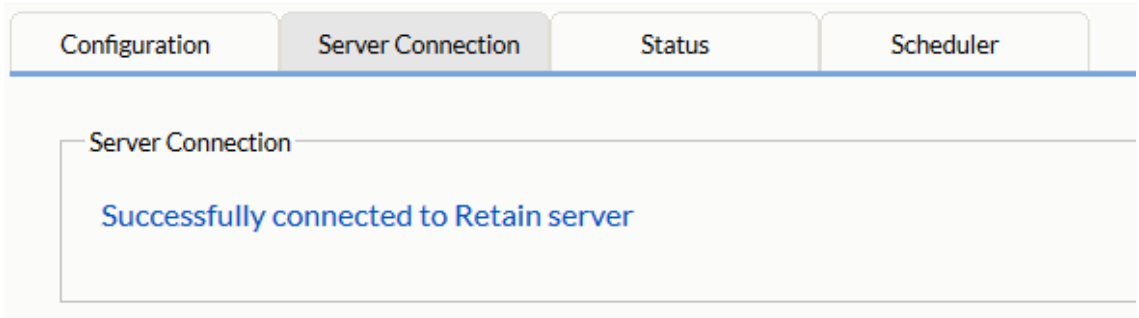


On all subsequent logins, you should be shown the configuration page first.



You should be able to see your configuration on the screen, as shown. The “Check for New Version” button checks for updates. The “Refresh job cache now” button tells the worker to poll the server for configuration changes and new jobs. After this button is selected, the Worker page will refresh automatically after the Retain Server has been polled. It takes just a few seconds.

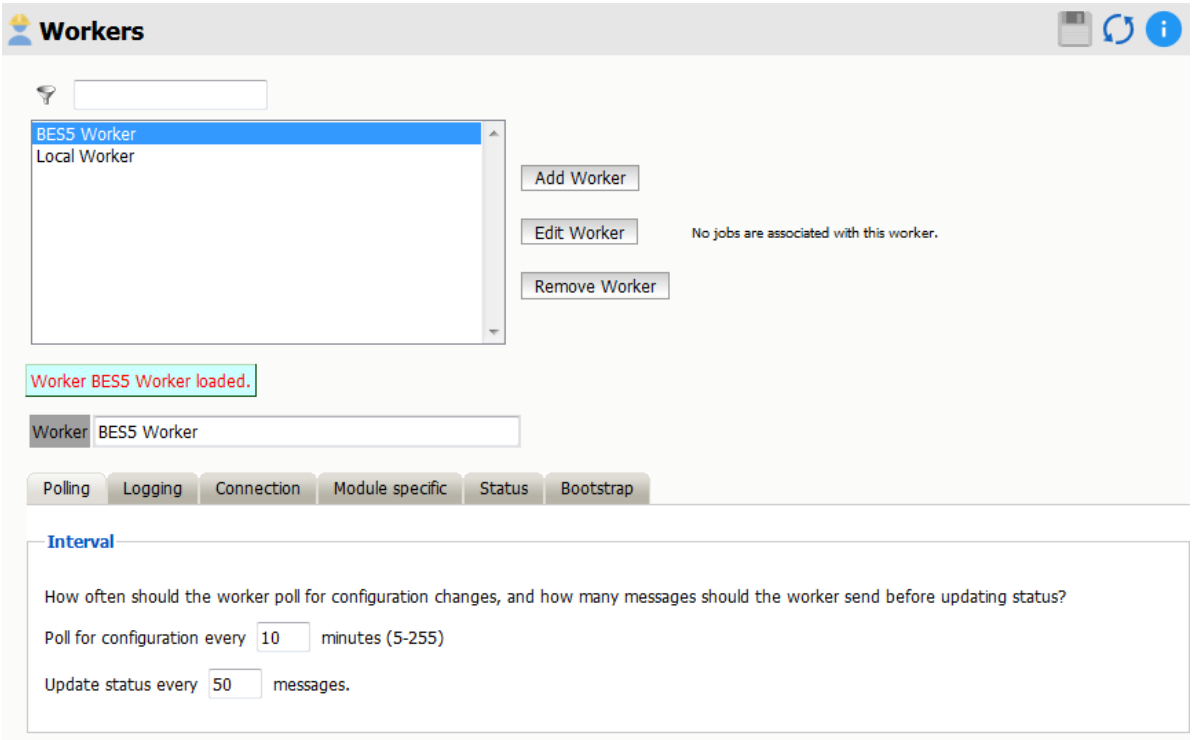
17. Verify the Retain Server connection success by selecting the Server Connection tab.



18. Go back to the Retain Server.

19. Click “Workers”.
20. Click on the worker you just created.

You should see the following:



Pay close attention to “Poll for configuration every ‘n’ minutes. This is how often Retain will check for new jobs.

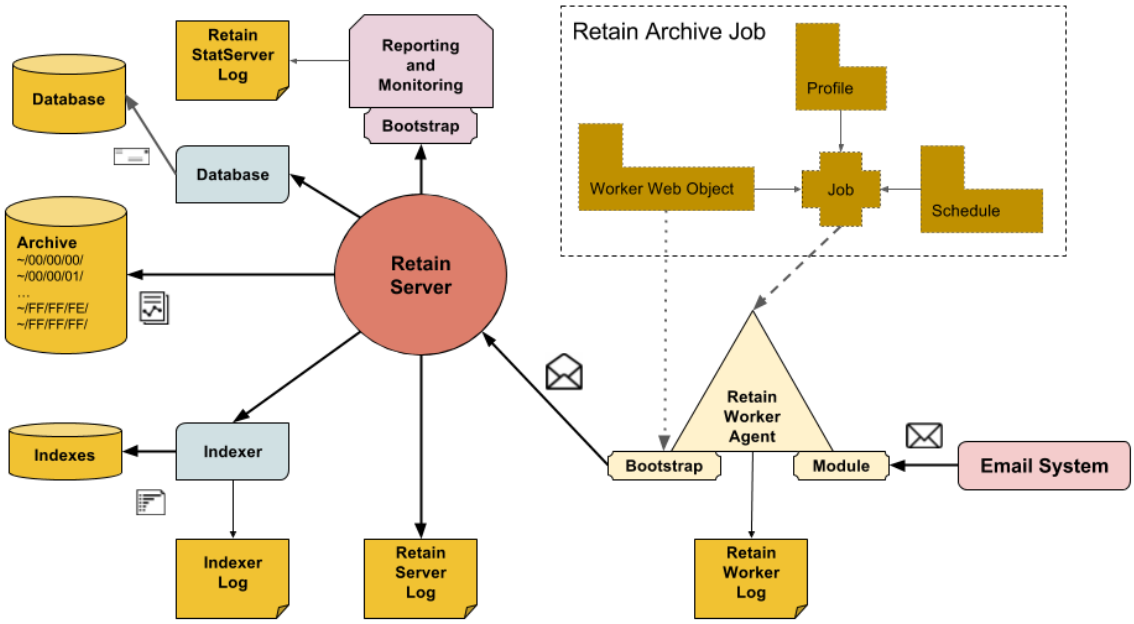
Also, status updates are displayed on the admin page every ‘n’ messages.

Now, set logging to a level you would like. “Diagnostic” writes extensive amounts of data to the logs. If using this level, only use it until satisfied that the system is fully operational. Please be aware, however, that using this setting will write a huge amount of data in the log files.

The profiles, schedules, and jobs must be setup, if they already have not been, in order to begin actual data collection.

4.25 Jobs

The Jobs section is where the Worker, Profile, and Schedule are all brought together and associated into a comprehensive configuration. This is the final piece that will allow messages to be archived from an email system.

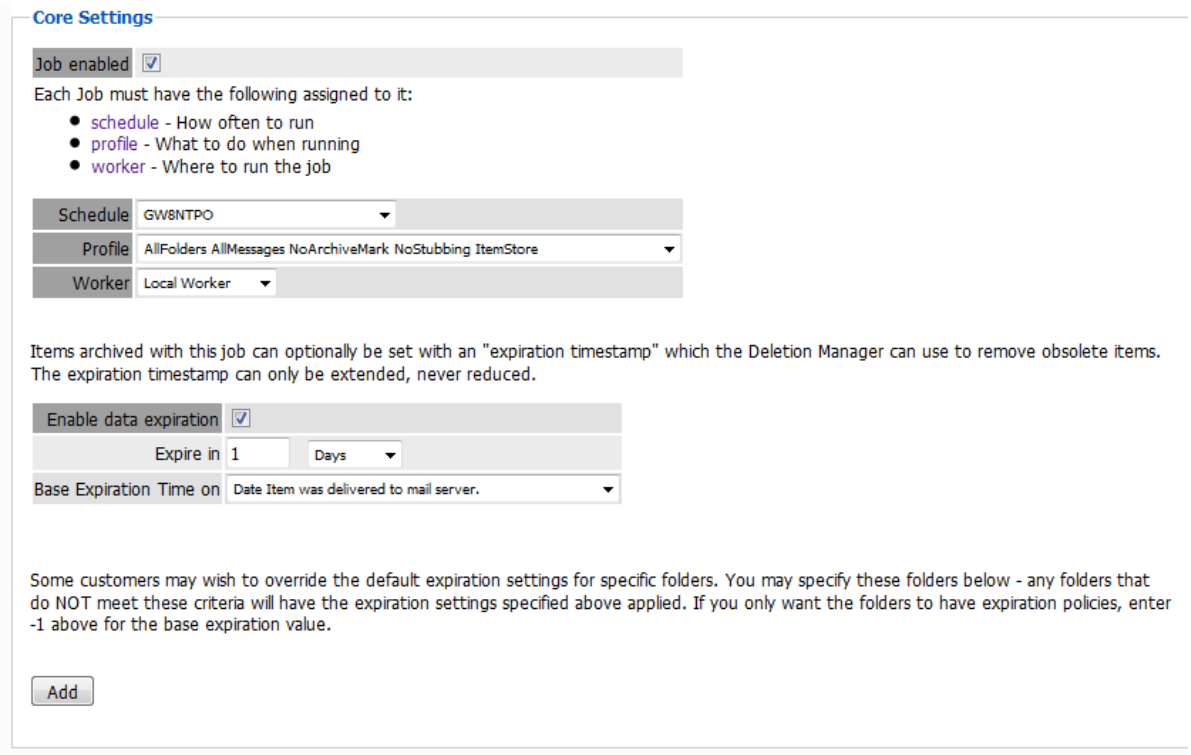


The Worker will not start archiving anything until a created job is bound to a Worker, Schedule, and Profile.

Jobs, like profiles, will not be displayed until the module has been configured. Jobs are module specific. While the majority of the job interface is general to all modules, there are some specific settings.

4.25.1 Core settings

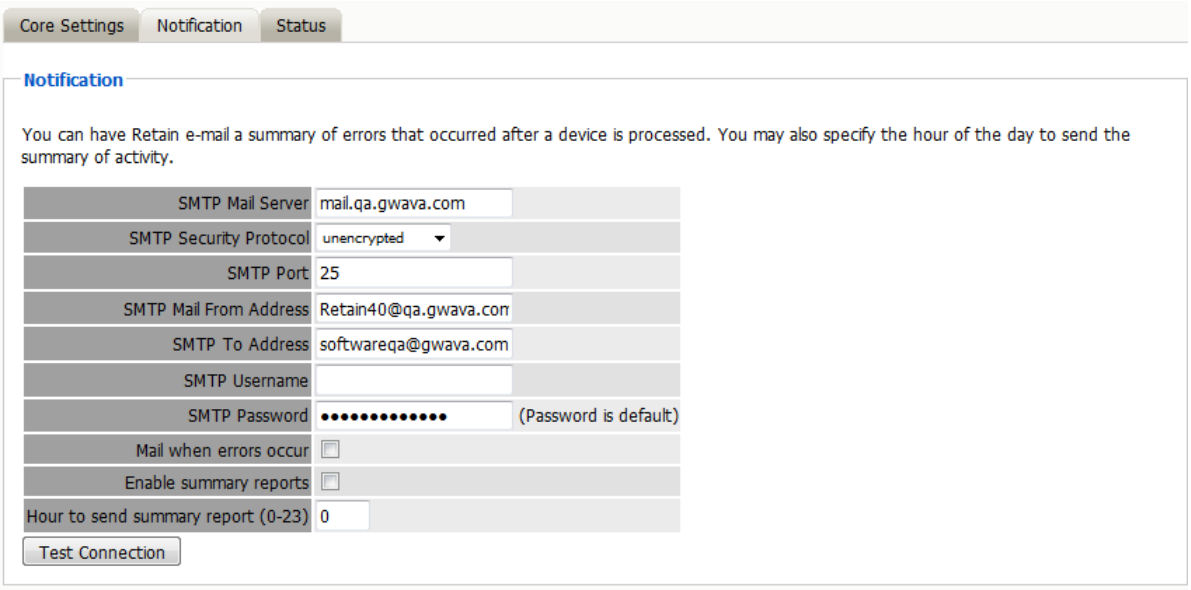
The Core Settings of a job contains configuration which must be set for the job to be saved and become active. A job must be enabled before it will run. Jobs must also have a specified schedule, profile, and worker. These are all selected from drop-down menus, and will not be populated unless those items are already configured in the system.



The Data Expiration setting is an option to place a time stamp on data in the Retain database, which allows for ease of automation for the deletion manager. In addition, devices such as NetApp, Centera, and Hitachi HCAP may use this number to enforce hardware level protection of the stored item so that no one (including Retain) may delete the item before its expiration date. Job Expiration is not retroactive for mail in the database, and only applies to mail archived by the job that it is active for. In order to have messages with custom job or folder expiration dates properly expire, the deletion management date scope must be set to delete messages with an Expiration Date older than 1 day.

4.25.2 Notification

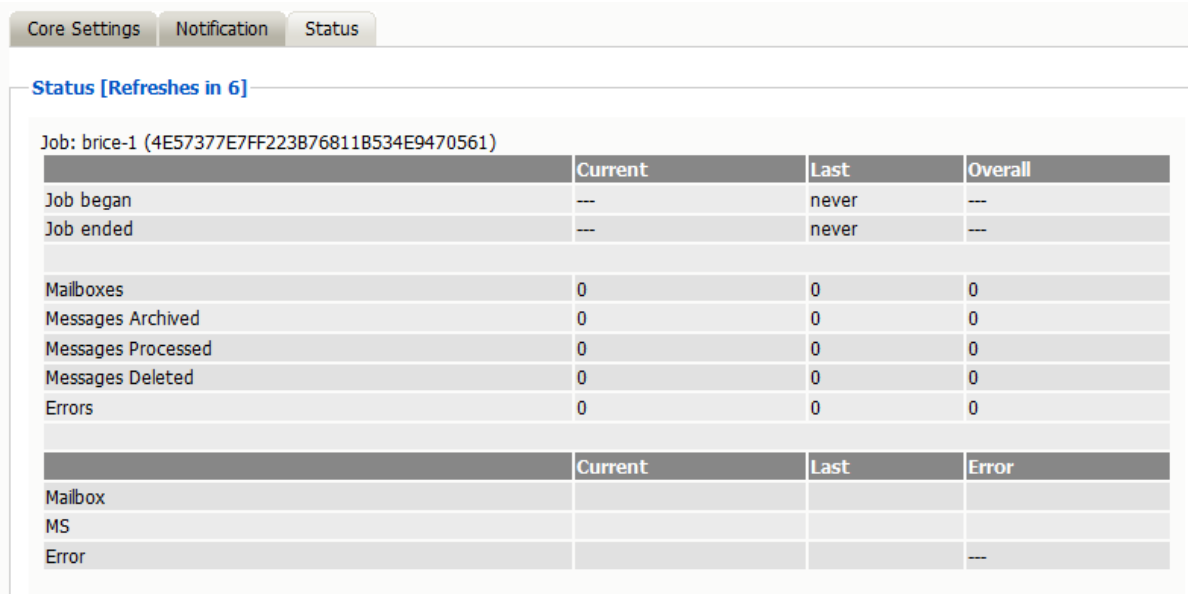
When a job is run, the Notification option allows the administrator to be emailed a summary and report of any errors, for each running job.



For notification to function correctly, the SMTP information for the desired SMTP server must be fully filled-out. How much information is required varies depending on the mail system used.

Status

The Status tab displays the status of any currently running jobs, as well as the stats of the last completed job.



On some modules, currently running jobs may be terminated here. For the rest, this tab is informational only.

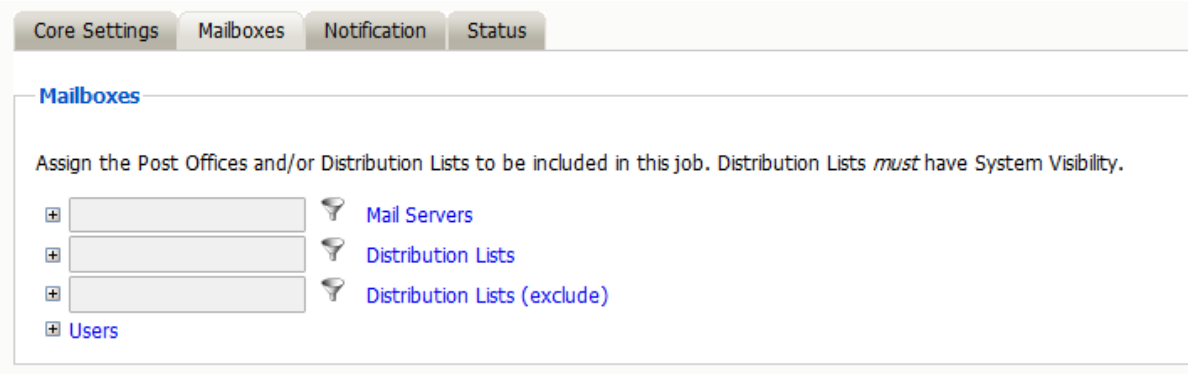
4.25.3 Mailboxes

The mailboxes tab is where the administrator specifies which entities (mail server(s) and/or Distribution List(s)) are to be scanned. This tab is not displayed for the mobile module.

Expand the Post Office and/or Distribution List trees, and check off the items you want to be dredged.

NOTE: If you desire to have a job backup a single user, or selected group of users, select the Users menu and assign the user(s) desired

---



The Distribution List selection allows you to include or exclude a group of users from an archive job. If you want to use GroupWise Distribution Lists, the visibility needs to be set to “system wide”.

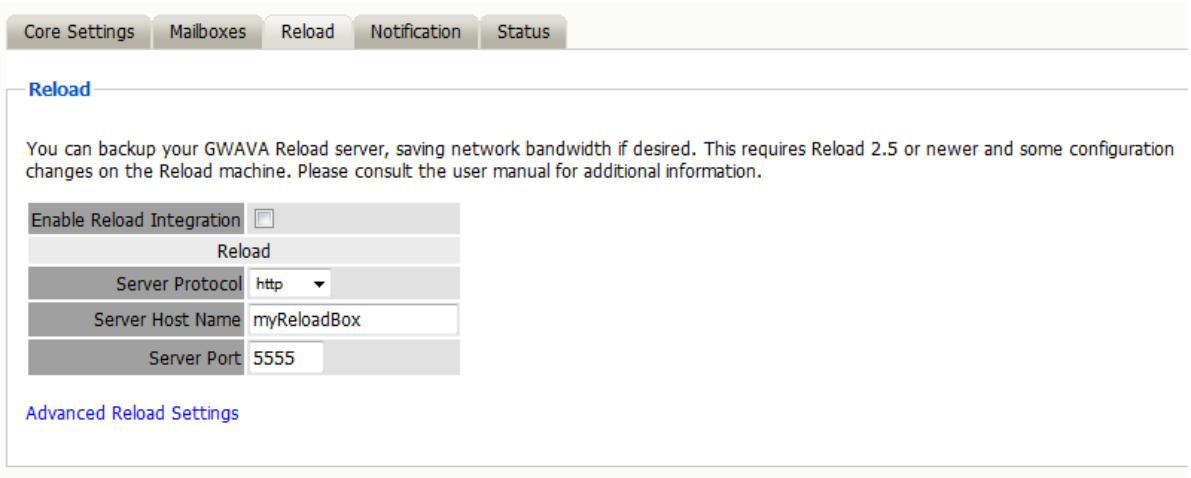
The users section allows you to select individual users to include, or exclude them from an archive job. For example: you can select an entire Mail Server to be archived, and then expand the users section to include or exclude users to the job.

This can also be used to select only certain users in the system for an archive job.

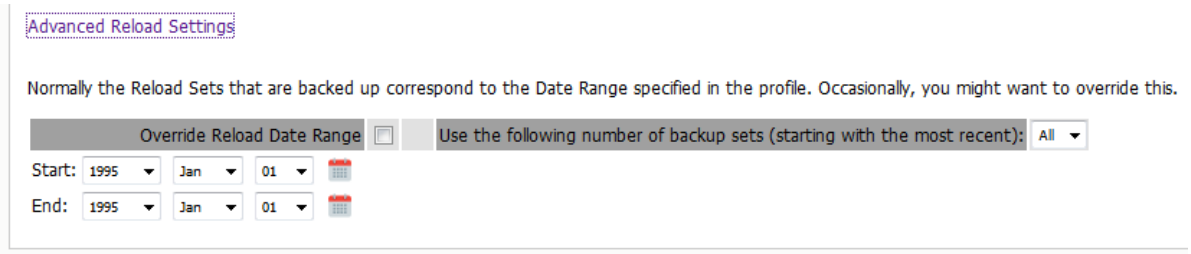
To add a user to the Include or Exclude list, select the respective ‘Add user’ button and search for the user. It can be helpful to unselect the ‘only show recently cached items’ option. Add the selected users to the list in the search window, then select ‘Ok’ to add them to the include or exclude list.

4.25.4 GroupWise Job option - Reload

Selecting this option tells the job to use the Reload integration for systems utilizing Reload. See the Reload Integration section for instructions on configuration. The Server URL must be edited with the cor-rect contact info for the reload server. The default Reload port is 5555.

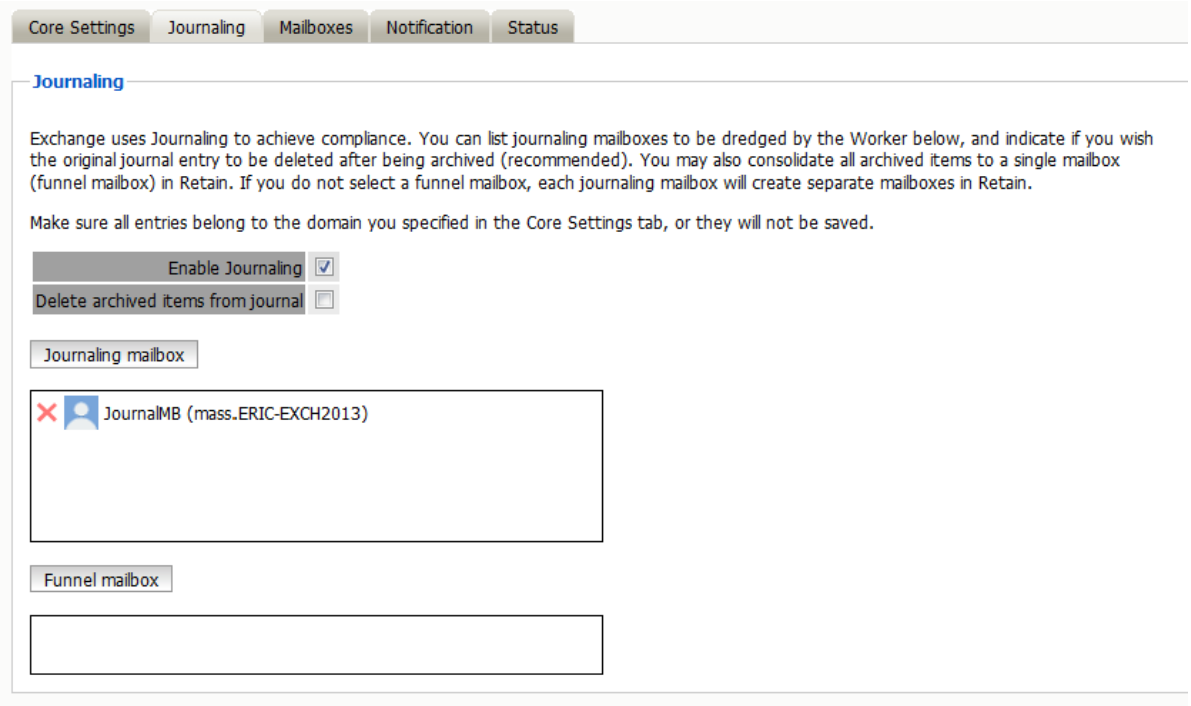


Retain instructs Reload to load the corresponding backup set for each date range in the profile. If a differ-ent set needs to be specified, the override setting specifies a date range specifically for Reload, which then will ignore the date range in the profile. Retain will still archive according to the profile date range or settings.



4.25.5 Exchange Job Option - Journaling

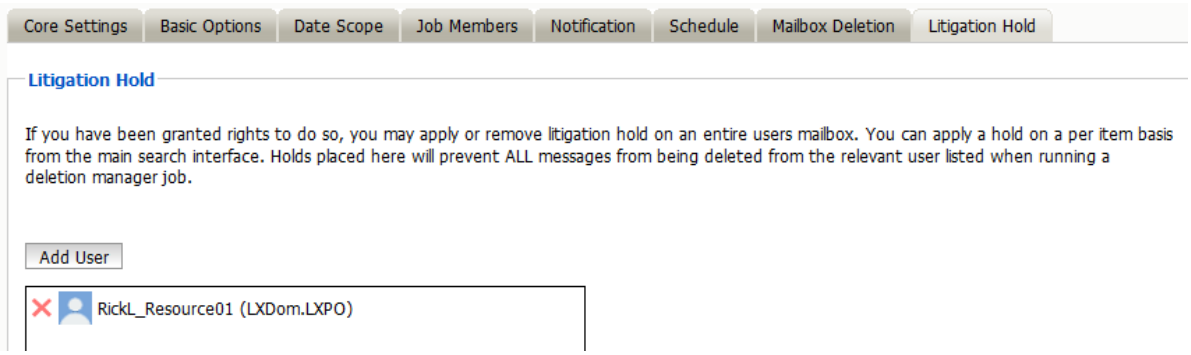
In order to achieve compliance, Exchange utilizes a Journaling mailbox. This mailbox can be set to be archived by retain to collect all messages on the system. The Journaling mailbox can rapidly grow in size if it is not cleaned out after messages have been archived. The Journaling option for Exchange jobs allows Administrators to set whether Retain will automatically clean out messages from the Journaling mailbox which have been archived.



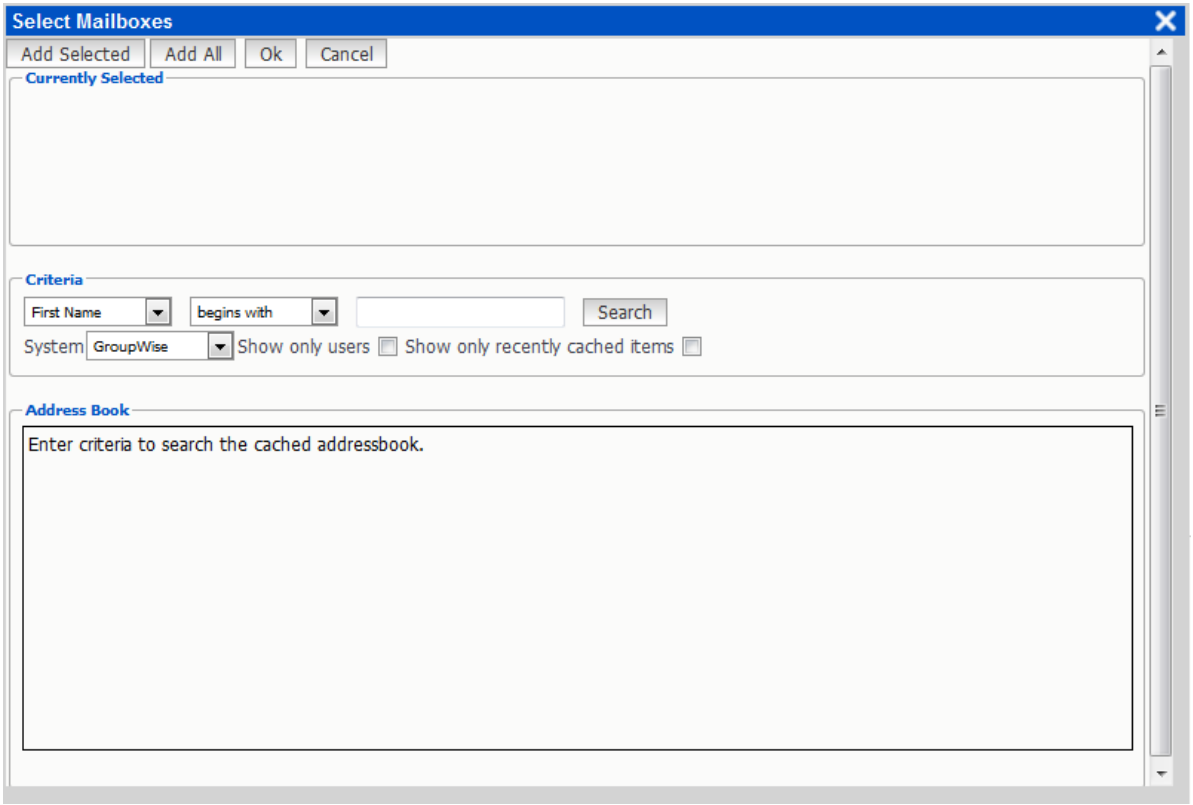
On larger systems where there are multiple journaling mailboxes, Retain will automatically create a mail-box for each of them in the archive. However, if desired, the journaling mailboxes may be all archived to the same specified mailbox in the archive. This is the funnel mailbox. If desired, specify the mailbox by selecting the ‘funnel mailbox’ button, search for and select the desired mailbox. Only existing mailboxes in the Retain system may be specified as a funnel mailbox.

4.26 Litigation Hold

The Litigation Hold tab provides the ability to exclude any specified user’s data from any deletion job, preventing any of their data from being deleted when the job runs.



To add a user to the litigation hold list, select the ‘Add User’ button to open the ‘Select Mailbox’ window.



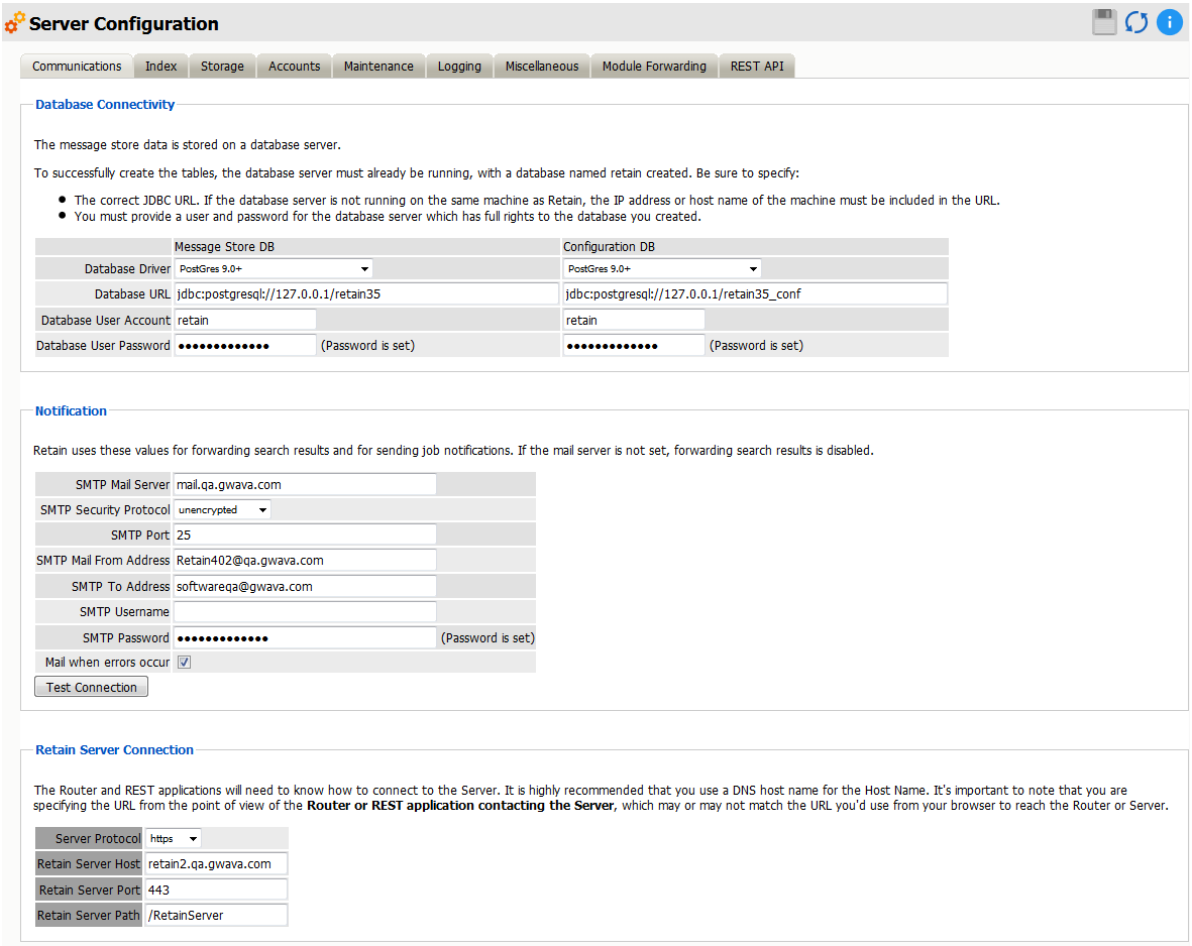
Select the source system for the user and enter search criteria. After searching, select the desired user or users and select the ‘Ok’ button to add them to the list. Save all changes.

5 Configuration

5.1 Configuration

5.1.1 Server Configuration

The overall operation of Retain is configured in this section.



Each tab will be explained below.

Along the top, you will see configuration options for the Retain server and workers known to this server.

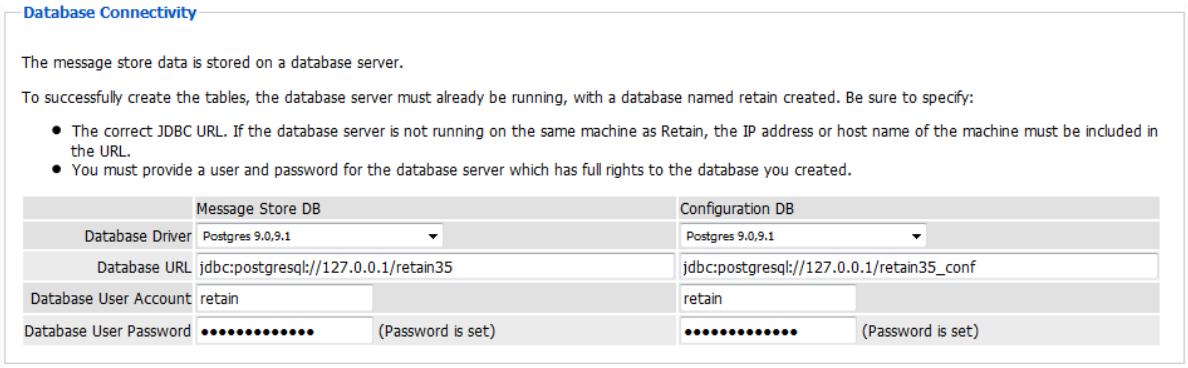
The most important topic is communications. These are the settings you have set when you initially installed Retain. You may make changes here.

After changing any communications options, it is strongly recommended that Tomcat be restarted imme-diately.

5.1.2 Communications tab

The settings for the Communications tab are set in the initial setup of Retain. These include the connec-tion to the SQL Database server, the REST and Router connection information, and the notification or SMTP information. If any changes in the system or corrections are needed, they should be configured here.

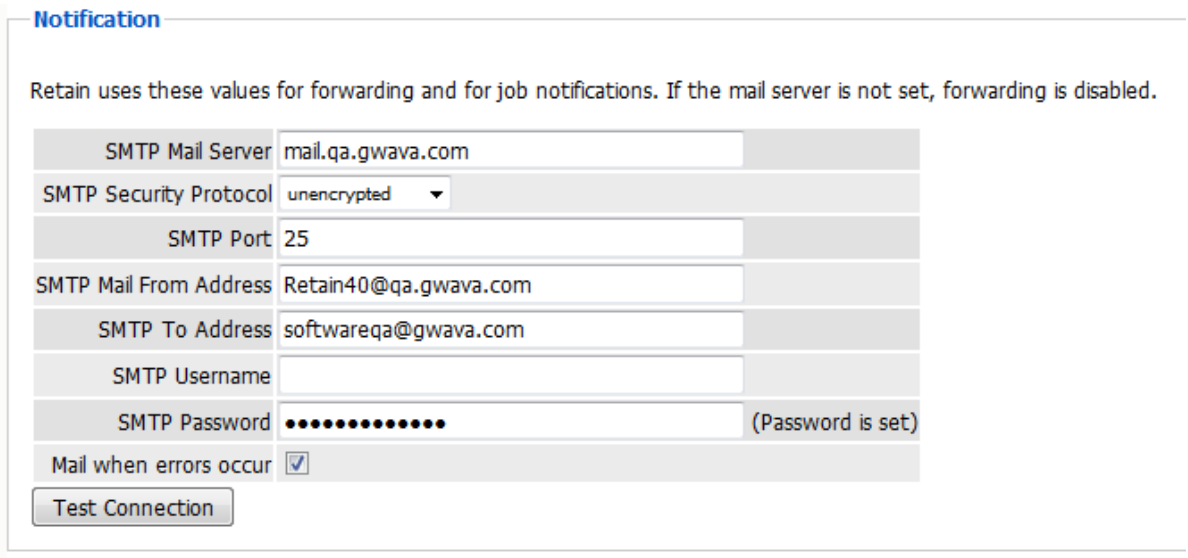
5.1.3 Database Connectivity



The Database Connectivity tab contains the connection information for both the configuration and stor-age databases. This information should never change unless the database server is being migrated to a new location.

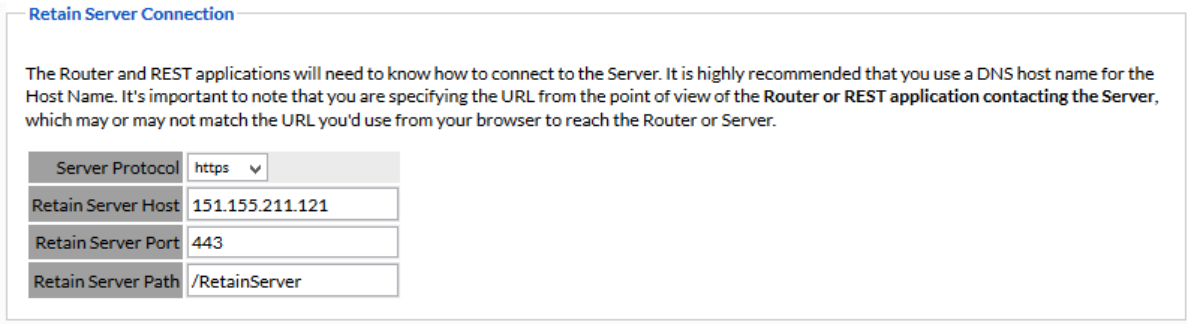
5.1.4 Notification

In this location, the admin user specifies what address receives notifications about the general system, as well as what SMTP Gateway to utilize to send these notifications. This is also used for forwarding messages from Retain to the SMTP Mail Server for transmission to the recipients.



5.1.5 Retain Server Connection

This location dictates the connection address, port, and protocol to which the Retain Server will listen for communications from the Router and from any other REST applications. The communication settings must be correct from the point of view of the applications trying to use it. The hostname or IP address must successfully connect to the Retain Server.

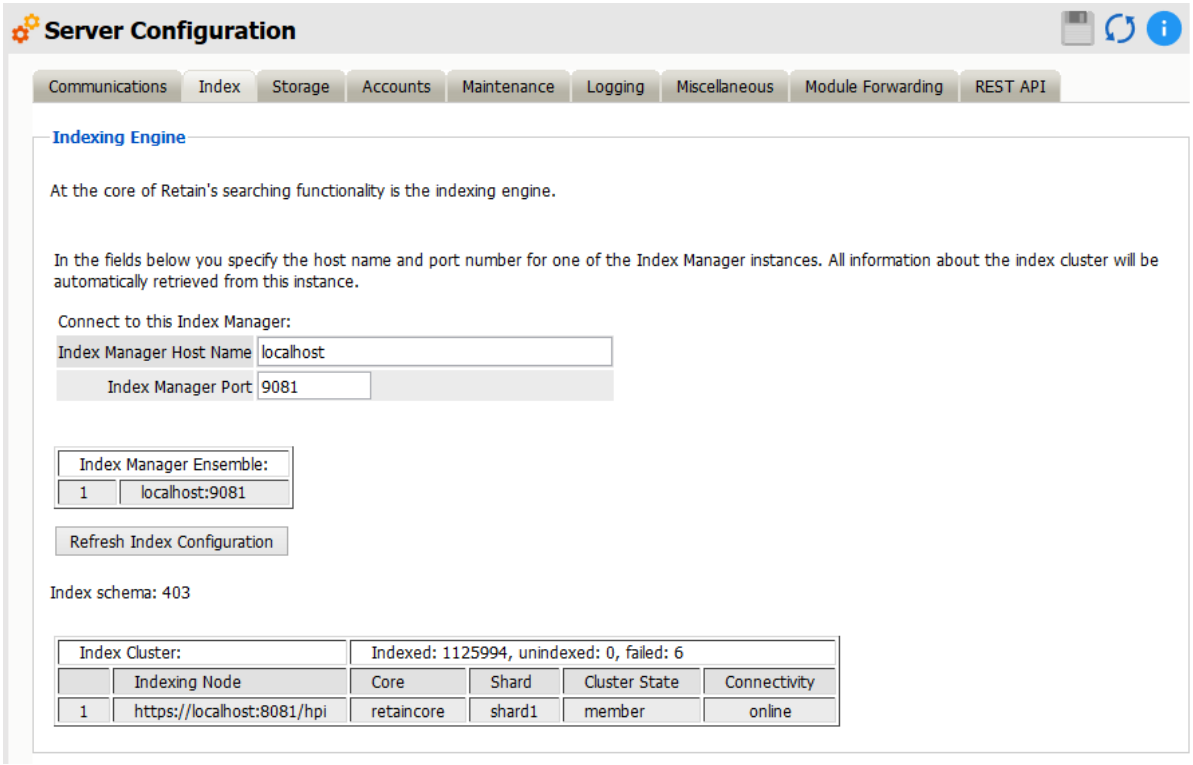


5.1.6 Index tab

The Index tab allows the management of the indexing engine as well as the ability to decide what kinds of attachments are indexed and what size.

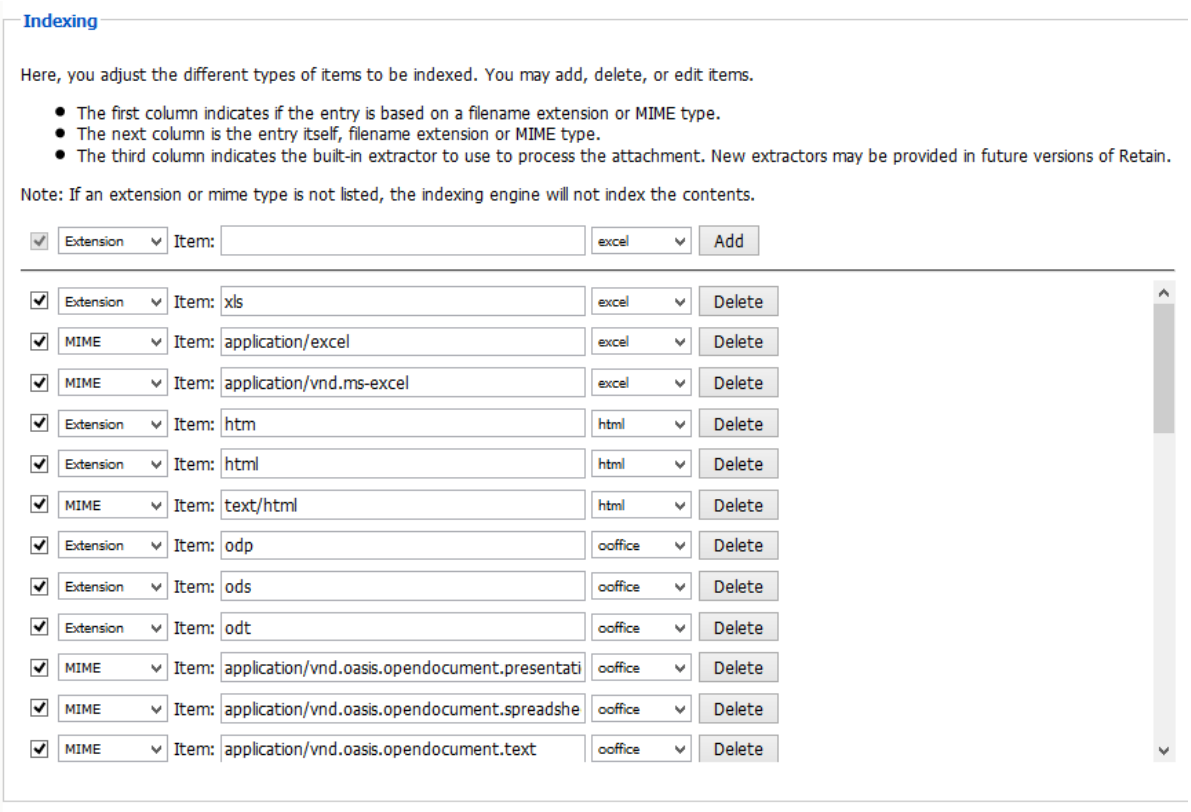
Indexing Engine

Retain utilizes two index engine choices. These choices are set at install or upgrade. The index engine is the same between the Standard High Performance Index and the High Availability Index options, (HAI). However, while the Stan-dard engine is nested with the Retain Server, the HAI is a cluster of index engines, managers, and replicas which provide scalability and disaster insurance. The standard index engine will perform better than the HAI under all but heaviest loads. (Index engine speed is limited by disk IO for the Standard, and Disk IO and network speed. The trade off of performance and bottleneck moved in the favor of HAI when the system is very large and under large demand.



Indexing

You can control what Retain indexes here. You may add as many items as you wish to the list of attach-ment types to index. Note the explanation at the top of the table. The items are listed (in order) by type, extension, archived form (extractor used), and maximum stream size and file size. Choose whether to index the attachment based on its filename extension or its MIME type (the content itself). You also choose which extractor to use to index the attachment. Select as many as you need. If an attachment type is common in the system and needs to be indexed, but is not already indexed, it may be added by using the ‘add’ row.



Rebuilding Indexes

Under rare circumstances - and usually due to hardware failure/issues - the indexes used by Retain when the search feature is invoked become damaged or files become missing. Resetting the index state of messages in the Retain database causes the Retain indexer to re-index the affected messages.

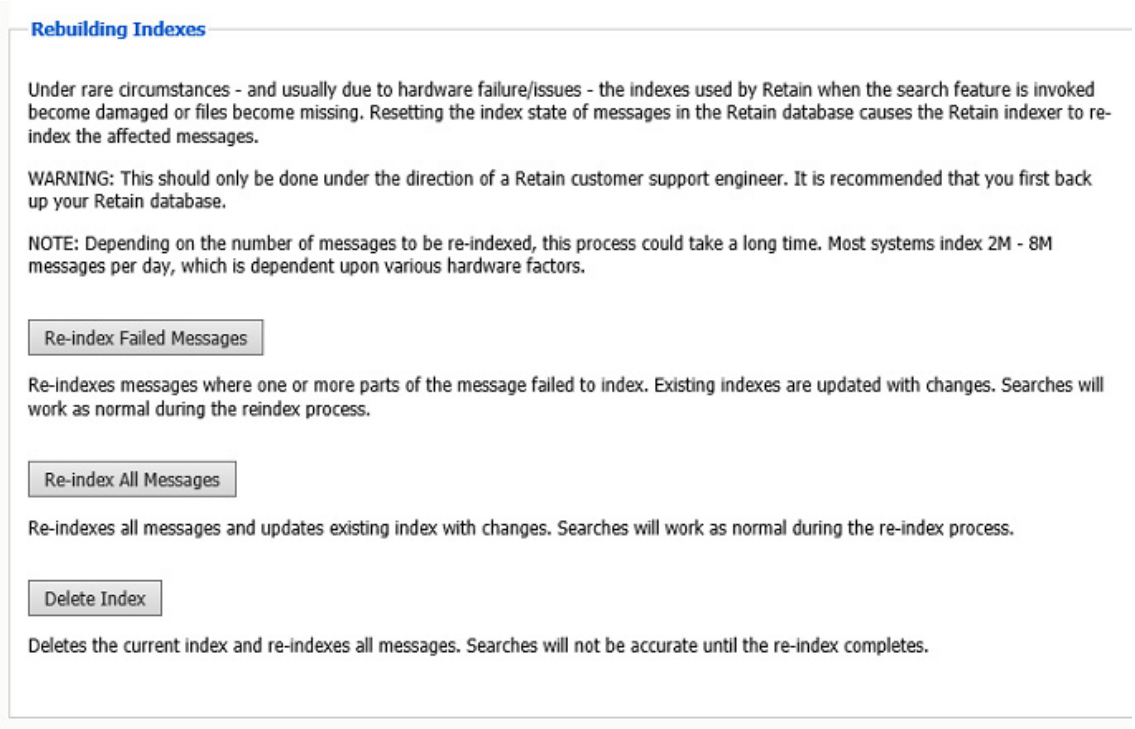
WARNING: This should only be done under the direction of a Retain customer support engineer. It is recommended that you first back up your Retain database. Rebuilding indexes will use significant amounts of CPU, RAM, disk IO and time.

NOTE: Depending on the number of messages to be re-indexed, this process could take a long time. Most systems index 2M - 8M messages per day, which is dependent upon various hardware factors.

There are 3 options:

1. Re-index Failed Messages: Re-indexes messages where one or more parts of the message failed to index. Existing indexes are updated with changes. Searches will work as normal during the re-index process.
2. Re-index All Messages: Re-indexes all messages and updates existing index with changes. Searches will work as normal during the re-index process.

3. Delete Index: Deletes the current index and re-indexes all messages. Searches will not be accurate until the re-index completes.



Index performance can be increased by the number of threads allowed to the indexer. Again this will use significant amounts of CPU, RAM and disk IO and any one of these can limit indexing and Web UI performance.

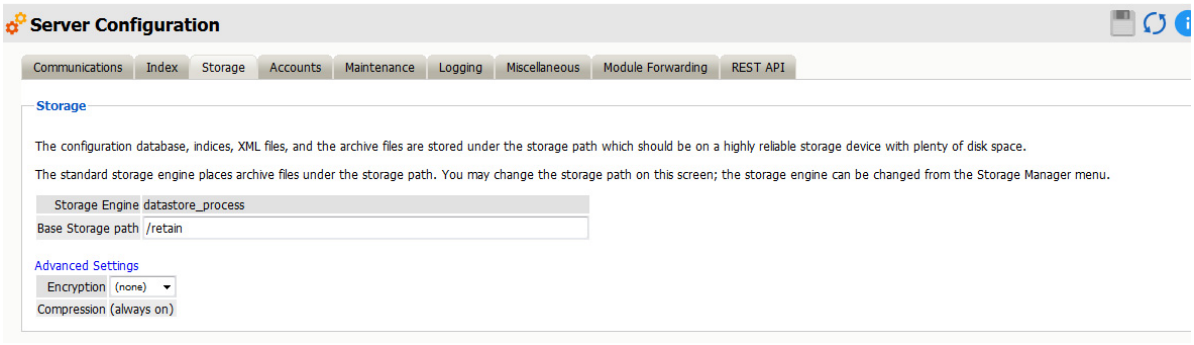
To increase indexing threads:

1. Stop tomcat
2. Edit ../RetainServer/WEB-INF/classes/config/solrcloud.indexing.properties
3. At the end of the file, change the setting: *indexingThreads=0* to *indexingThreads=[1 - 24]* (You should not go higher than *10* times the number of CPU cores; otherwise, the overhead of CPU slicing would actually work against you.)
4. Start tomcat

You will also need to keep in account users needing to access their Retain mailboxes and perform searches. Increasing threads could slow down the web UI response, so you'll need to experiment with what works best for your system.

5.1.7 Storage tab

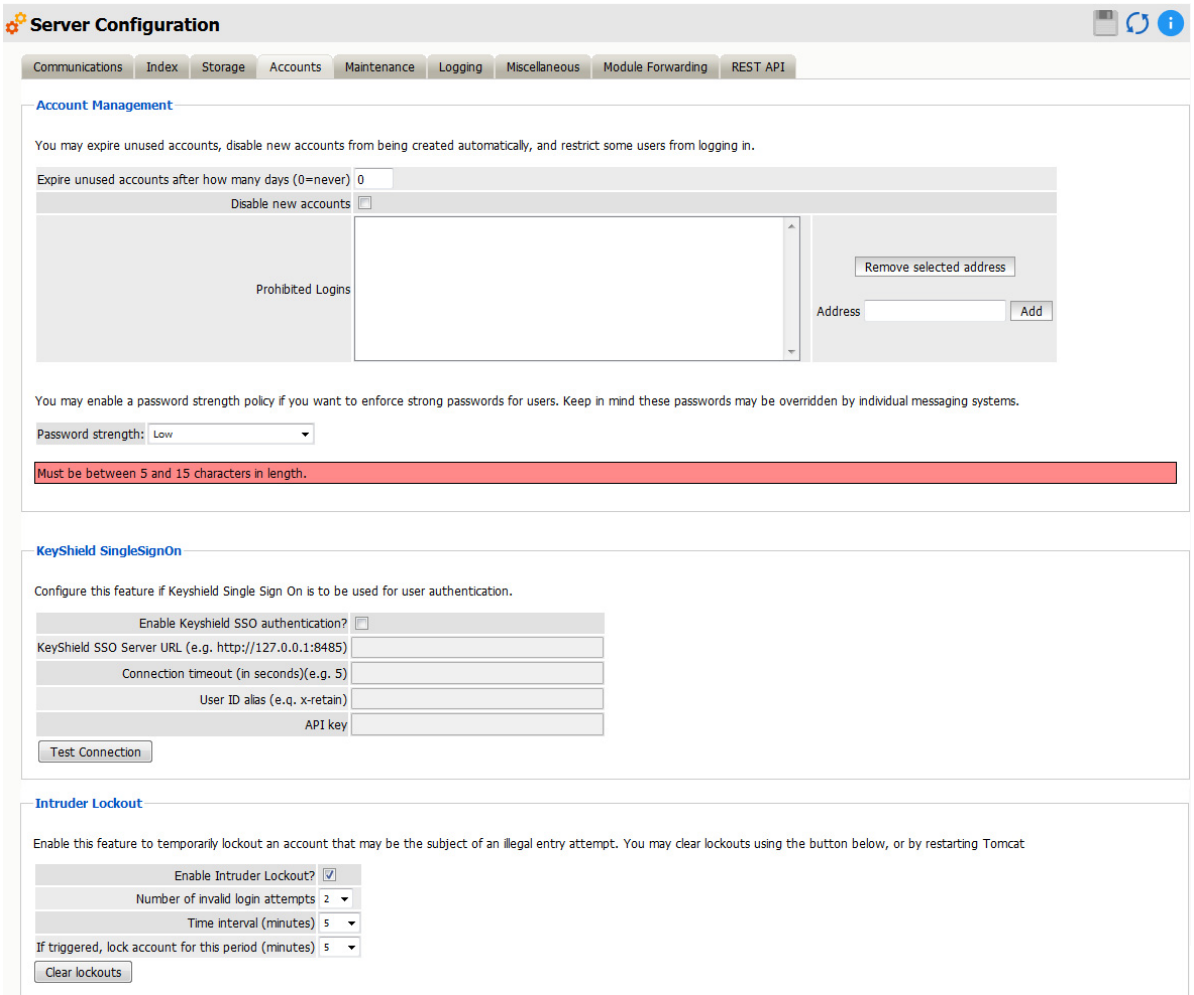
The default storage path for your Retain archives is listed here.



Retain uses a hybrid storage system. All message headers and metadata is stored in the database but all message text and attachments are stored as files on the file system in the location listed here (and subdi-rectories). This is done a little like GroupWise, where the OFFILES directory stores Binary Large Objects (anything larger than 2KB). In the case of Retain, this is transparent to the user and you may choose whether to encrypt this data. If it is stored in a location accessible ONLY to the Retain Server, then encryption is not strictly necessary.

5.1.8 Accounts tab

The accounts tab allows you to control the creation of accounts.



Open System vs. Closed System

Normally, Retain lets all mail system users log in. This is considered to be an “open” system. When that happens, Retain will check to see if a Retain account already exists and if not, it will create a new account for them and assign them to the group *default*.

Sometimes, you don’t want certain users to have access to the Retain archives. In this case, you may add these users to the list of Prohibited Logins. You do so by entering their name in the Address field and click “Add”.

To make a “closed” Retain system, simply click on “Disable New Accounts”. If you use this option, it means that you will have to manually create accounts in Retain for authorized users. In other words, the only people who can access your system will be people who you specifically create an account for.

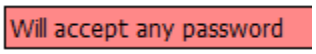
In Retain, user accounts expire after 30 days of inactivity by default. You may choose the number of days or choose 0 for “accounts never expire”.

See “User Rights” for more information.

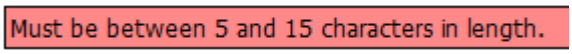
Password Strength

User-created passwords may be controlled for strength. By default, Retain accepts any password set by users. To require a higher security password, select the higher level desired. Requirements for the low, medium, and high settings are defined as:

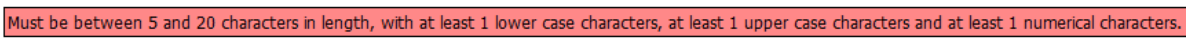
Will accept any password



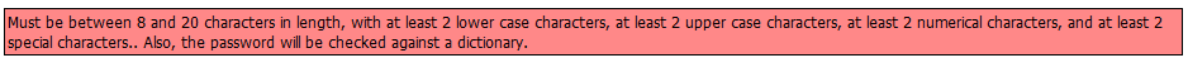
Low: Must be between 5 and 15 characters in length.



Medium: Must be between 5 and 20 characters in length, with at least 1 lower case characters, at least 1 upper case characters and at least 1 numerical characters.

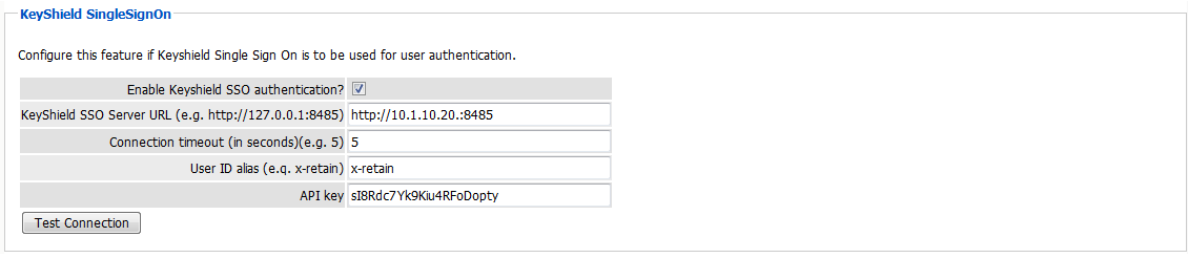


High: Must be between 8 and 20 characters in length, with at least 2 lower case characters, at least 2 upper case characters, at least 2 numerical characters, and at least 2 special characters.. Also, the password will be checked against a dictionary.



KeyShield SingleSignOn

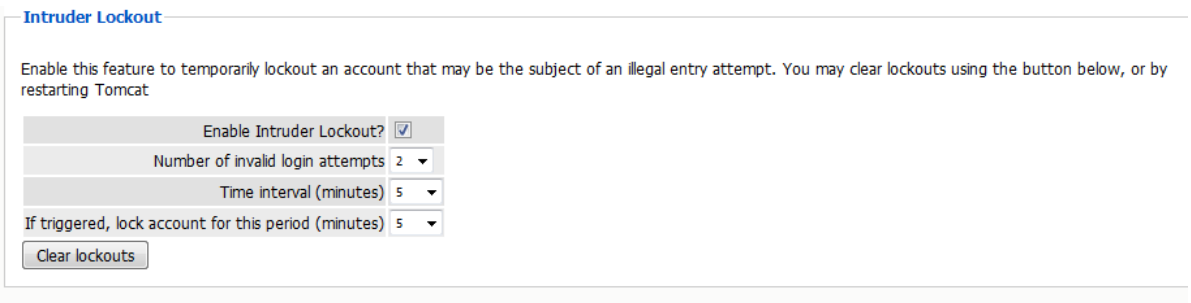
Retain supports the use of KeyShield SSO for users. To use the KeyShield client in coordination with Retain, Retain needs to have an open connection to the KeyShield server, the User ID alias, and the API key. Specify the KeyShield SSO Server URL, Alias, and API key. The Timeout is set in seconds, and may be anything required, 5 is recommended. Test the connection to ensure proper function.



When configured, Retain checks to see if the KeyShield client is running and if the user is currently logged in. If they are logged in, Retain checks the user against the specified KeyShield Server and then either fails authentication and sends users to the login page, or immediately passed them to their interface. The effect is that users who are already logged into the KeyShield client will not be required to login to Retain, but will be immediately taken to their appropriate interface.

Intruder Lockout

Accounts may be locked if multiple failed attempts are detected, according to the specified time window. This is useful to deny password cracking attempts on the server.



To enable Intruder Lockout, select the checkbox next to the ‘Enable Intruder Lockout’ option and save the changes. All changes will be immediate as soon as the save button is selected.

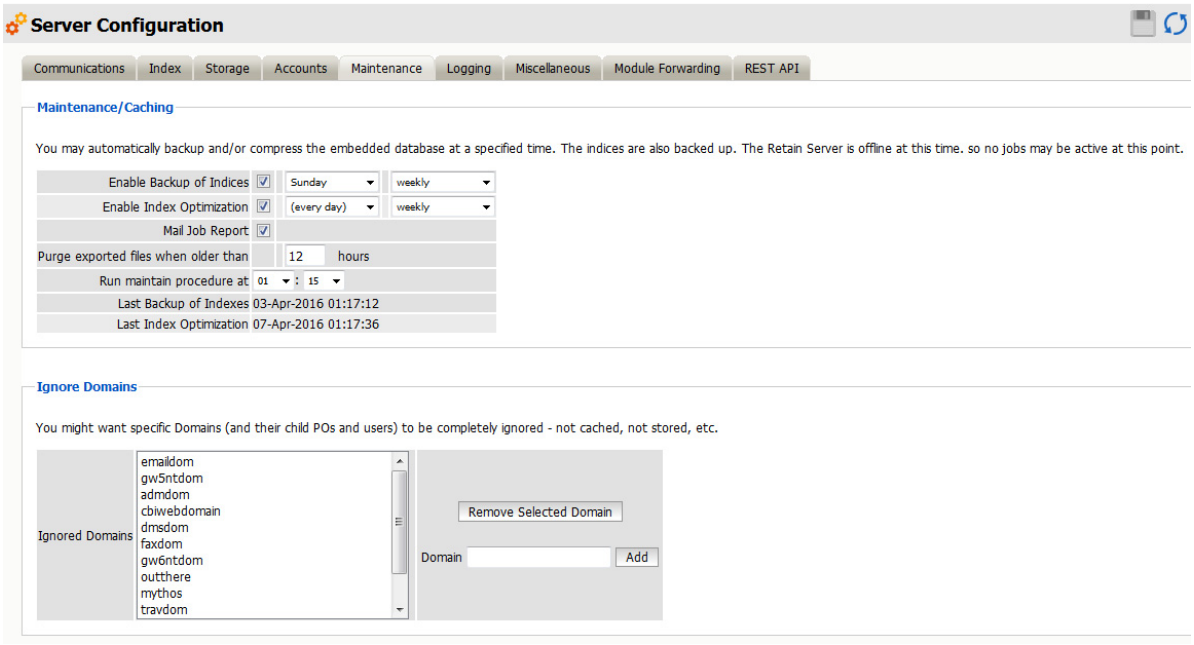
If a user has locked their account and requires immediate access to the system, all lockouts may be cleared. To clear any locked accounts, select the ‘clear lock outs’ button at the bottom of the page. There is no need to save changes; the clear command is immediate.

5.1.9 Maintenance tab

Maintaining a Retain system involves many tasks:

- Backup Configuration (Embedded) Database
- Compress Configuration Database
- Indexing (All Indexes)

These items are all controlled here.



This configuration database – the “Embedded” database:

- Should be backed up regularly
- Contains your system address book
- Also contains key information used by Retain
- Configure how often it gets backed up here.

Index Maintenance

You can also configure how often indexes get backed up and how often the indexes are optimized for speed.

WARNING: Index optimization will temporarily *double* the amount of index storage used as new index files are written. If a job is running the space use can *triple*. It is recommended to run index optimization on the weekend while no jobs are running.

Retain also lists a history of backups and maintenance here.

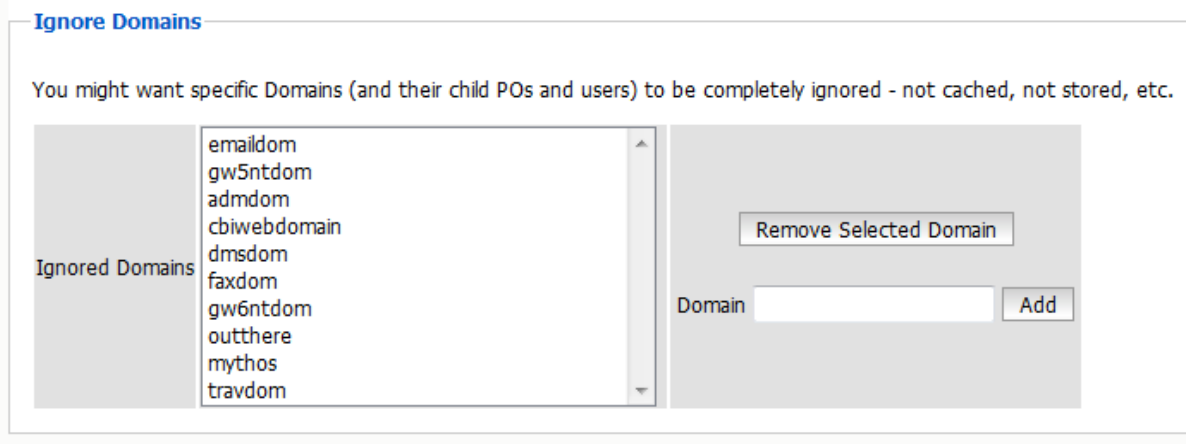
NOTE: This maintenance section does not involve backup or maintenance of the data store. You must back that up yourself with your normal backup tools.

You must backup:

- SQL Database
- File System where your data is stored

5.1.10 Ignore Domains

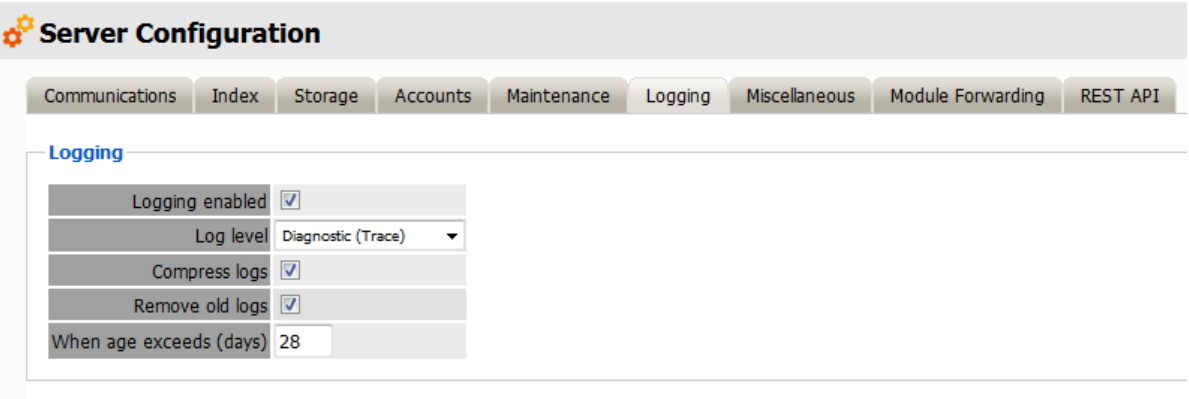
Some mail systems allow for the addition of non-system, or external domains. Sometimes, you might do this to add external domains to the mail system address book. However, the mail system cannot pull up e-mail from these domains and you may want to tell Retain to ignore references to these external domains to avoid wasting time during data collection.



Specify these domains in this window.

5.1.11 Logging tab

Configure the system logging here. Normally, you want verbose logging. Diagnostic is ordinarily done just for troubleshooting purposes.

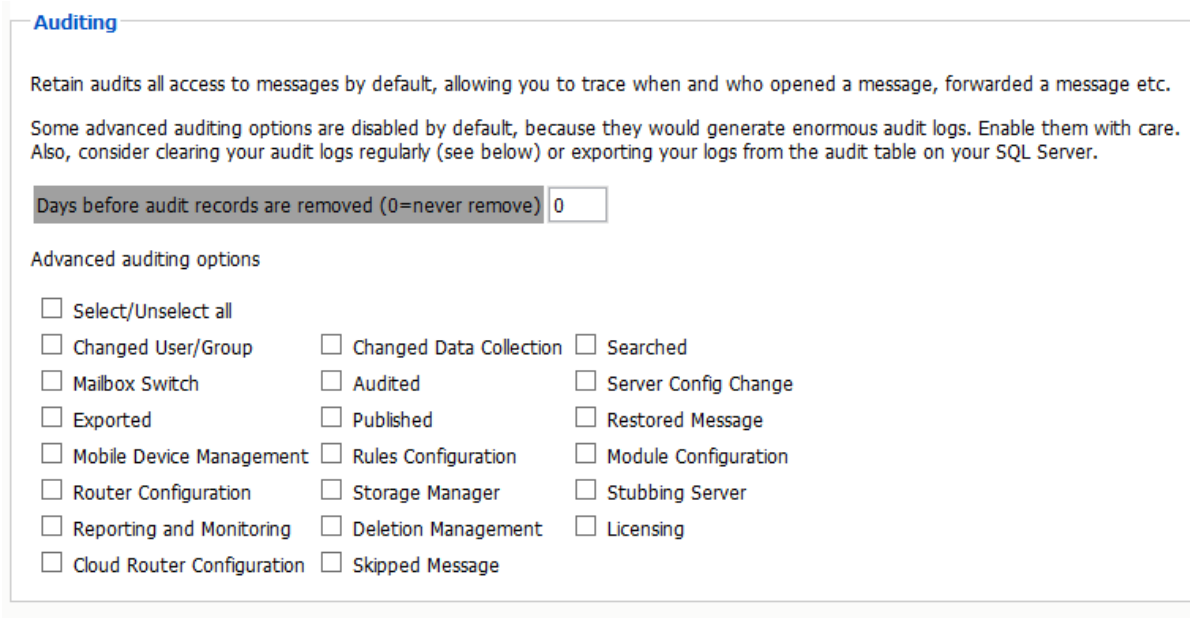


You also specify how long to keep logs here. By default, logs are deleted after 10 days. Logs may be com-pressed to save disk space.

5.1.12 Auditing Records

Retain now will create an auditing record of all actions, specified by the user, which are taken on a spe-cific piece of mail. Auditing records can be removed automatically after a specified length of time.

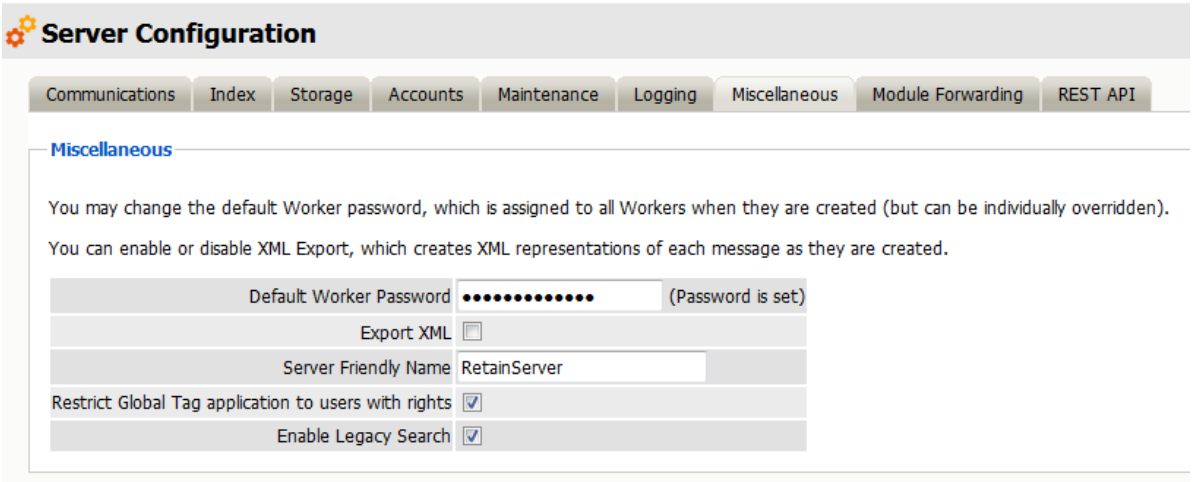
The variable is set in days. All of the options associated with this feature are found under the Server Con-figuration page, Logging tab.



This logging option creates very detailed activity logs for the options selected. It is very important to know that if every option is turned on the logs can become extremely large. However, the audit log cannot be searched for any items or activity which is not configured to be logged. If activity on any of the offered items it must first be enabled here. It is highly recommended that an expiration date is set for the logs so that they are automatically removed from the system to avoid filling up your disk space. Selecting all options for logging will also adversely affect performance. Do not select all the options at one time unless requested by Support.

5.1.13 Miscellaneous tab

You may choose the default worker password here. One is automatically generated for you when you create a worker but you may change the password here if you like.



5.1.14 Export XML

This XML Export function is included in Retain in case you have an XML compliance mandate. You enable it here.

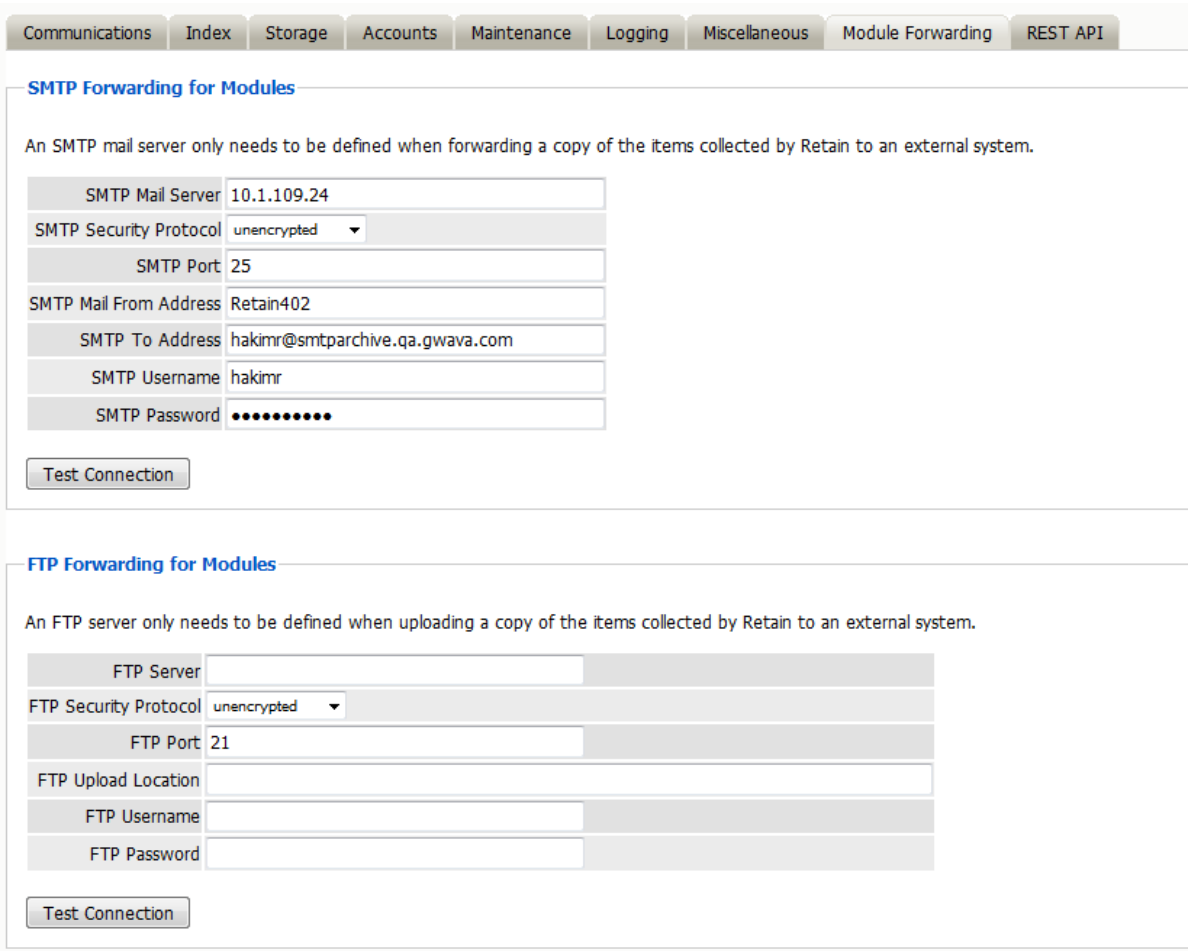
When selected, each attachment will have an XML export file of its parent message. In other words, an XML representation of the metadata is created and linked to the blobs as the messages are stored. There's nothing more done.

Ordinarily, you would not want to do this because it consumes enormous quantities of disk space, loads up your file system and degrades performance.

When you do NOT use the XML export function, you will benefit from Retain's single-instance message storage and data compression to save disk space and improve performance.

5.1.15 Module Forwarding

Modules can be set to forward all new items to another location.



SMTP Forwarding

SMTP Forwarding is a feature which instructs Retain to forward a copy of any archived message data to the specified address and domain, when the message item is archived. For Module Forwarding to function, it must be configured here on the Server Configuration and also enabled on the specific desired module. Any module which has the SMTP Forward option enabled will send messages according to the settings configured here.

This setting is designed for exporting data to an external SMTP system for redundant archival. This is most often used for mobile, social, or blackberry message systems. SMTP Forwarding will take a text, pin, sms, or posted message and message data, convert it into a MIME file and send the message to the specified SMTP system. The Forward process checks for and sends any queued data once every 10 minutes. If a message is unable to be sent for any reason, after 5 days the MIME file will be saved to the local disk and may be reviewed and repaired by an administrator.

This feature is not recommended for use with any SMTP system which Retain is currently archiving. If Module Forwarding is enabled and configured to use the SMTP system which Retain is currently archiving, duplicate data will be archived; Mobile, Blackberry, and Social data will be doubled in the Retain archive.

DO NOT use SMTP Forwarding with any email module if Retain is archiving the destination SMTP Mail System. This will cause a feedback loop which will rapidly fill the archive and email system.

Requirements

- Separate, working SMTP system
- Mobile module requires the mobile app 4.0.1 or higher

Configuration

The SMTP Forwarding feature requires a SMTP Mail Server connection configured. Input the destination Mail Server's DNS name, protocol, and port.

The SMTP Mail From Address will be the address which displays as the 'from' address for the forwarded messages.

The SMTP To Address is the destination account for messages. If a single address is to be used, it should be a journaling mailbox. Retain can also send each message to a destination mailbox for each originating user. If it is desired to send the messages to each respective owner's mailbox, configure the 'SMTP To Address' with:

{userid}@<yourdomain>.com

Retain will automatically use the userID of the originating device or account. Use of this feature requires that the SMTP server has an existing mail account for each user which matches the userID Retain shows for the mobile or Blackberry device.

If the destination SMTP server requires a login, provide an appropriate login. The login username and password will have no bearing on the 'from' address.

FTP Forwarding

This feature instructs Retain to take a copy of any archived message data to the specified location, when the data item is archived. This feature must be configured here, in Server Configuration as well as selected on the desired system module configuration page.

FTP Forwarding is designed for exporting data to an external FTP server for redundant archival. While the FTP Forwarding service does not have the inherent danger of duplicating data as the SMTP Forwarding feature does, it is important to note that should the exported data be sent through any system the Retain system archives, it will duplicate data in the Retain archive. FTP forwarding simply sends a copy of the data in the format that it is received, (Mime, text, etc.). The FTP Forwarding service is run every 10 minutes, checking for any data queued for delivery.

Requirements

- Separate, working FTP Server
- Module and Server Configuration must both be configured

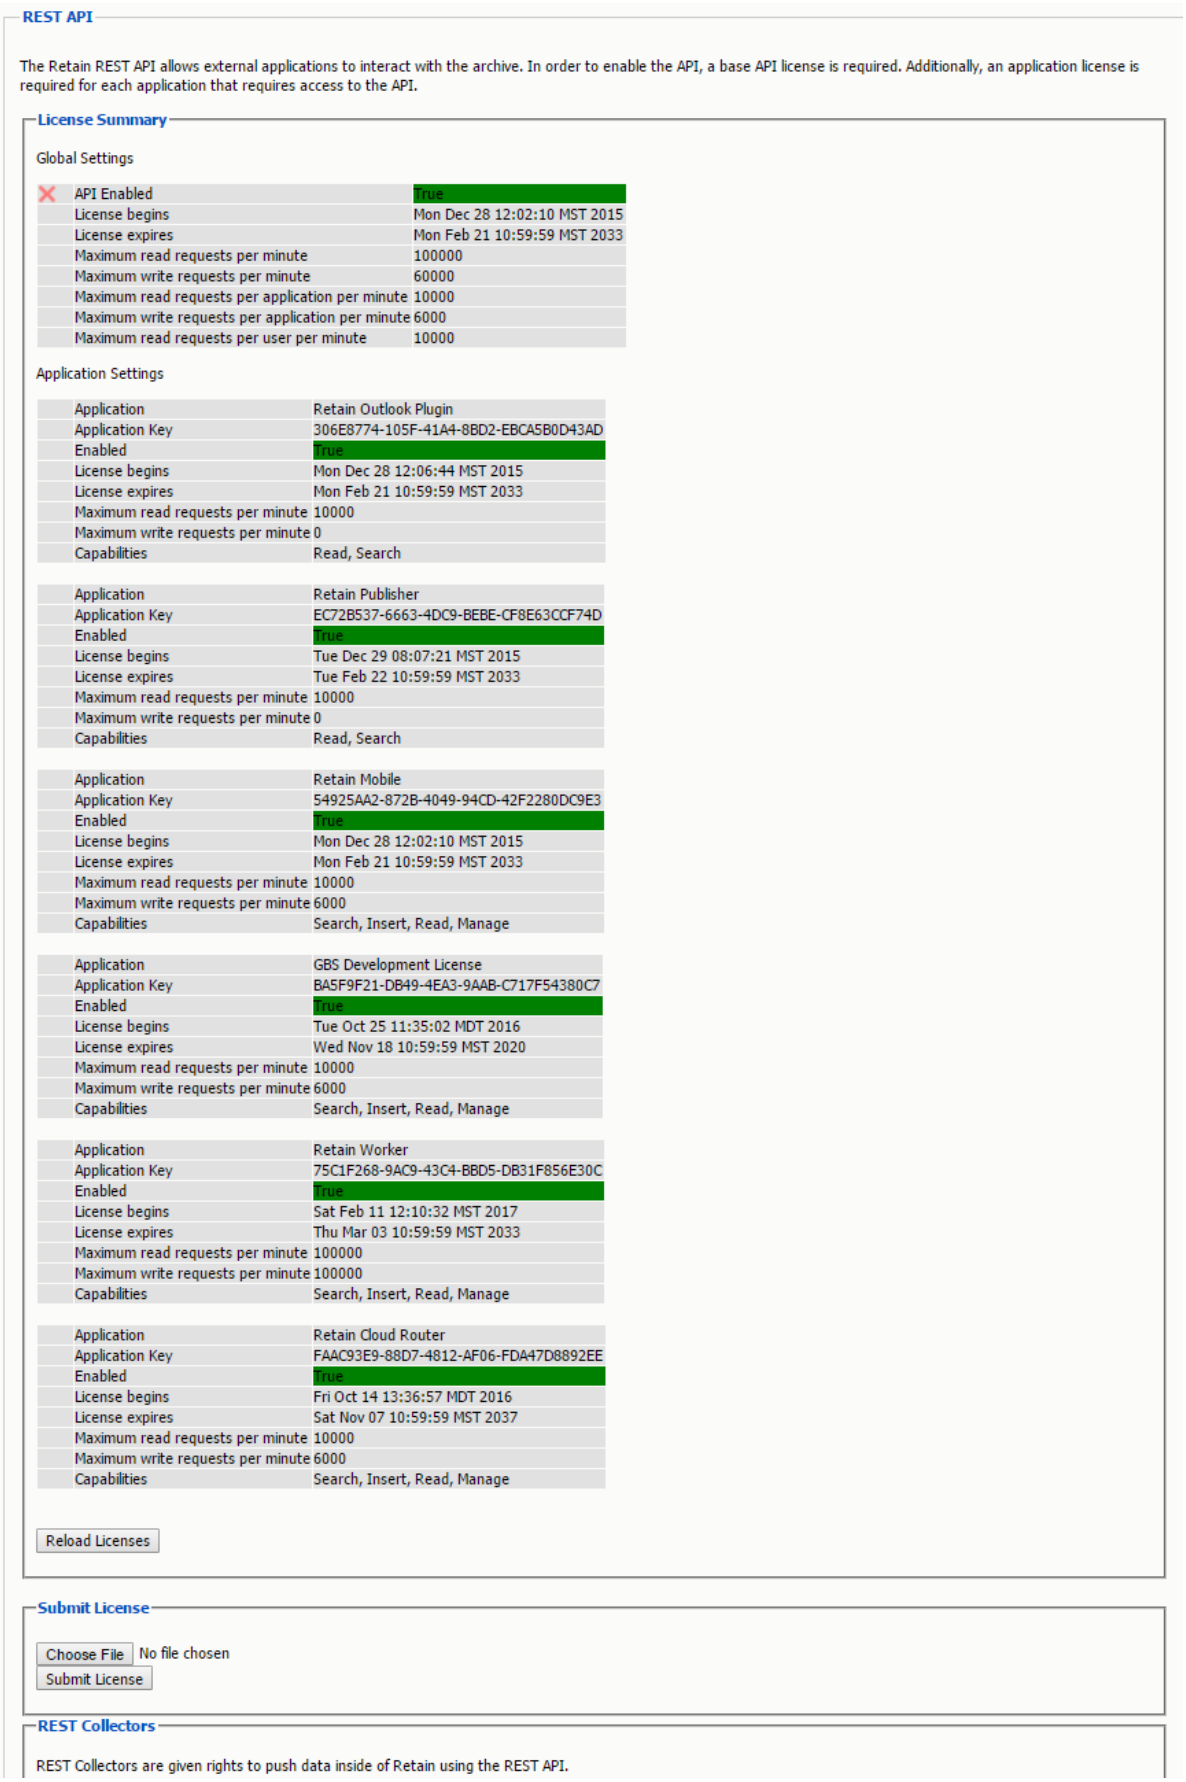
Configuration

- The FTP Forwarding feature must have a configured FTP Server connection before the forward option will appear in the Module configuration pages. To configure the FTP Forwarding option fill out all the required information and then test the connection. Test results will be displayed after the test completes.
- FTP Server is the connection address or DNS name of the FTP Server
- FTP Security Protocol is the security of the receiving FTP server. Options are Unencrypted, SSL, or TLS.
- FTP Port is the listening port of the FTP server.
- FTP Upload Location is the path to the desired destination directory of the FTP Server. This must be the location as you would view it in an FTP client.
- FTP Username and Password are the login credentials to the FTP Server.

Once configured, select 'Test Connection' and then save the settings by selecting the disk icon at the top right of the page. Now the FTP Forwarding option has been configured the FTP option will appear in the module forwarding section of each individual module configured in the system. Only the modules which have had the FTP option saved will utilized FTP Forwarding. The feature must be configured in both places in order to function.

5.1.16 REST API

The REST API is an interface built into Retain which allows third-party applications to input data into the Retain archive. This API was developed to open new horizons to the Retain Archive for any application which conforms to the API requirements. Because of the access the API grants to applications, an addi-tional API license file is required for each application. This license contains credentials and access rights. In order for a third-party application to utilize the REST API, the application must have a credentialed key provided by Micro Focus.



REST API

The Retain REST API allows external applications to interact with the archive. In order to enable the API, a base API license is required. Additionally, an application license is required for each application that requires access to the API.

License Summary

Global Settings

API Enabled	True
License begins	Mon Dec 28 12:02:10 MST 2015
License expires	Mon Feb 21 10:59:59 MST 2033
Maximum read requests per minute	100000
Maximum write requests per minute	60000
Maximum read requests per application per minute	10000
Maximum write requests per application per minute	6000
Maximum read requests per user per minute	10000

Application Settings

Application	Retain Outlook Plugin
Application Key	306E8774-105F-41A4-8BD2-EBCA5B0D43AD
Enabled	True
License begins	Mon Dec 28 12:06:44 MST 2015
License expires	Mon Feb 21 10:59:59 MST 2033
Maximum read requests per minute	10000
Maximum write requests per minute	0
Capabilities	Read, Search

Application	Retain Publisher
Application Key	EC72B537-6663-4DC9-BEBE-CF8E63CCF74D
Enabled	True
License begins	Tue Dec 29 08:07:21 MST 2015
License expires	Tue Feb 22 10:59:59 MST 2033
Maximum read requests per minute	10000
Maximum write requests per minute	0
Capabilities	Read, Search

Application	Retain Mobile
Application Key	54925AA2-872B-4049-94CD-42F2280DC9E3
Enabled	True
License begins	Mon Dec 28 12:02:10 MST 2015
License expires	Mon Feb 21 10:59:59 MST 2033
Maximum read requests per minute	10000
Maximum write requests per minute	6000
Capabilities	Search, Insert, Read, Manage

Application	GBS Development License
Application Key	BA5F9F21-DB49-4EA3-9AAB-C717F54380C7
Enabled	True
License begins	Tue Oct 25 11:35:02 MDT 2016
License expires	Wed Nov 18 10:59:59 MST 2020
Maximum read requests per minute	10000
Maximum write requests per minute	6000
Capabilities	Search, Insert, Read, Manage

Application	Retain Worker
Application Key	75C1F268-9AC9-43C4-BBD5-DB31F856E30C
Enabled	True
License begins	Sat Feb 11 12:10:32 MST 2017
License expires	Thu Mar 03 10:59:59 MST 2033
Maximum read requests per minute	100000
Maximum write requests per minute	100000
Capabilities	Search, Insert, Read, Manage

Application	Retain Cloud Router
Application Key	FAAC93E9-88D7-4812-AF06-FDA47D8892EE
Enabled	True
License begins	Fri Oct 14 13:36:57 MDT 2016
License expires	Sat Nov 07 10:59:59 MST 2037
Maximum read requests per minute	10000
Maximum write requests per minute	6000
Capabilities	Search, Insert, Read, Manage

Reload Licenses

Submit License

Choose File No file chosen

Submit License

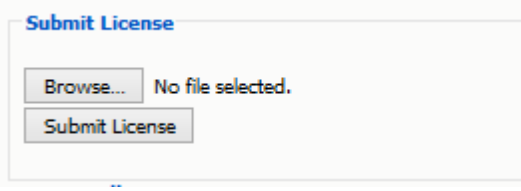
REST Collectors

REST Collectors are given rights to push data inside of Retain using the REST API.

The REST API is an input protocol only. For Archive security, the API prohibits migrating data out of the Retain Archive. There are also limits placed in each applications API license file to protect the integrity of the data being input into the archive, based on the application.

Submit License

To gain a credentialed key for a third party application, contact a Micro Focus sales representative. Use of the REST API will require a REST API base license, and an application specific license for each desired applica-tion.



REST Collector

Some modules use the REST API for data collection. You need to create a collector to connect to the message system. You may have multiple REST collectors.



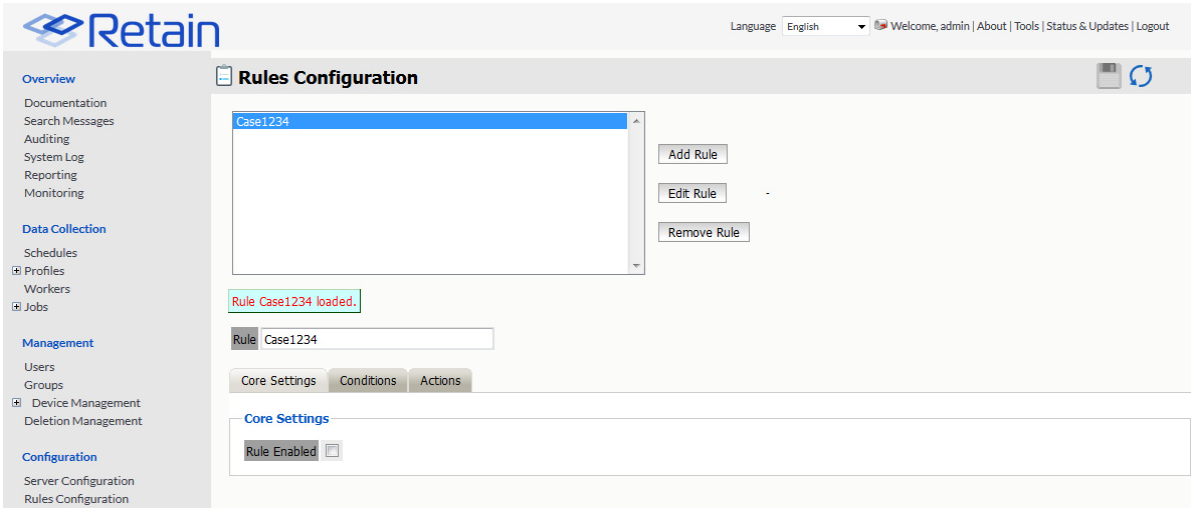
Create a REST Collector by providing:

1. The *Collector Name* can be anything you want, but we recommend something easy to identify with the system it connects with. A Key and Secret will be created for the collector. The Key and Secret will be used on the message system side of the system for authentication.
2. Then press "*Add collector*" to enable the collector.
3. *Export* the collector allows you to export the Key and Secret to disk for use with the message system.

Delete removes the collector.

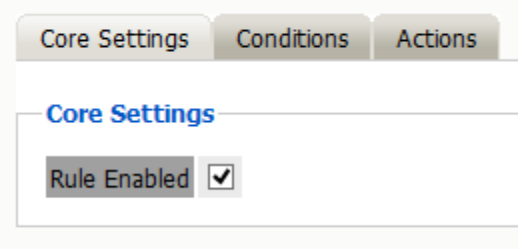
5.2 Rules Configuration

Rules for Retain is a way for administrators to automate tasks and administration based on criteria. This means that any message containing items or information specified in the criteria to have the described action applied automatically. A rule must be enabled, have at least one condition, and one action defined before a rule comes into effect.



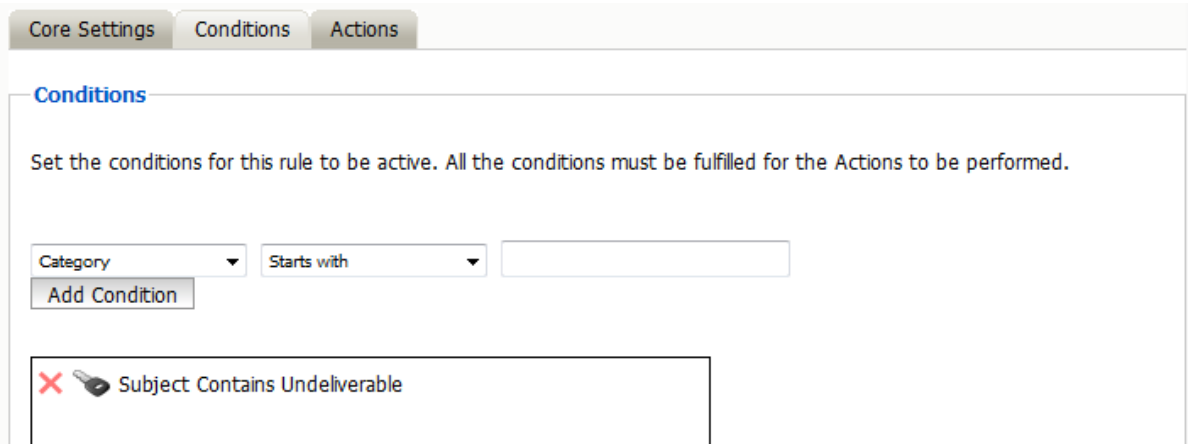
5.2.1 Core Settings tab

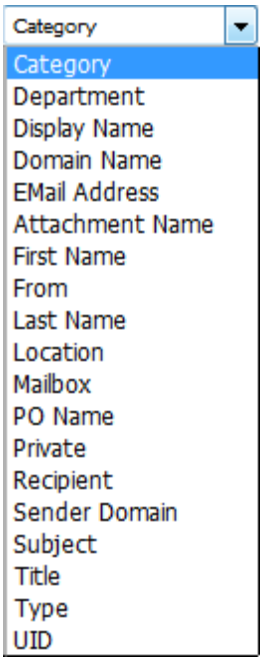
To create a rule, select the ‘add rule’ button and name the rule. To make the rule active, select ‘Rule Enabled”.

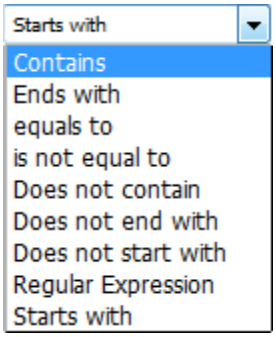


5.2.2 Conditions tab

The conditions are the criteria which determine whether the action is applied to message data or not.





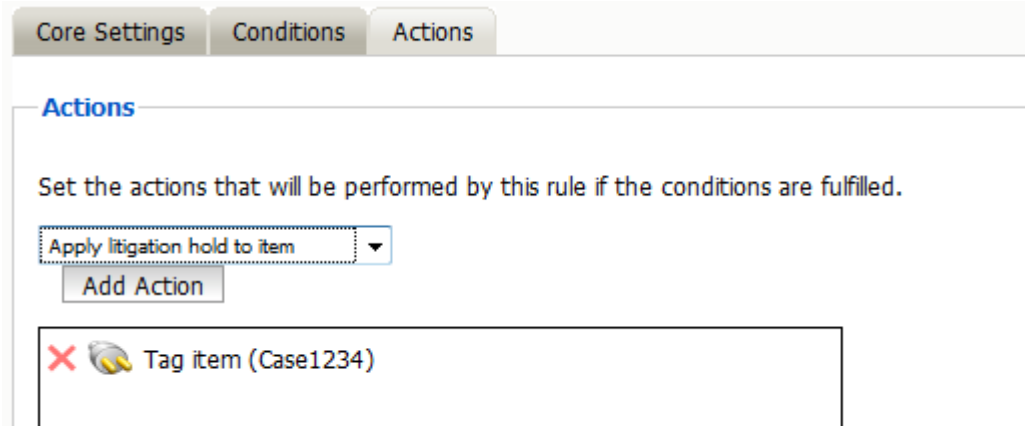


Criteria is entered in plain text or in regular expression, and are separated into categories. Select the cat-egory that the criteria is to be restricted to, from the menu.

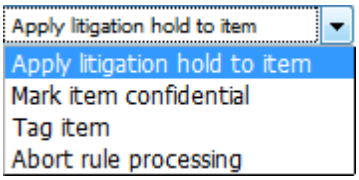
Once a category has been selected, the criteria to be applied to that needs to be specified. The ‘Starts with’ menu further defines the criteria. The ‘Starts with’ menu also allows the administrator to either relax some constraints by selecting a vague entry like ‘Contains’ or ‘Starts with’, or specify the criteria with ‘Regular Expression’. Select the appropriate limiter and then select the ‘Add Condition’ button. Multiple conditions may be added. Make sure to save changes before leaving the page.

5.2.3 Actions tab

The actions tab defines which actions are to be taken. There are four different options to choose from.



All four actions may be applied to any one rule. Once the desired actions are applied to a rule, save changes before leaving the page.

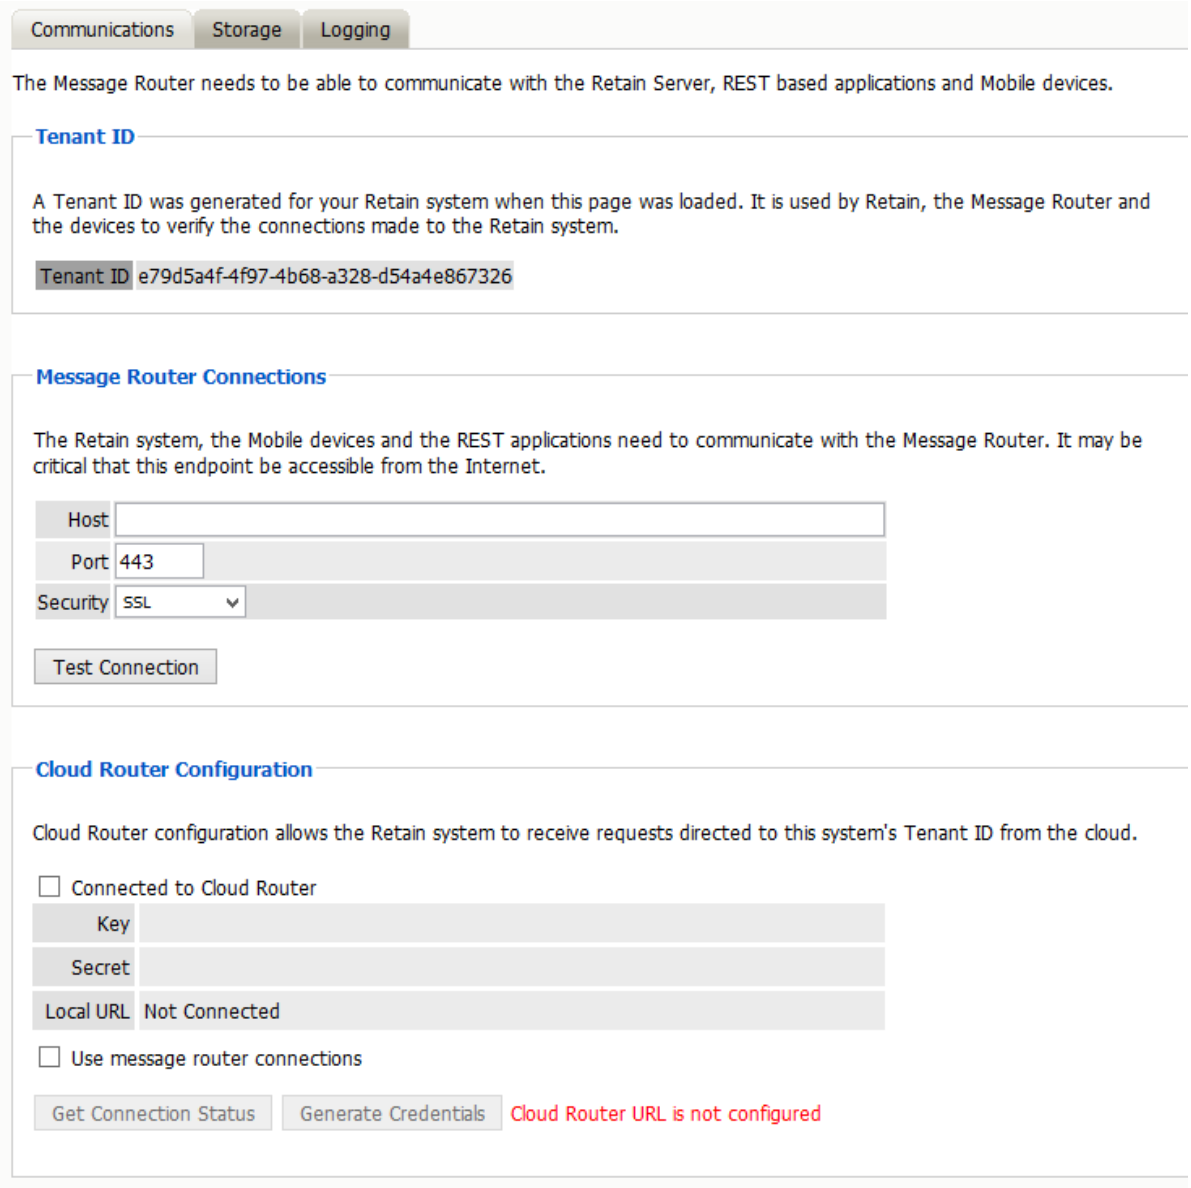


5.3 Router Configuration

Configuration for the Message Router may be changed or checked on this page. The Router is used in conjunction with the mobile module, the BBM module integration, and any other REST API integrations. In contrast with the Retain Server, the Router is best placed in a network's DMZ, enabling it to be fully accessible to communication from the Internet as well as the network.

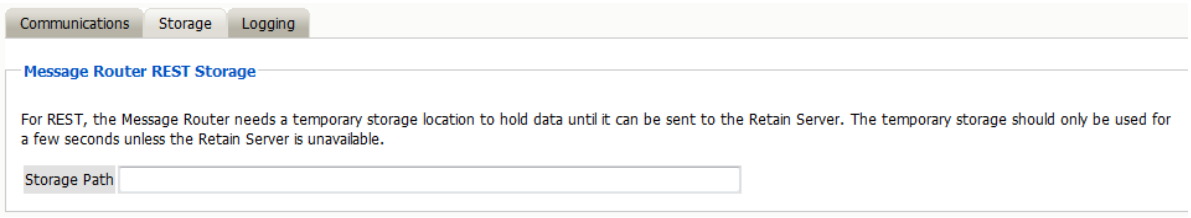
5.3.1 Communications tab

The Router's communication settings and port for both the Retain Server and the Internet is kept here. This is the address that the mobile devices will use to talk to the router and should be accessible to the outside world.



5.3.2 Storage tab

The storage tab contains the path to the temporary storage location local to the REST Router.



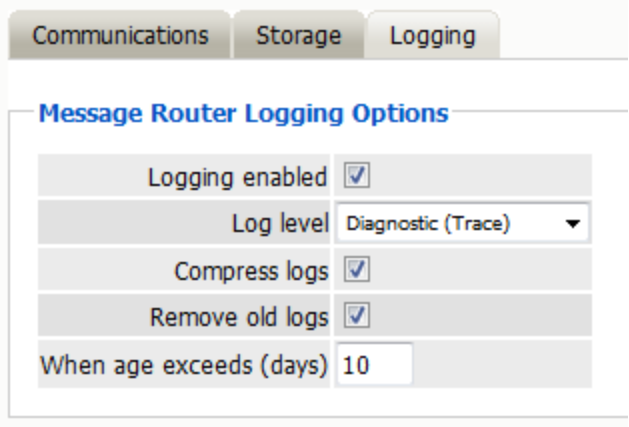
The storage is used by the REST Router to store data for a few seconds before it is sent on to the Retain Server. In the event of a failure to access the Retain Server, the REST Router will house the data until the Server is accessed. This storage area is not required to be large, and a few gigabytes is

more than sufficient. When specifying the Storage path, note that the Retain System automatically appends the storage directory. So specifying a path of "/Retain/storage" will result in an actual path of "/Retain/storage/storage".

NOTE: The Storage Path must not be the same as the Mobile Module Data path.

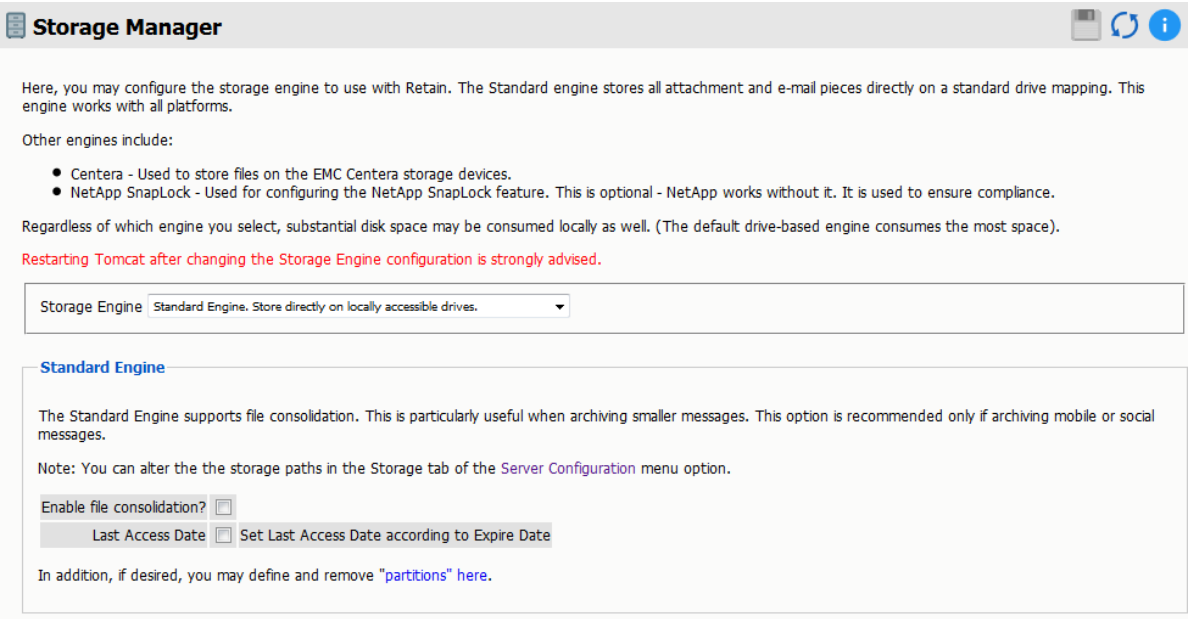
5.3.3 Logging tab

The logging tab contains settings to enable logging, dictate the logging level of the Router, and settings to purge old logs from the system after a specified age.



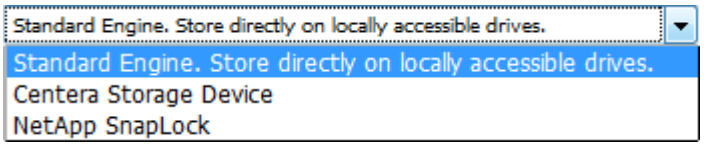
5.4 Storage Manager

Retain Supports the ability to change the location or engine used to store the archived mail.



To change the storage option settings which were created during the initial setup of Retain, simply select the different storage engine of your choice.

If the Standard Engine is selected, then the Standard Engine window is displayed. This details the setting location for the default storage location, under Server Configuration |Storage Tab.



The other two options are:

- Centera Storage Device
- NetApp SnapLock

Standard Engine

The standard storage engine is the standard storage option which should be selected for all new installations of Retain. The storage engine utilizes locally accessible drives and media. When selecting this option, make sure that storage requirements will be easily met, size or capacity increased, and managed.

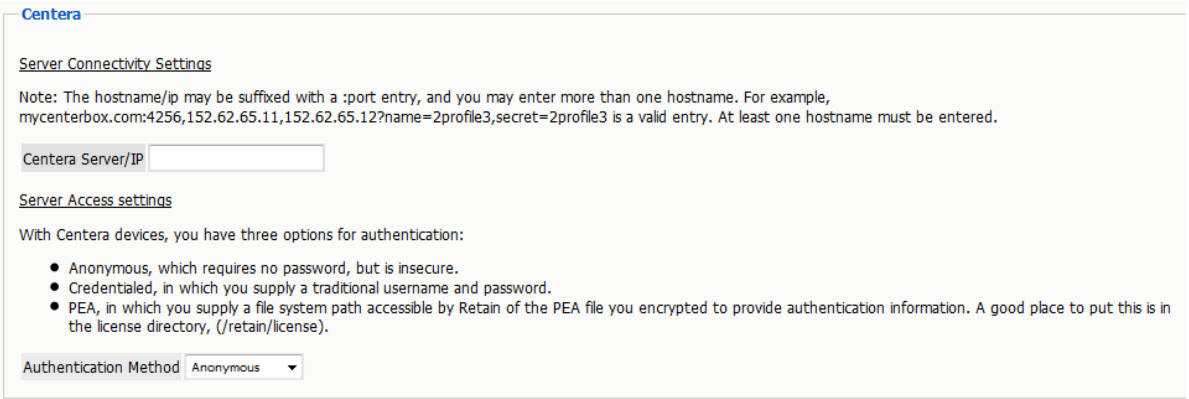
The engine also allows for background compression to be used. Background compression is recommended for social media and mobile modules, to help keep the size of the data store down.

Legacy Engine

The legacy engine works as a seamless connection to the 2.x line of Retain, and is available as a bridge for upgraded devices and legacy support systems.

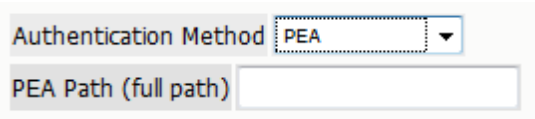
Centera Storage Device

The Centera Storage option simply requires you to put the storage server IP address or DNS name, and the authentication method.

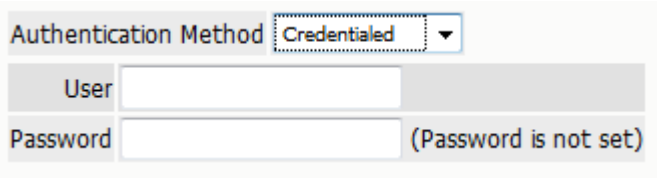


The different Authentication options are PEA and Credentialed.

The PEA requires the full path to the PEA file.

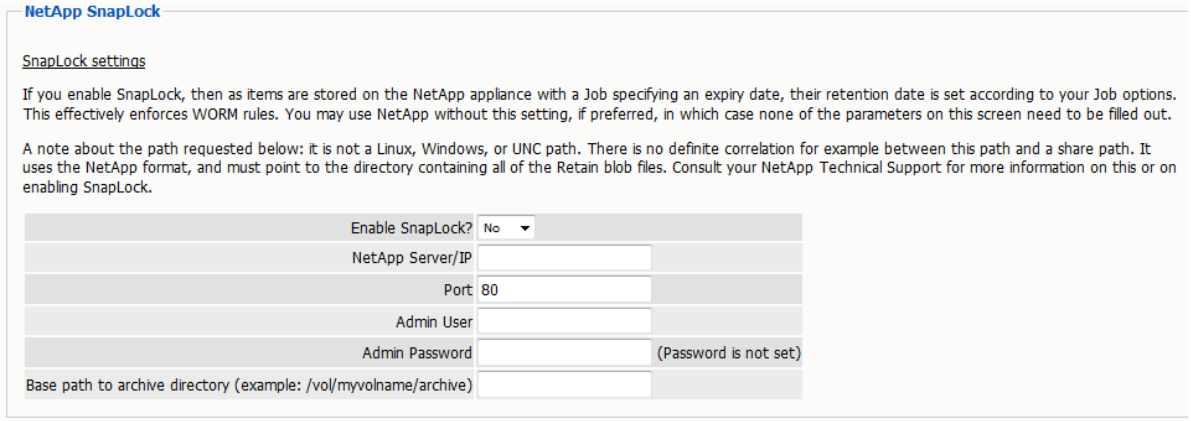


For a credentialed authentication, simply input the username and the password for the Centera server.



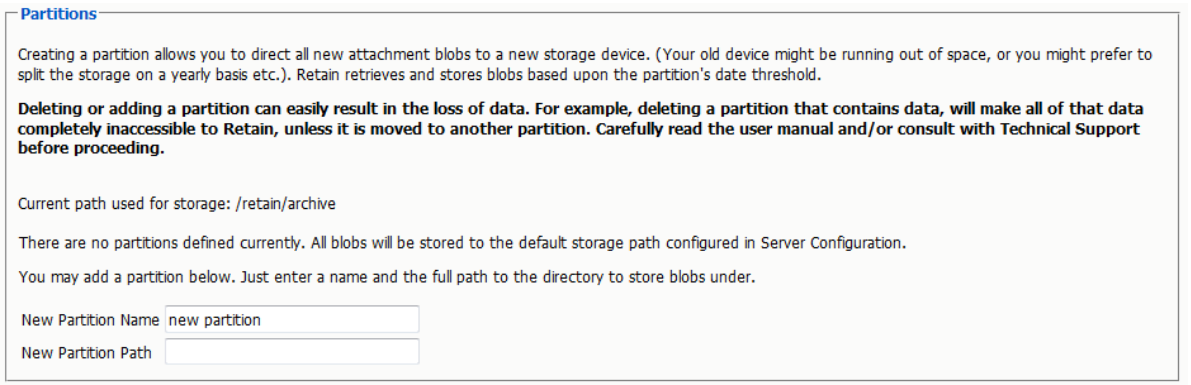
NetApp SnapLock

NetApp is fully supported with SnapLock for storage. Note the information in the option. To enforce WORM rules you MUST set the Job Expiration feature. Also set the IP address, user name, password, and the correct path to the server. The path must be set in netapp format.



5.4.1 Partitions

Retain also supports the ability to define a different storage location, through the partitioning option. Partitioning is only available when Storage is set to the ‘local’ drive. To access the partitions option click on “partitions” in the Enhanced or Standard Engine window. Ideally, this option would only be used to archive to a new, larger, location when the current storage path has been filled, or to separate mail stor-age between years. (I.e. /archive/2008; /archive/2009)



A Retain system without a specified partition location will store all mail in the default storage location. This is defined in the Server Configuration menu, under the Storage tab mentioned previously. The default location is set during the initial install.

Read the bold warning on changing partitions!

You may only specify one new partition at a time. After you have entered the new partition information, click ‘save changes’. Creation or changes to a partition requires Tomcat to be restarted. When adding a Partition on Linux, the path needs to be created manually on the Linux Server before specifying the path within Retain. Ensure that the permissions for the new path are given to Tomcat.

Once Tomcat has been restarted, the new partition can be viewed, edited, or deleted from this window.

Partitions are listed with their Begin Date, Name, and Path. The Begin Date is the creation date. A parti-tion is used to store all archived mail from the moment it is created

Because of how Retain indexes and tracks messages that are archived, deleting a partition can lead to catastrophic data loss. Do not delete an active storage partition or location without first consulting sup-port.

If a partition is deleted Retain will look for the archived mail that was stored in the deleted partition to be stored in the default storage location. (This is defined in Server Configuration, under the Storage Tab.) Removing a designated partition in the management console does not delete the storage directories or data on the hard drive or SAN, but it does disconnect Retain from the storage location.

All partition changes are logged in the partition.log file, located in the backup directory specified in Server Configuration. (By default this is located in <your storage location>/backup)

5.5 Stubbing Server

Retain has the ability to use the Stubbing capabilities of GroupWise. Stubbing was designed as a way to reduce disk space usage in a time when the cost of disk space was at a premium. This is less of an issue today.

5.5.1 Overview

The Stubbing server has two abilities:

1. Enabling Retain Archive search from the GroupWise desktop client.
2. Removing items from GroupWise and leaving stub pointers to the data in the Retain Archive. This functionality CANNOT be reversed. *We do not recommend enabling this functionality.*

How Retain Interacts With GroupWise Stubbing

When GroupWise is set up for stubbing and the Stubbing Server is configured in Retain, the archive job that runs will archive all the items and place in their stead in the GroupWise mailbox a "stub". Since the actual message gets placed into the Retain system, GroupWise replaces the original message with a stub that points to the message residing in Retain. A user's GroupWise mailbox will look the same - the user will see a list of messages; however, the icon is different.

When a user clicks on a stubbed message, the GroupWise POA will send a request to the Retain Stubbing Server for the item. Retain sends the item to the POA and the client displays it; however, this does NOT place the message back into GroupWise. When the user closes the message, the stub is all that remains.

Why Stubbing Is Good

- Save Disk space. Mail can be removed from the system without being deleted.
- Users have easy access to items stored in Retain.
- Transparent access, no need to login to Retain for common tasks.
- The "Restore" feature in Retain can be enabled. However, it does not actually restore the message into GroupWise. It only restores the stub if the stub is missing (user deleted it).

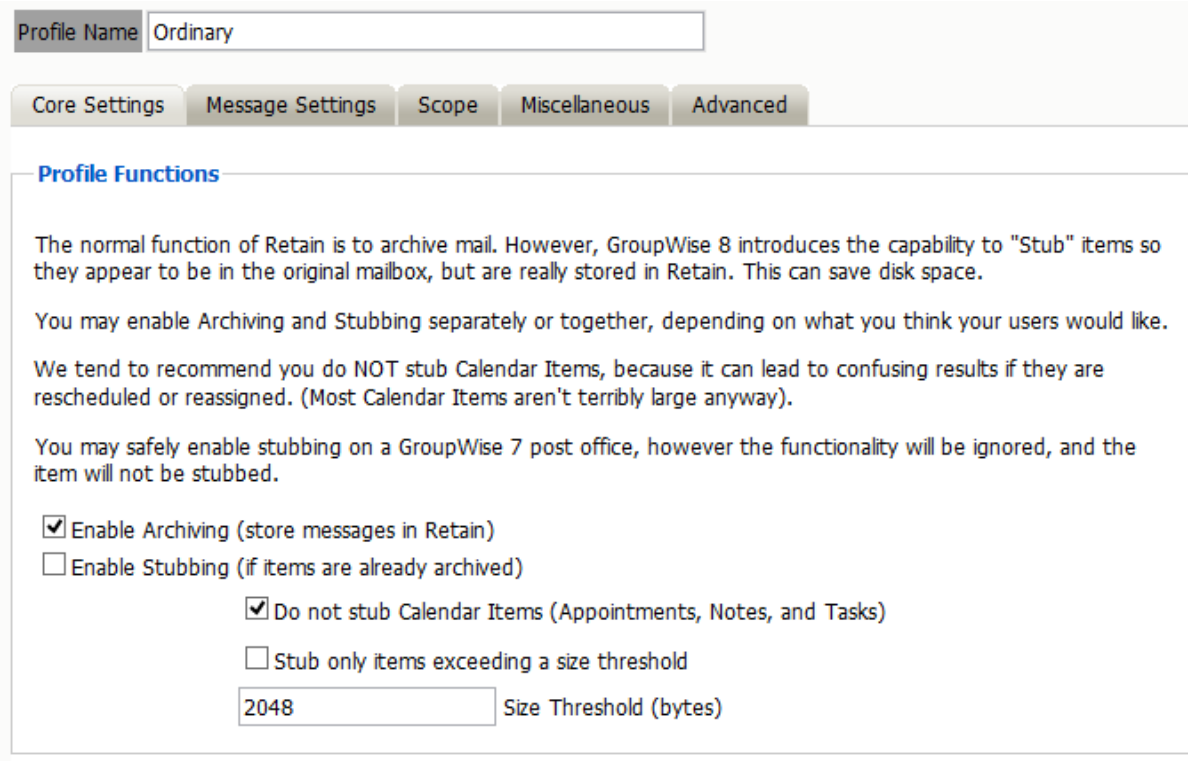
Why Stubbing May Not Work for a System

- *Stubbing CANNOT be reversed.* Once a message is stubbed, it cannot be returned to the same position in the system.
- *Stub items are NOT COMPATIBLE with third party APIs*. These applications view the stubs as posted items only, with no text or attachments. Mobile, Web, IMAP, and POP3 clients will only see posted empty items, not actual contents.
- User training for stubbed items may be necessary due to decreased access speed.
- No method for protecting stubs from deletion in the GroupWise client (re-creation is difficult and complicated).
- Non-mail items should not be stubbed (calendar and appointment items behave erratically when rescheduled).

Determine the correct option for your system. If you wish to enable Stubbing, you need to install and configure the Stubbing Server as well. If you wish to only stub large items in your mail server, you may set a size threshold. Any item larger than the specified size will be stubbed. If this option is not set, all messages will be stubbed regardless of size.

5.5.2 Prerequisites

- You must be using GroupWise 8.0HP1 or greater.
- The GroupWise module must be configured. This means that a Trusted App Key exists for Retain to use already.
- Confirm that stubbing is *disabled* in *all* the GroupWise profiles:



5.5.3 Stubbing Install

If you did not install the Stubbing Server during the initial install of the Retain Server you can install it now.

Run the installer and click through to the point where the installer detects what you have installed.

```
Retain Installer 4.1.0.0


Retain has already been installed on this system with the following components:

Retain Server
Retain Reporting & Monitoring Server
Retain Worker

Choose one of these options:
1) Upgrade current installation of Retain
2) Upgrade current installation of Retain and add new components
3) Exit this installer

Choose (1-3): 
```

Choose "2) Upgrade current installation of Retain and add new components."

```
Retain Installer 4.1.0.0


Retain has already been installed on this system with the following components:

Retain Server
Retain Reporting & Monitoring Server
Retain Worker

Choose one of these options:
1) Upgrade current installation of Retain
2) Upgrade current installation of Retain and add new components
3) Exit this installer

Choose (1-3): 2

Choose components...
Install Retain Stubbing Server (y/n)?y
Install Retain Message Router (y/n)?n
```

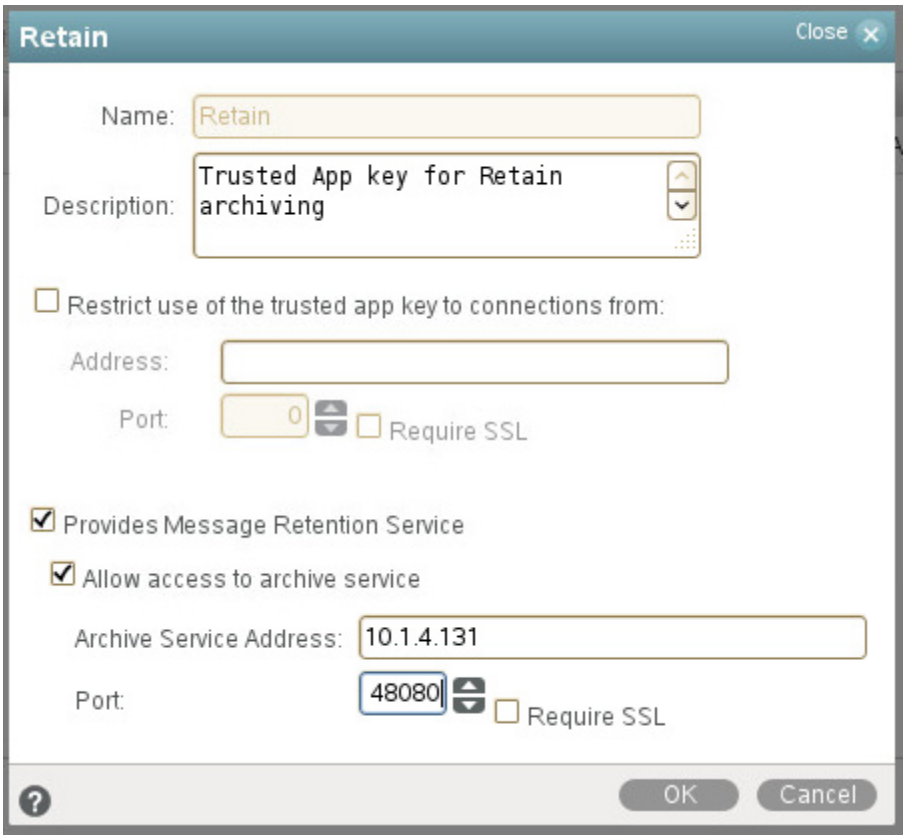
Then tell the installer to Upgrade Retain, preserving configuration. The installer will stop tomcat, install the Stubbing server and restart tomcat.

5.5.4 Stubbing Configuration

GroupWise

Update Trusted Application Key

1. Go to the Retain Trusted App Key and enable "Allow access to archive service"
2. Set the Archive Service Address to the Retain server IP Address or hostname
3. Set the port to 48080 (the Retain tomcat port)



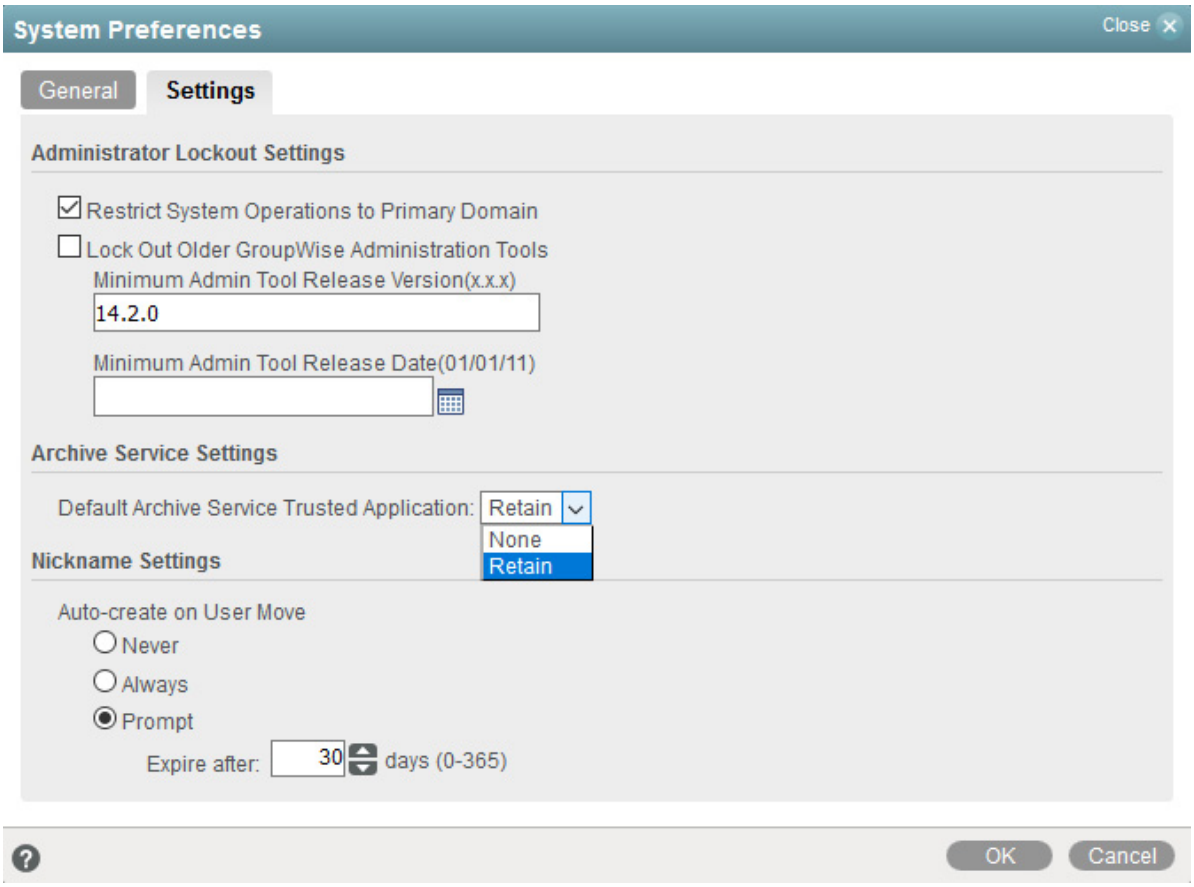
Now you have to decide if you are going to use stubbing on a particular Post Office or System wide.

Archive Service System Setting

If you are enabling stubbing across the entire system you will need to override the Archive Service Trusted Application setting in System Preferences.

Go to System | System Preferences.

Under Settings and Archive Service Trusted Application, enable the Override and select the Trusted Application Key that Retain is using.

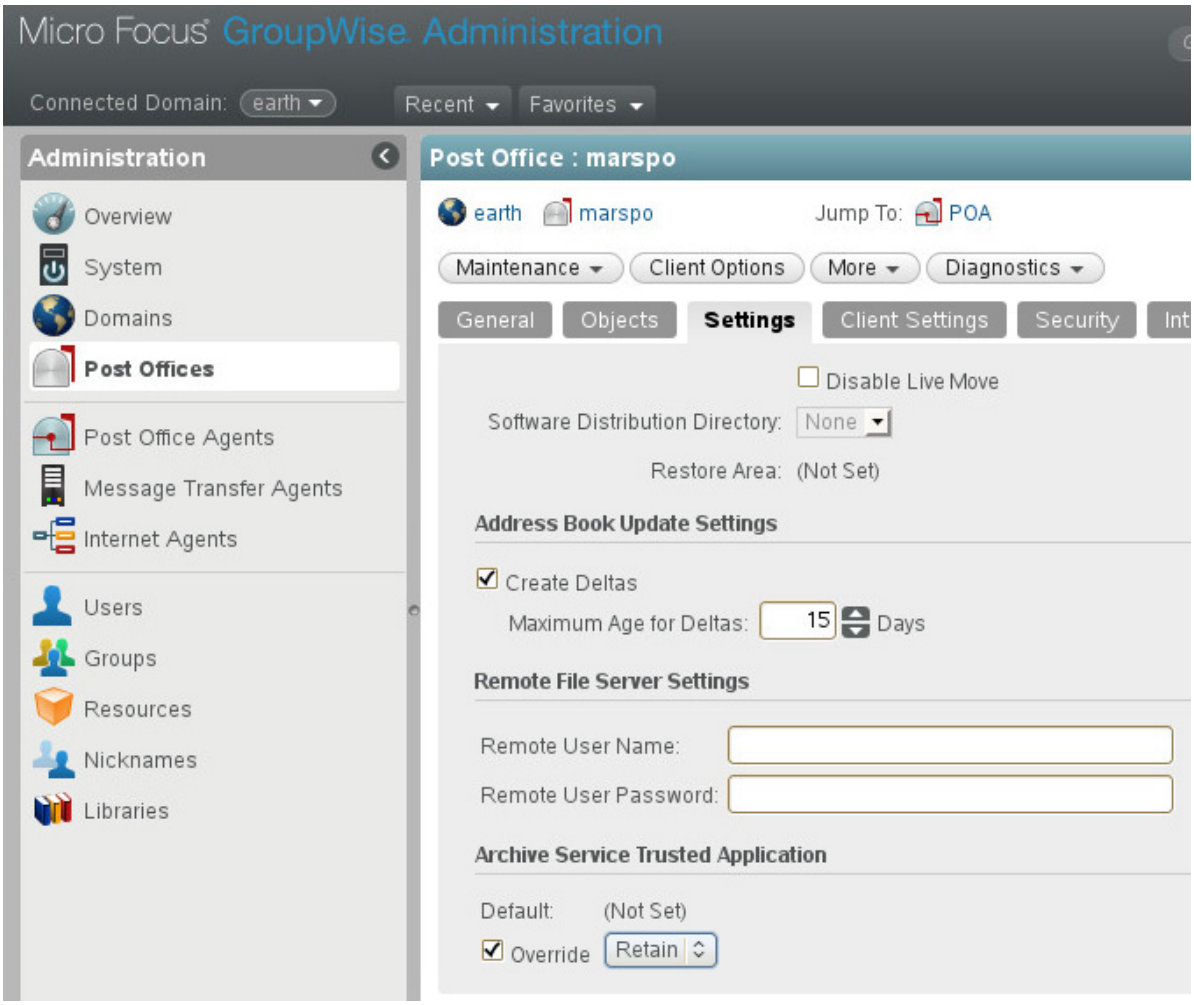


Archive Service Post Office Settings

If you are just stubbing one, some but not all Post Offices you will need to override the Archive Service Trusted Application setting for each Post Office.

Go to the Post Office that will be stubbed.

Under Settings and Archive Service Trusted Application, enable the Override and select the Trusted Application Key that Retain is using.



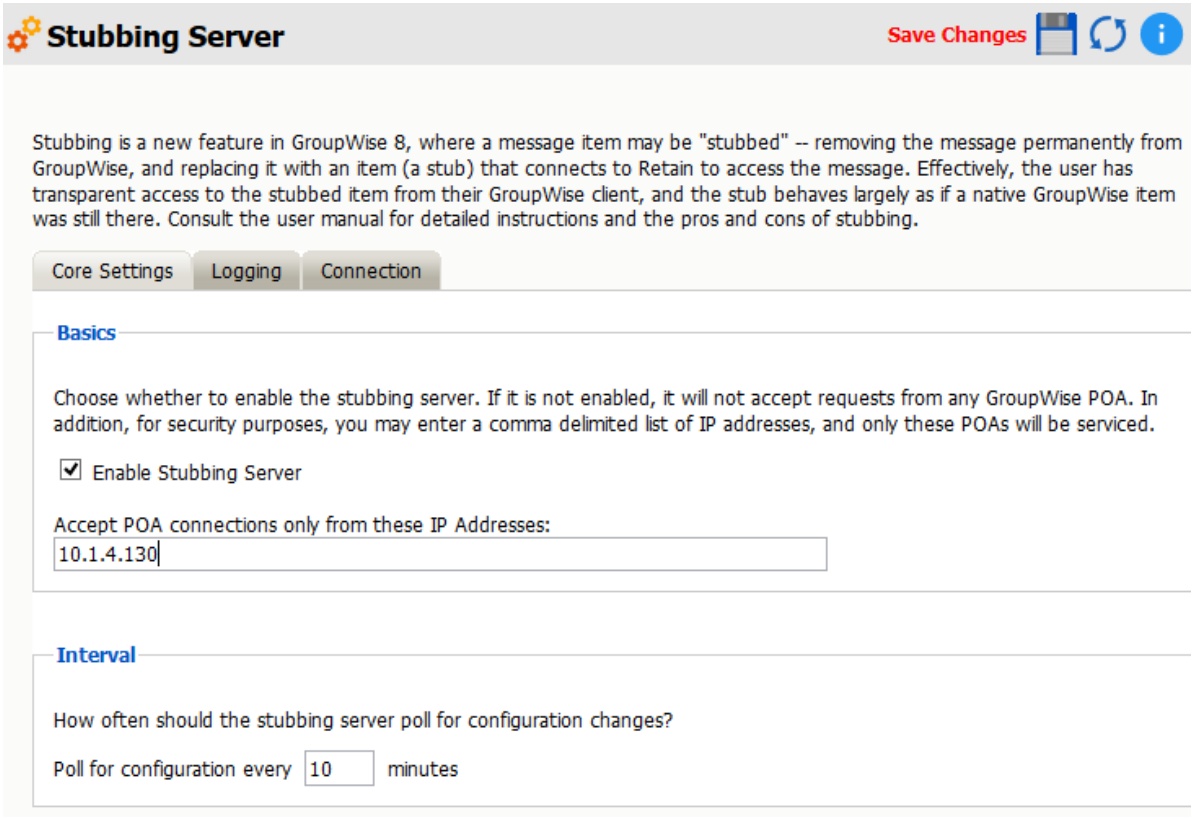
5.5.5 Retain

In Retain, open the Stubbing Server page from the left-navigation bar.

Core Settings tab

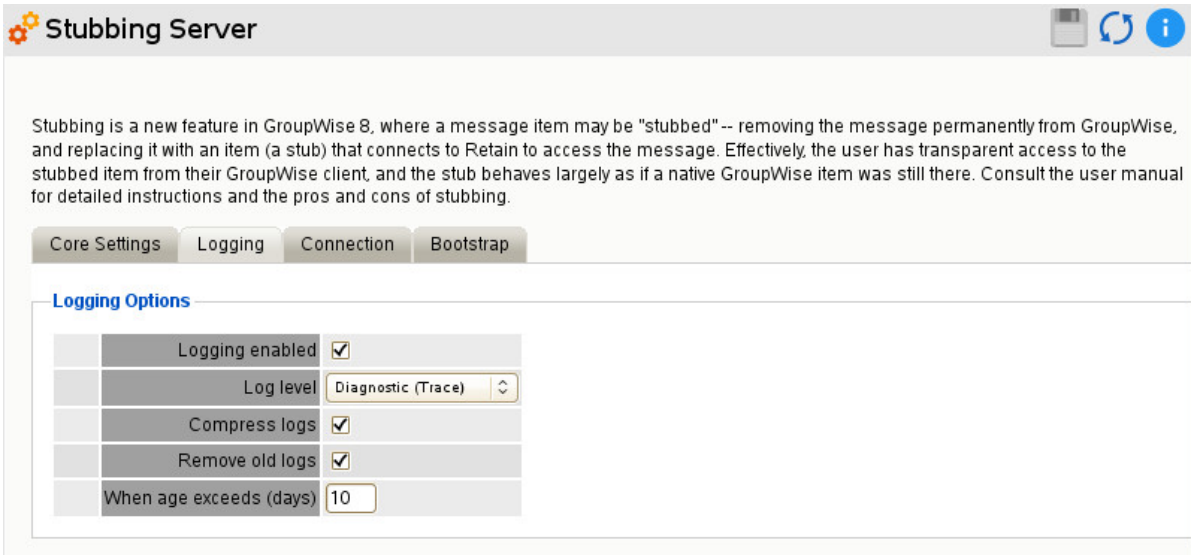
Check the box next to "Enable Stubbing Server"

You may restrict the IP addresses POA connections will be accepted from. Use commas to separate multiple addresses.



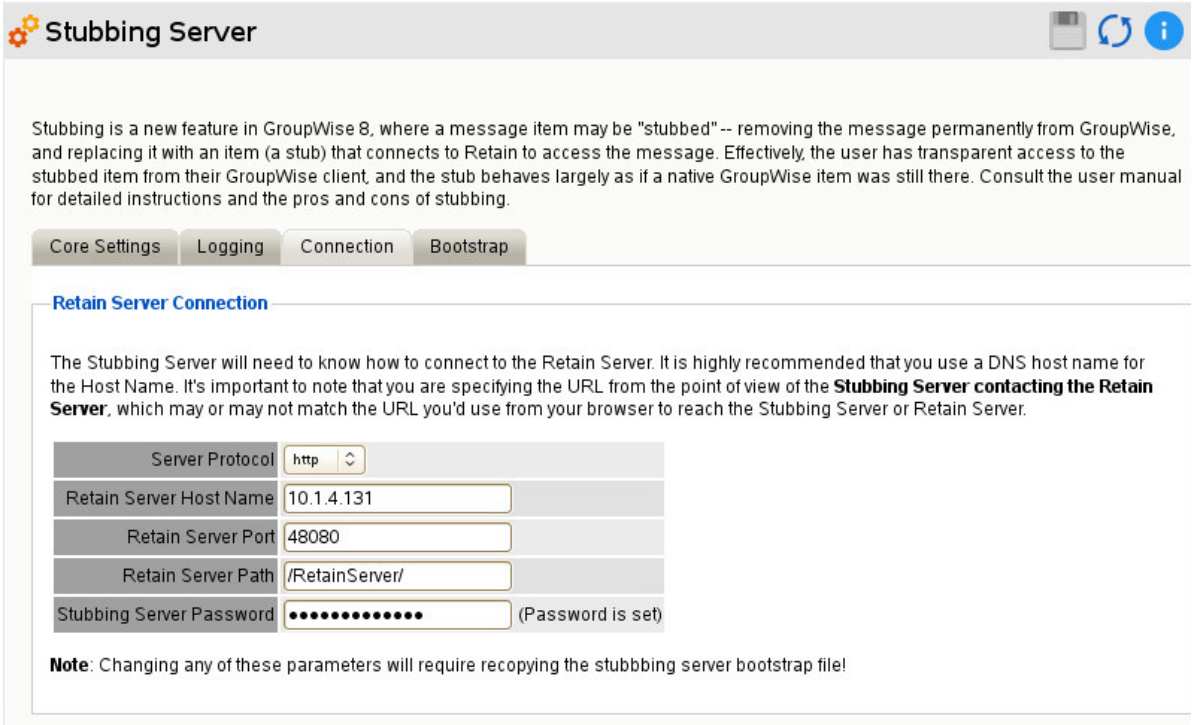
Logging tab

Set the logging level to Diagnostic (Trace)



Connection tab

If you have installed the stubbing server on a separate server you will need to specify the location of the Retain server so that it can connect.

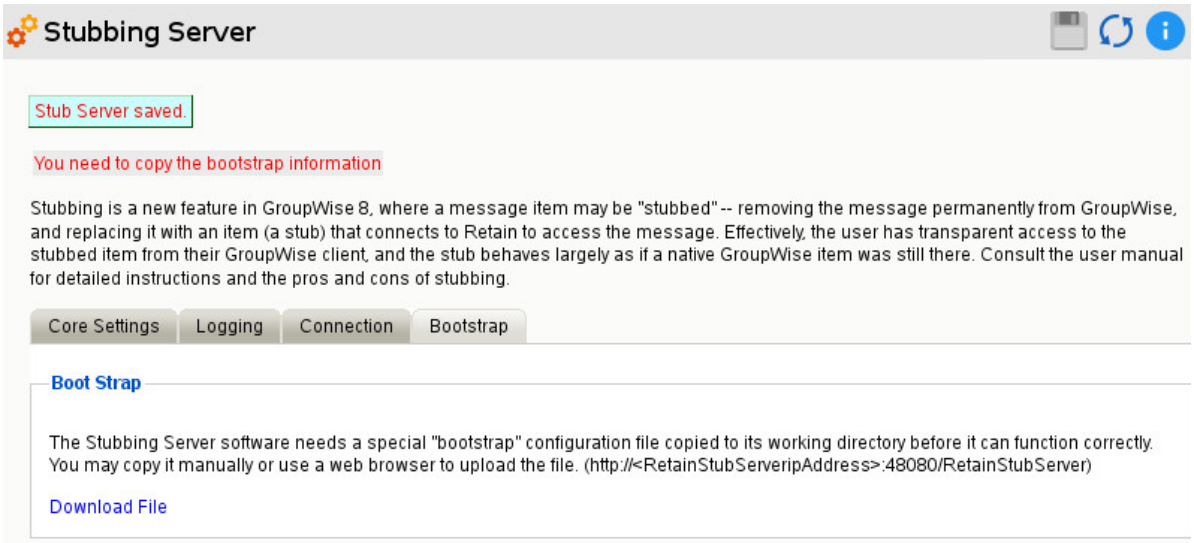


The connection to the Retain Server is very important. This must be set to the address of the Retain Server, and is automatically filled out for you. Verify that the Retain Server is located at the specified address and is accessible. Generally you may leave this section default as the Stub Server should be installed local to the Retain Server. If your Retain Server IP address is assigned by DHCP, which we do not recommend, change the Host Name to the DNS name or localhost (127.0.0.1).

Press "*Save Changes*" and the Bootstrap tab will appear.

Bootstrap tab

Click on "Download File" to save the bootstrap file.

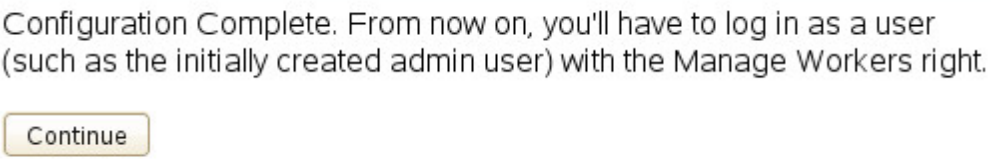


5.5.6 Configure Stubbing Server

Browse to http://<Retain_Stubbing_Server_Address>/RetainStubServer or http://<Retain_Stubbing_Server_Address>:48080/RetainStubServer and upload the bootstrap.



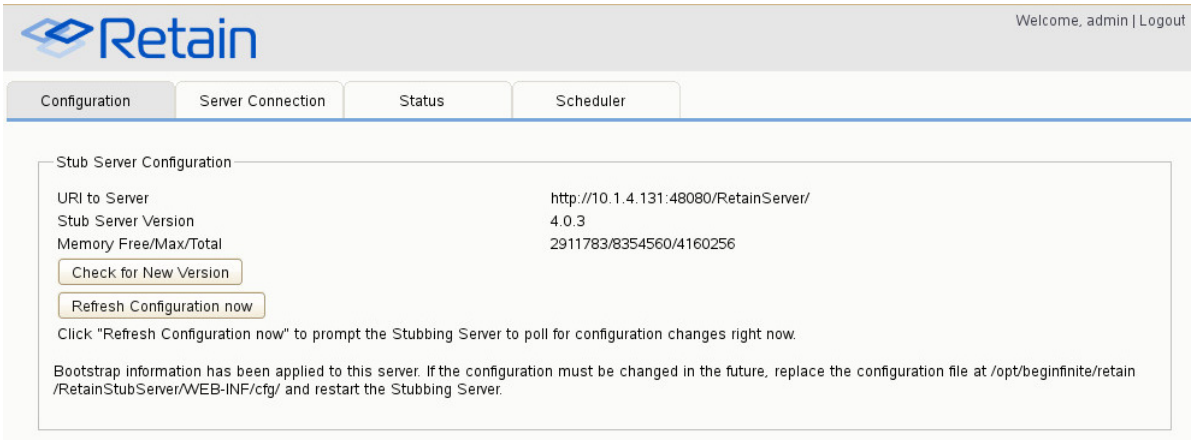
When the Configuration is complete, press Continue.



5.5.7 Stubbing Server

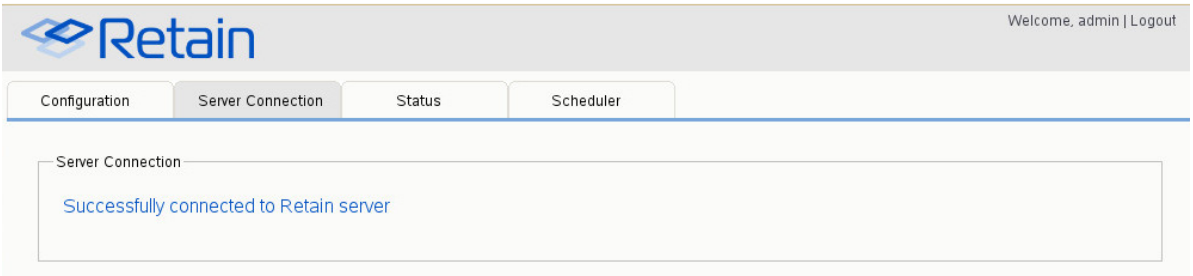
Log into the console with the admin user.

You will then see the Stubbing Server Configuration tab page, which looks much like the Worker page.



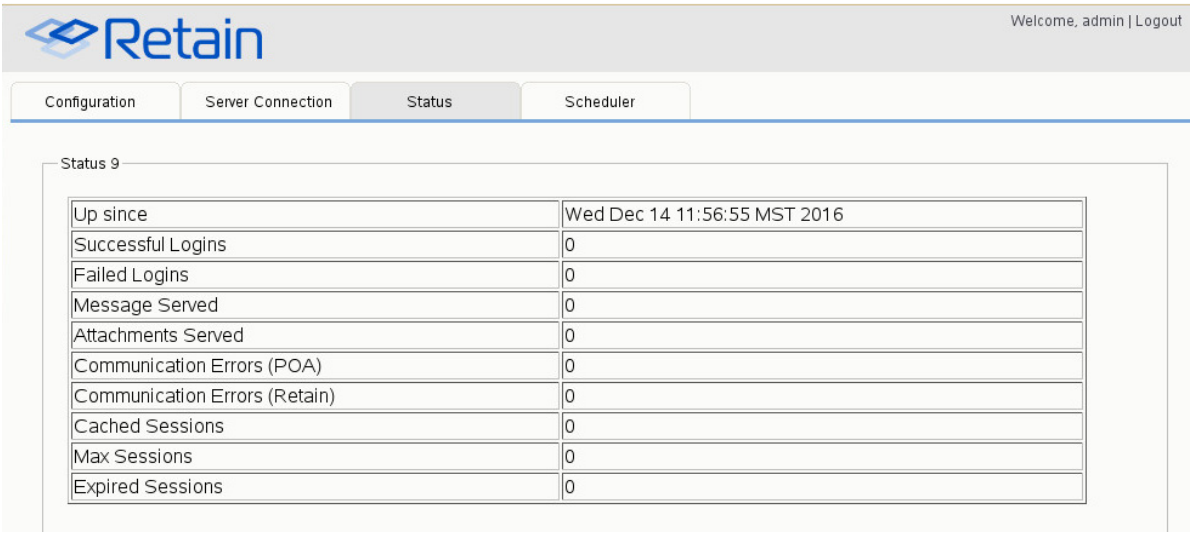
Server Connection

This tab will show you if the Stubbing server was able to connect to the Retain Server.



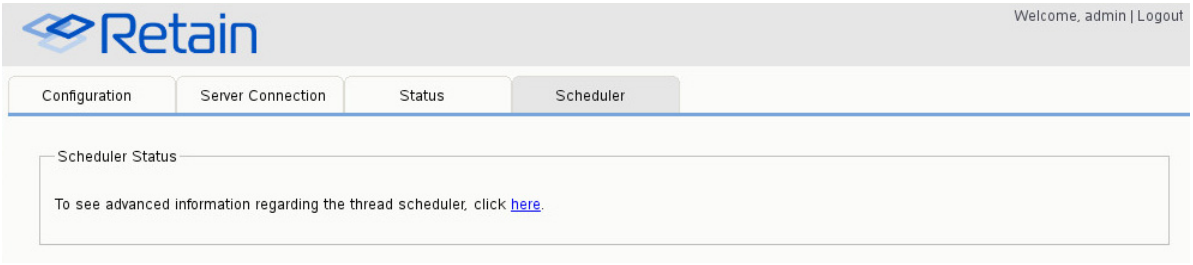
Status

This tab provides statistics about what has been happening with the Stubbing server.



Scheduler

This tab is mainly used by support.



5.6 Performing Retain Searches Within The GroupWise Client

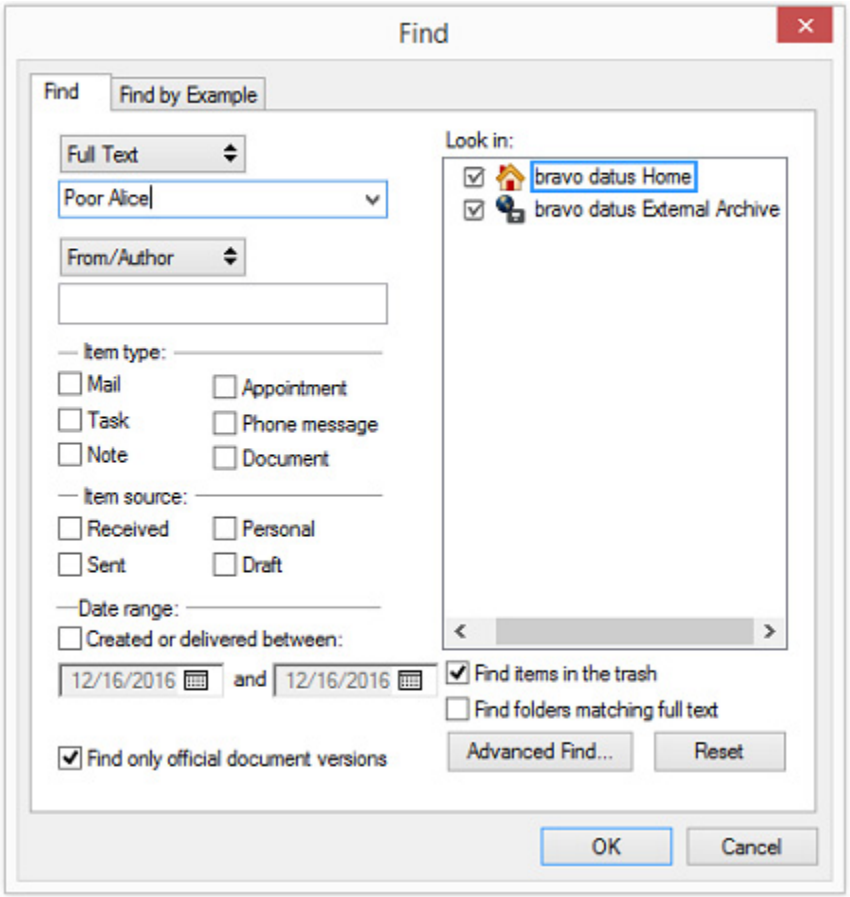
Once GroupWise and the Stubbing server is configured to this point, it is possible to use the Retain Stubbing Server to allow GroupWise to search the Retain archive from the GroupWise client.

5.6.1 Prerequisites

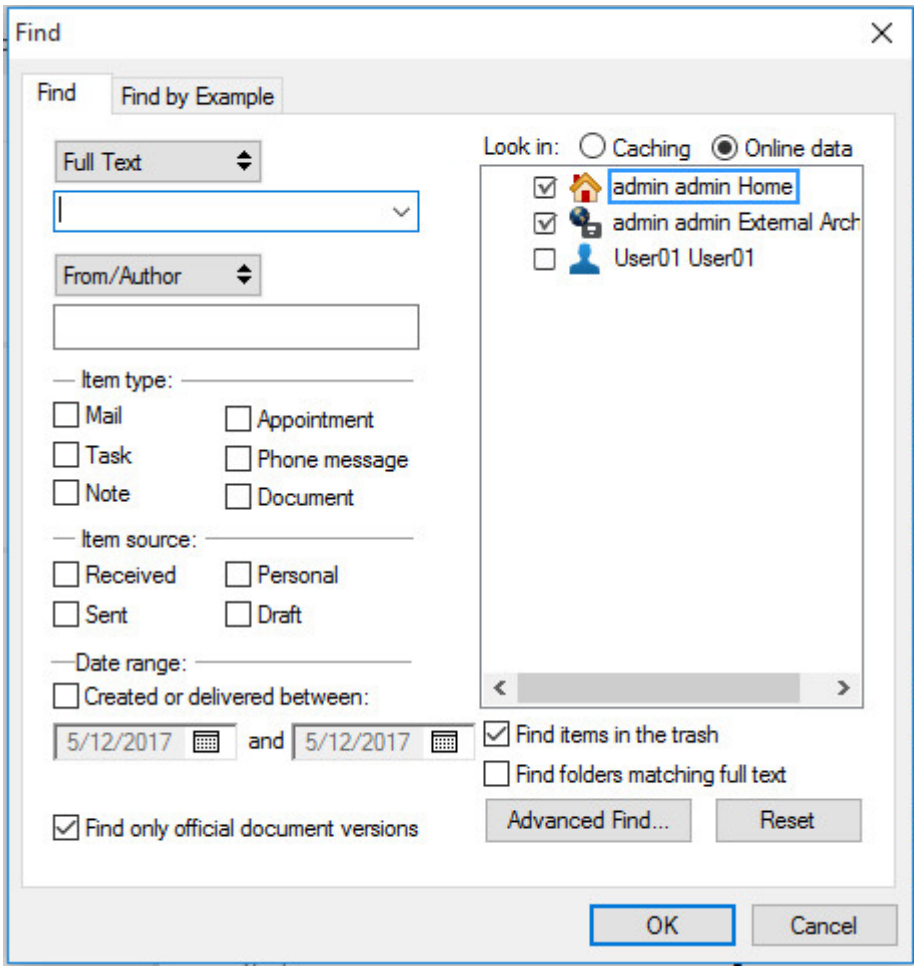
Setup Stubbing Server as above. However, do NOT enable stubbing in the Retain Profile.

In the GroupWise client click on the search magnifying glass icon. Under the Find tab you'll see an External Archive for the user that is the Retain archive.

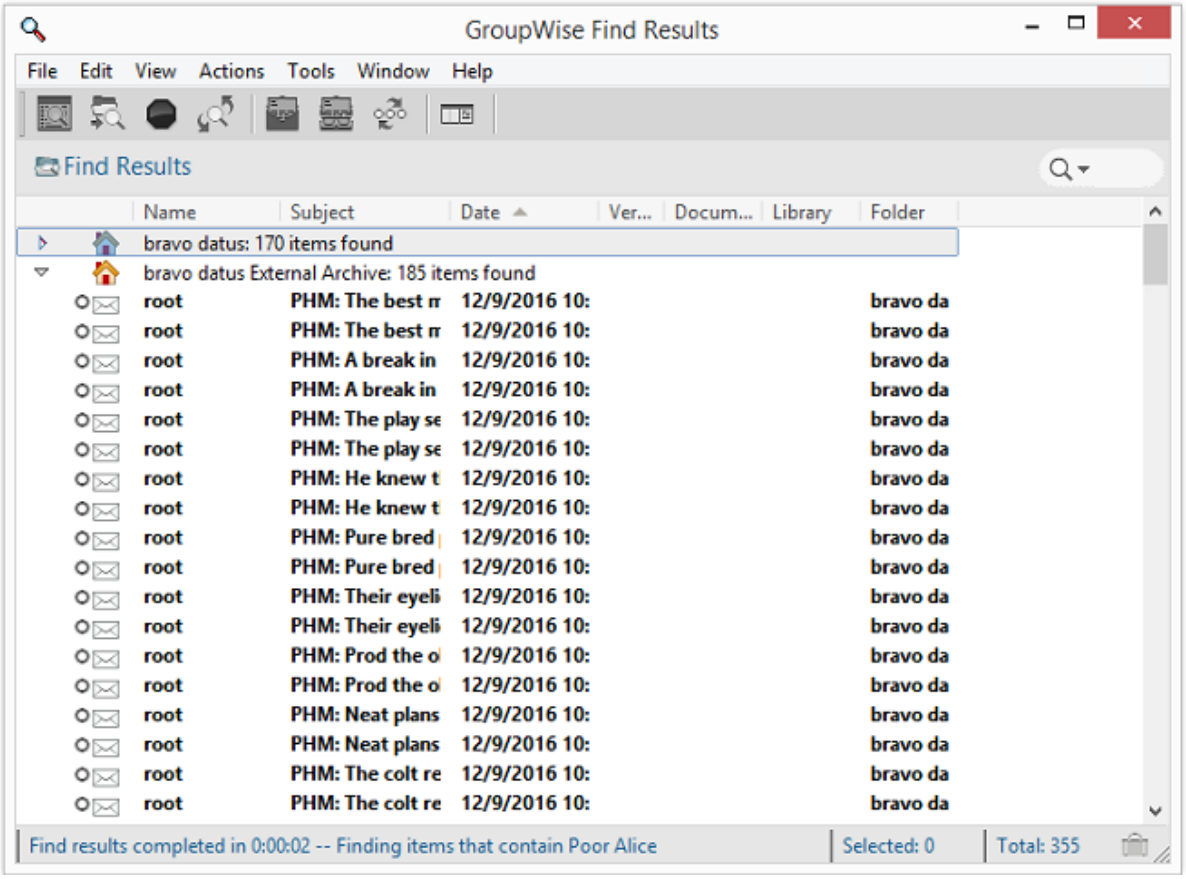
In Online External Archive mode, click the checkbox to enable the *External Archive* search.



In Caching mode, you will also have to set Look in to *Online data*.



Now a search will search not only the GroupWise mailbox but also the Retain mailbox. The Retain ones will show the empty circle stubbing icon for the messages it has.

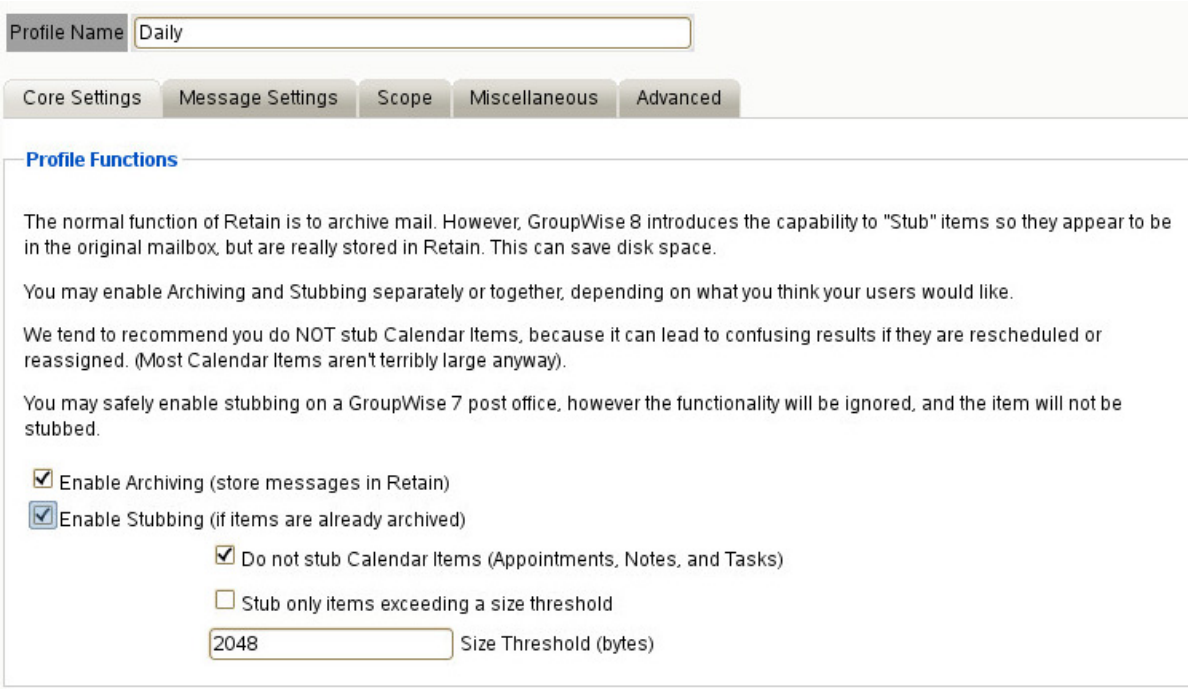


5.7 Configure Stubbing To Remove Messages from GroupWise and Leave Only Stubs

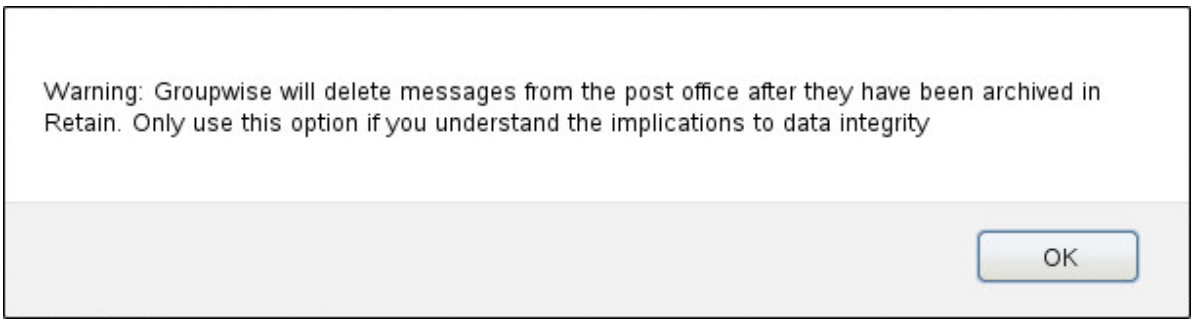
Enabling Stubbing in the Profile will cause Retain to automatically delete items from GroupWise and replace them with a stub pointer to the Retain Archives. This will have effects, for example, you will be able to see the header data in GroupWise but not the message bodies or attachments, on viewing the items in anything other then the desktop client.

Once you start stubbing in the profile you cannot reverse it. We do not recommend this use of stubbing.

Once GroupWise and the Stubbing Server have been configured as above, you need to enable stubbing in the Profile.



Retain will show a popup warning of the effects of stubbing.



WARNING: GroupWise will delete messages from the post office after they have been archived in Retain. ONly use this option if you understand the implications to data integrity.

The next job that runs will delete items from GroupWise and replace them with stubs that point to the Retain Archives. Users will not be able to see anything other then the header data (Subject, to, from, date, &etc.) in WebAcc, or mobile clients. User will be able to read the whole message only in the desktop GroupWise client.

You would want to set the Date Range to Scan to "All Items in Mailbox" because what would be the point of stubbing only new messages.

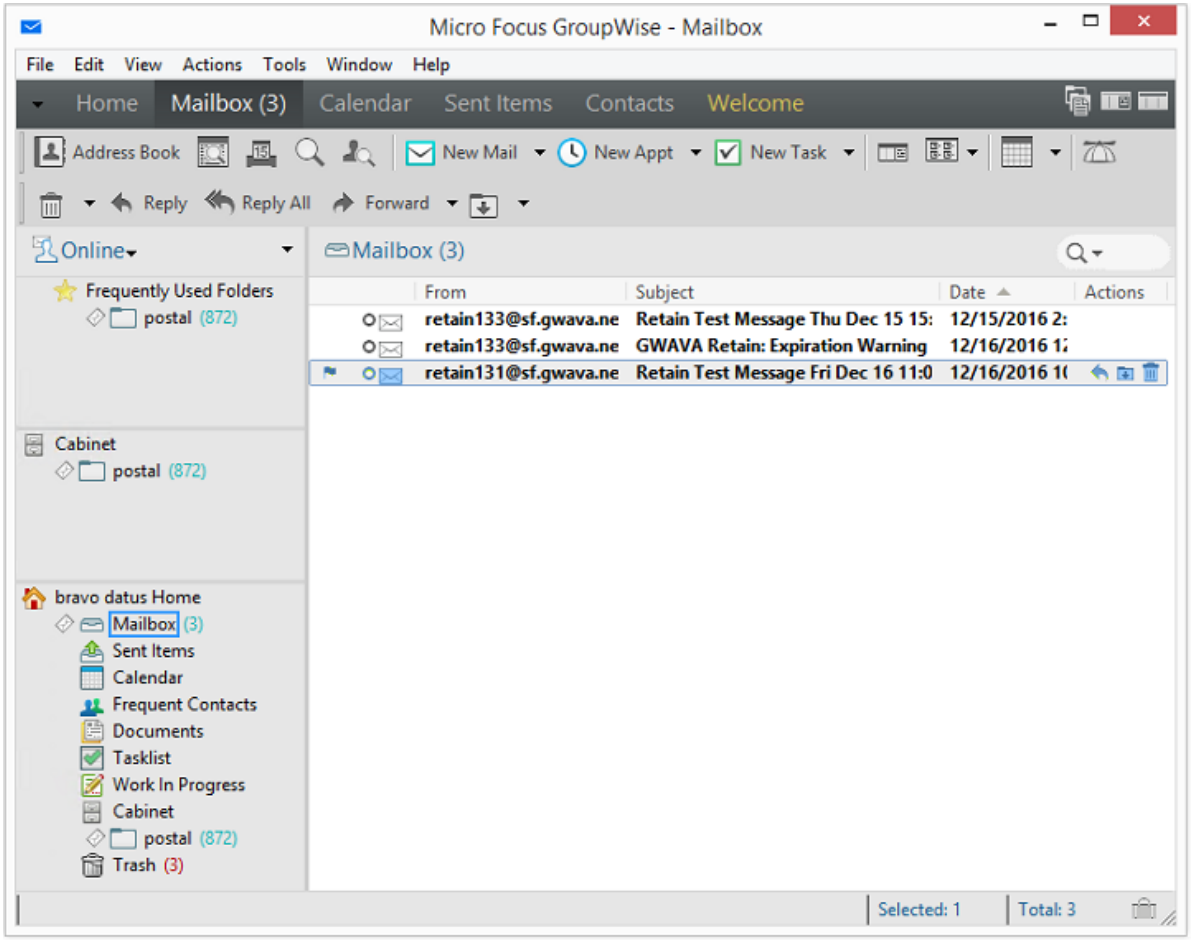
Set up the Job

You can set up the job to use the Stubbing profile and select the Mailboxes to a particular Post Office or a particular User.

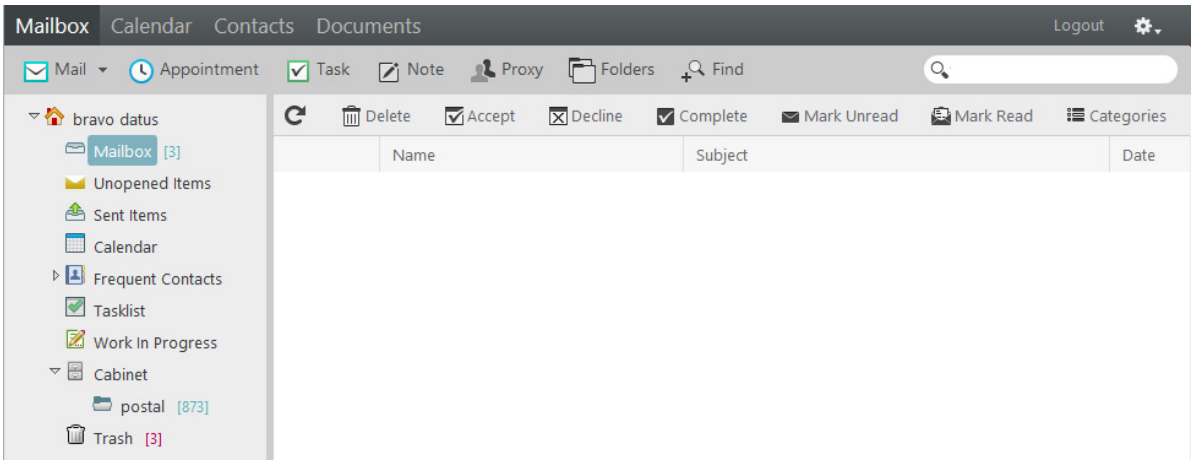
When the worker successfully dredges each item in the mailbox the item will be deleted and stubs placed.

5.7.1 GroupWise Client

In the GroupWise Client, messages that have been stubbed will be marked with an open circle icon. Opening messages may take a moment as the data is retrieved from Retain.



But in web access and third party applications, you'll see nothing:



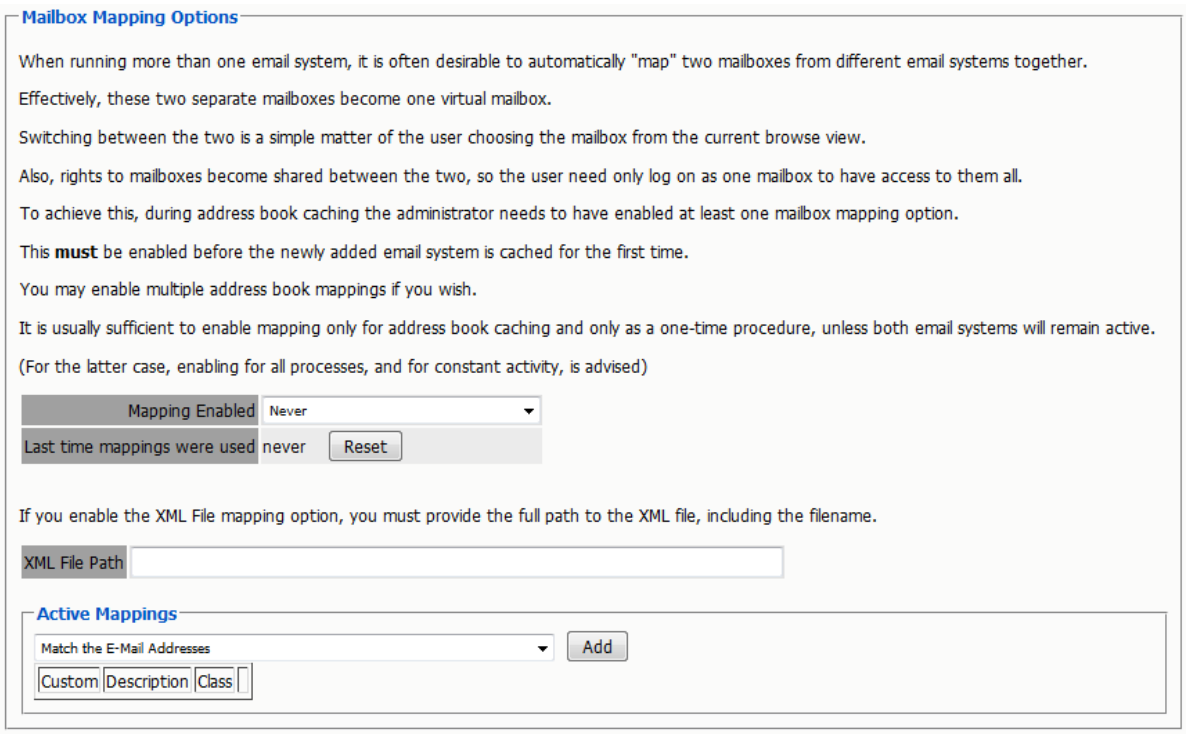
5.8 Mailbox Mapping

5.8.1 Mailbox Mapping Options

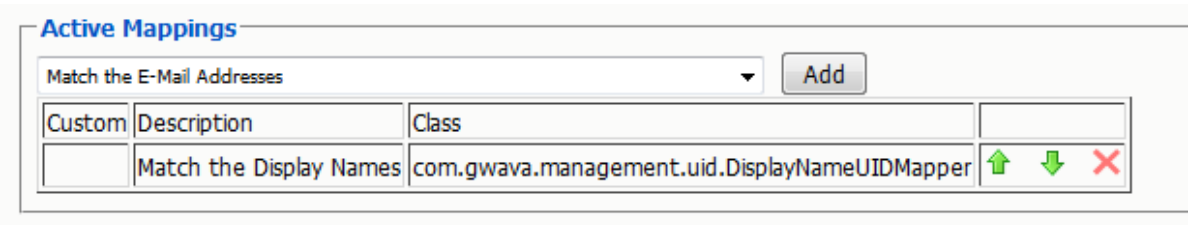
Mailbox mapping automates the merging of users' login names that have both GroupWise and Exchange accounts, combining the GroupWise and Exchange archives under the same user in Retain. Retain connects the internal GUID's of accounts to reflect the same user. After mailbox mapping has been activated, an immediate caching of the address books of both systems is required. Mailbox mapping is not permanent; however, splitting must be performed manually via the Mailbox Merger / Splitter tool found on the 'About' screen. If a refresh of the mailboxes is not performed before users attempt to login or other activ-ity is present on the server, the merge may fail, leaving extra entities created in the Retain archive which must be removed manually. Furthermore, the merge option is a once-off option. Please contact support for further options.

To map one mail system to another, both message systems, Exchange and GroupWise, must be function-ing and have the desired accounts created on the system which allows for identification of the different users.

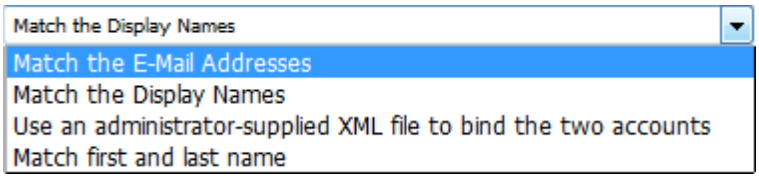
Go to Module Configuration and click the Mailbox Mapping Options button to reveal the tool.



Accounts may be matched via email address, display names, or specified by an administrator supplied XML file.



Most organizations maintain a naming scheme from one system to the next which results in the email address being the same in both systems. For this situation, the email address match is the easiest and quickest way to map the users. For systems, which have listed the first and last name of each user, but have different email addresses, the name matching is the best option. The final option is for an admin-istrator provided xml file to bind the accounts together is time consuming and requires the administrator to build a list of matching addresses with UUID's in both systems.



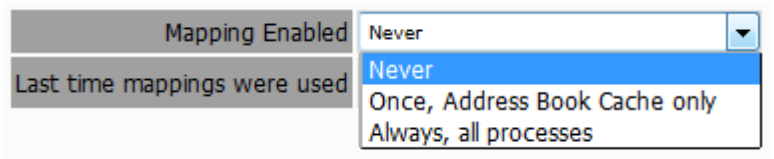
Any mis-mapped addresses may be corrected by using the Mailbox Merger/splitter utility found on the 'About' page. see the 'About and Diagnostics Page' section for more information. A sample XML file is provided and can be download from the 'Mailbox Merger/Splitter' utility page, it is also shown below.

The Address book MUST be refreshed immediately in order to merge the users. The merge will not com-plete unless the Address Book is cached afterwards, and the merge cannot be re-run.

How the merge runs depends on the different users in the local system. If both message systems are to remain active, with users being added and removed from the system, it will be necessary to have mailbox mapping on all the time. Merged addresses are also required for the GroupWise to Exchange migration to function.

5.8.2 Mailbox Mapping Enabled

The settings here are designated for the different situations which a system migration may be in.



5.8.3 Never

This setting disables mailbox mapping, and is to be used in most cases.

5.8.4 Once, Address Book Cache only

To migrate a system, or for systems which have already been switched over, where the old system is no longer being used, but is still running.

5.8.5 Always, all processes

For situations where both mail systems continue to have incoming, new mail, the mapping needs to be completed repeatedly, to keep the message data mapping to be accurate for newly archived mail.

5.8.6 XML File

The supplied XML file is the way for an administrator to manually dictate which addresses are to be mapped or tied together in the Retain system. This process is not recommended as it is time consuming and somewhat complicated. In general, it is likely quicker to use the Mailbox Merger/ Splitter tool from the ‘About’ page instead. Both situations require the administrator to know which email addresses are to be tied together before beginning.

The XML file requires the administrator to specify, in pairs, which addresses are to be tied.

The source address is the address from which the system is being migrated.

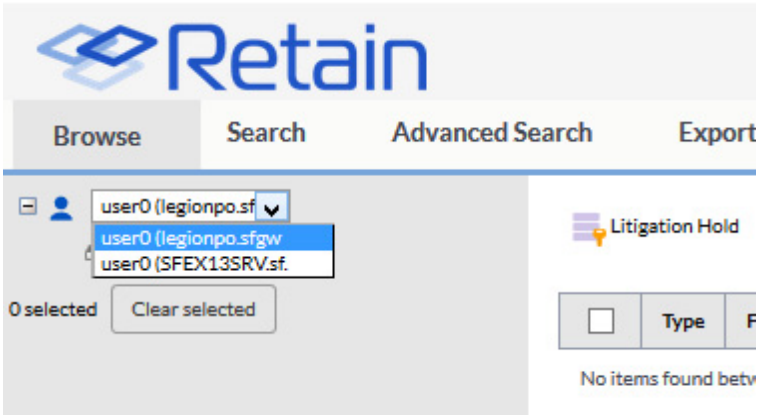
The destination address is the address which the user is the new system, or where the data is being migrated to.

The ‘removeOrphanedEntity’ option tells Retain to remove any unconnected entities left over. Default for the XML file is ‘true’, while the default in the Mailbox Merger/Splitter tool is ‘false’. By rule of thumb, if you are uncertain, don’t remove it.

```
− <MergeTasks>
  − <tasks>
    − <MergeTask>
        <sourceMailbox>source0</sourceMailbox>
        <destinationMailbox>destination0</destinationMailbox>
        <removeOrphanedEntity>true</removeOrphanedEntity>
      </MergeTask>
    − <MergeTask>
        <sourceMailbox>source1</sourceMailbox>
        <destinationMailbox>destination1</destinationMailbox>
        <removeOrphanedEntity>true</removeOrphanedEntity>
      </MergeTask>
    − <MergeTask>
        <sourceMailbox>source2</sourceMailbox>
        <destinationMailbox>destination2</destinationMailbox>
        <removeOrphanedEntity>true</removeOrphanedEntity>
      </MergeTask>
    − <MergeTask>
        <sourceMailbox>source3</sourceMailbox>
        <destinationMailbox>destination3</destinationMailbox>
        <removeOrphanedEntity>true</removeOrphanedEntity>
      </MergeTask>
    − <MergeTask>
        <sourceMailbox>source4</sourceMailbox>
        <destinationMailbox>destination4</destinationMailbox>
        <removeOrphanedEntity>true</removeOrphanedEntity>
      </MergeTask>
  </tasks>
</MergeTasks>
```

5.8.7 Mailbox Mapping View From The User Mailbox

Once mailbox mapping is enabled the user will see a dropdown menu that will allow them to switch easily between their mailboxes.



5.9 Reload Integration

Micro Focus Reload and Micro Focus Retain perform very different functions. Retain is an archiving product whose main feature is the storage of data in one place for later search and retrieval. Reload is a Hot Backup, Quick Restore and Push-Button Disaster Recovery product whose main feature is the storage of instances of GroupWise post offices for the purposes of restoring items to their original location in their original form or providing disaster recovery of domains or post offices.

So, why would you want to integrate Reload and Retain?

- Reload is very good at moving data efficiently from point A to point B.
  - It copies your post office data in its original form.
  - It can make what is effectively a full backup by moving and storing as little as 12% of the total amount of data in the post office.
  - By having the backed up data available in its original form, it can serve as a data source for Retain.
  - Reload's backups are available the moment the backup job is complete.
- Retain moves a lot of data and needs strong network links to do so rapidly.
  - An archiving job moving "everything" will move all of the data. This may seem self evident but when you combine Reload with Retain, you can achieve the same thing by moving only 12% of the data.
  - If you don't integrate them, you will pull data twice over the link – once for Reload, and once for Retain. On top of that, if you don't have Reload and you only have Retain, you will definitely have to move your data twice.
  - By integrating Reload and Retain, you can centralize your archives and ensure good backups and achieve a single data pull.

5.9.1 A Brief Review on How Reload Works

No Helper Software Needed.

Reload runs on a Linux server. It does not use agents or helper software on the source post offices to work. In other words, no agents or TSA's are required. Reload simply connects to the server where a post office or domain is stored and then copies the data to its backup storage location.

Backups are instantly available.

Because the data is copied in its original format, the data becomes available as soon as a backup job is complete by simply running a post office agent (POA) against it (for post office backups) or a message transfer agent (MTA) (for domain backups).

Backups Have Little or no Impact on Users.

Because Reload does not use the Post Office Agent to make backups, there is very little impact on users. The POA will continue to run and service users as normal. Reload also does not use TSA software or helper agents on the live post office server. Backups can be made while the users are logged in and work-ing.

Reload Leverages GroupWise's Architecture to Save Bandwidth.

A GroupWise post office is composed of databases and overflow files. Databases contain users' mailbox layouts and indexes and other databases contain users' authentication information. For any GroupWise item exceeding 2KB in size, such as e-mail with attachments, overflow files are stored, commonly called BLOBs (Binary Large Object).

While the contents of the databases changes almost constantly, the BLOBS are static. Therefore, in a Standard Backup, Reload grabs the databases in their entirety but only those BLOBS that have been newly created since the last backup.

Generally, the BLOBs take up almost 90% of the space occupied by the whole post office. Therefore, with a standard backup, Reload can get away with copying only 12% of the data – the databases and only those BLOBs which have recently been added. For those BLOBs that have been backed up in prior backup instances, Reload links to a master backup directory, taken the first time a backup was run, using a Linux feature called symbolic links. A symbolic link is like a Windows shortcut except that it looks, feels, and acts like the real thing.

This is how Reload can achieve VERY fast backups. In addition, using Reload to move data will save tre-mendously on network bandwidth compared to traditional backup systems which grab all of the data.

Backups Can be Made of Backups Allowing Centralization and Redundancy

Reload was made to backup live post offices and domains and it can also make backups of other Reload backups. The following two cases can help illustrate how useful this can be:

1. Consider client "A" who has two physical locations, one post office in each. This client wants redundant backups – a primary backup plus a secondary in case the primary fails.
2. This client installed a Reload server in each location. The servers backed up the local post office in addi-tion to making a backup of the Reload server in the other location. Thus, each Reload box effectively had backups of both servers.
3. Consider client "B" who has one central data center and four branch offices. This client wants the head office to have backups of all post offices in all locations.

Branch offices 1, 2, and 3 have fast WAN links to head office but branch office 4 has a very weak connec-tion to head office. However, branch office 4 has a strong WAN link to branch office 2.

So the client installed a Reload server in each office and one in the head office. The Reload server in the head office was set to back up the Reload servers in branch offices 1, 2, and 3. For Branch office 4, the Reload server in Branch office 2 was set up to back up the data from the Reload server in branch office 4 and then the head office was set to back this data up from the Reload server in branch office 2.

Thus, backups can make as many hops and can be backed up in as many places as you need.

Using the ability to backup one Reload server with another, you can achieve data centralization and redundancy. The redundancy also gives you the ability to use Reload for off-site disaster recovery.

Additionally, for client “B”, their old backup system moved all of the data every day. Using Reload, they managed to cut their network traffic by 88%.

5.9.2 How Retain Takes Advantage of Reload’s features

Consider client “B” from the earlier example who has four branch offices and a head office. They want their Retain Server to be located at head office. So they need to centralize their data.

Without Reload, they would have Retain Workers on the branch office POA servers and the data would be sent over the WAN links. For a data collection involving “everything”, all data would surely saturate the WAN links.

Plus, their backup/restore software would use the WAN links too, if they were centralizing their backups.

Adding Reload to the mix, they are able to achieve huge bandwidth savings and performance gains.

Reload would be set up to centralize the data to one Reload server in head office, saving immediately 88% of their bandwidth compared to their existing backup/restore system.

Next, A Retain Worker would be set up on the central Reload box to draw data from all backed up post offices.

One Retain Worker can only run one job at a time, so the post offices would be archived one at a time.

5.9.3 Multiple Workers on One Server

It is possible to install more than one Retain Worker on one server but this would double the hardware requirements, Tomcat memory tuning, and is limited to Linux as the platform OS. This option is built into the Linux installer and is activated by using the ‘addworker’ switch to the install command. (For example: ./Retain-Install.sh addworker)

You would normally only add additional workers if you wanted to dredge more than one post office at a time.

On a Reload server, it might not be so time critical to dredge the post offices on it since there is no impact on the end users. On top of that, Reload has a special feature made especially for Retain, a special post office agent that stays up all the time, except to move to the latest backup. This way, it is always available to Retain.

So you will have to decide if it is acceptable to have the post offices dredged one at a time or if you would prefer to dredge many at a time. To do many at a time requires multiple workers.

5.9.4 Timing

It's important to time the data collection on Retain so that the Reload backup will be complete long before the Retain job is scheduled to start. This is set in the schedule section under the Data Collection menu in Retain.

5.9.5 Retain Settings

The three tasks to configure Retain to work with reload are to assign the reload password for the worker, assign the running jobs to use the Reload integration, (this setting is found in the Jobs configuration page in the Retain management console), and configure the Profile to use the Item store flag for duplicate checking.

Enter the management console, and select Jobs from the Data Collection menu.

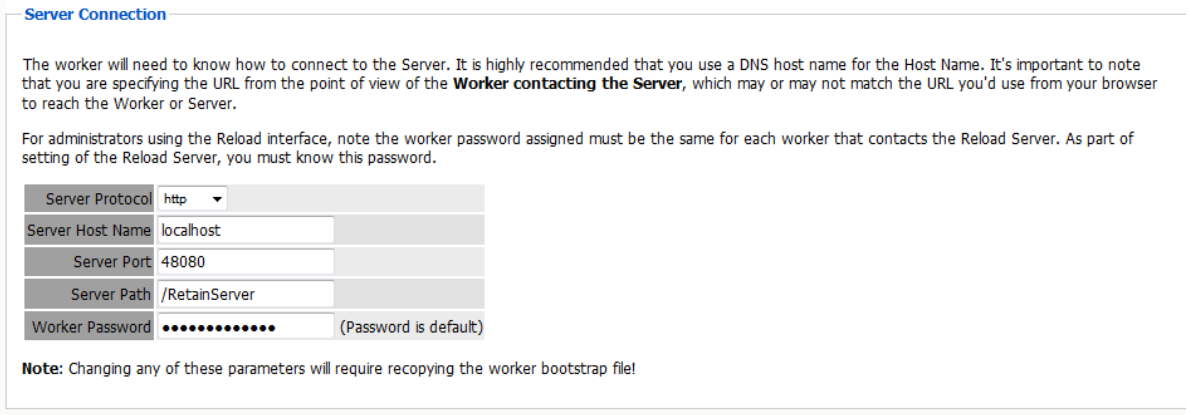


Create or select a job which you desire to use against the Reload system, and select the Reload Tab. You must select the Enable Reload Integration option, as well as supply the correct connection address for the Reload Server URL. (Both IP address and DNS name will work, but DNS is recommended wherever possi-ble.)

Set the rest of the Core Settings, Notification, and Status as you would normally for your Retain system, but note that in the Mailboxes section you MUST assign the mailbox that Reload is backing-up.

Save the changes.

To specify the Reload –Retain password to the worker, open the specified worker in the worker settings page, and click on the Connection tab. Specify the new Worker Password by entering it into the provided field and then click 'Save changes' in the top corner of the page. You must re-upload the bootstrap file to the worker after creating a new password. (See the worker section to get instructions on correcting the bootstrap file.)



5.10 How to Setup Reload to work with Retain

This part assumes that you have already set up and configured one or more Retain workers to collect data from your Reload box.

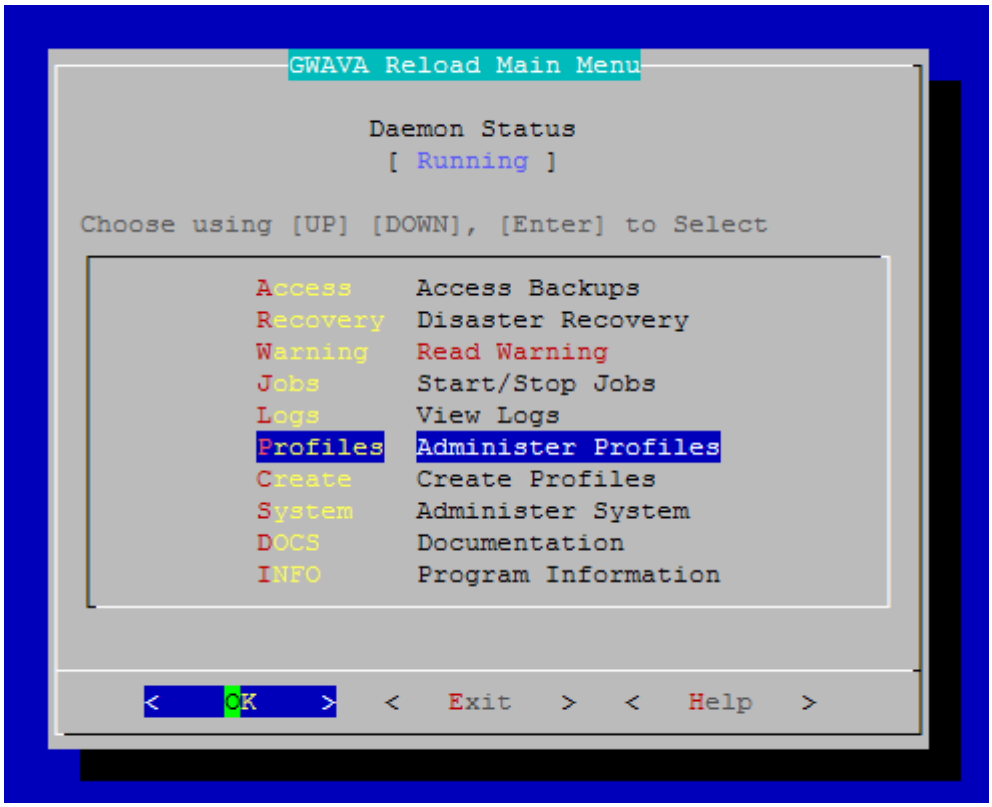
First, Reload must be set up so that the backups are available. There is a special feature in Reload for this. It calls up a post office agent that stays up all the time and it only goes down long enough to change to the most recent backup. So it will always be there with very small interruptions as the POA is brought down then up.

Setting up Reload is done on a Profile-by-profile basis. Each post office that you set up for Retain to dredge from must be configured within the profile configuration menu.

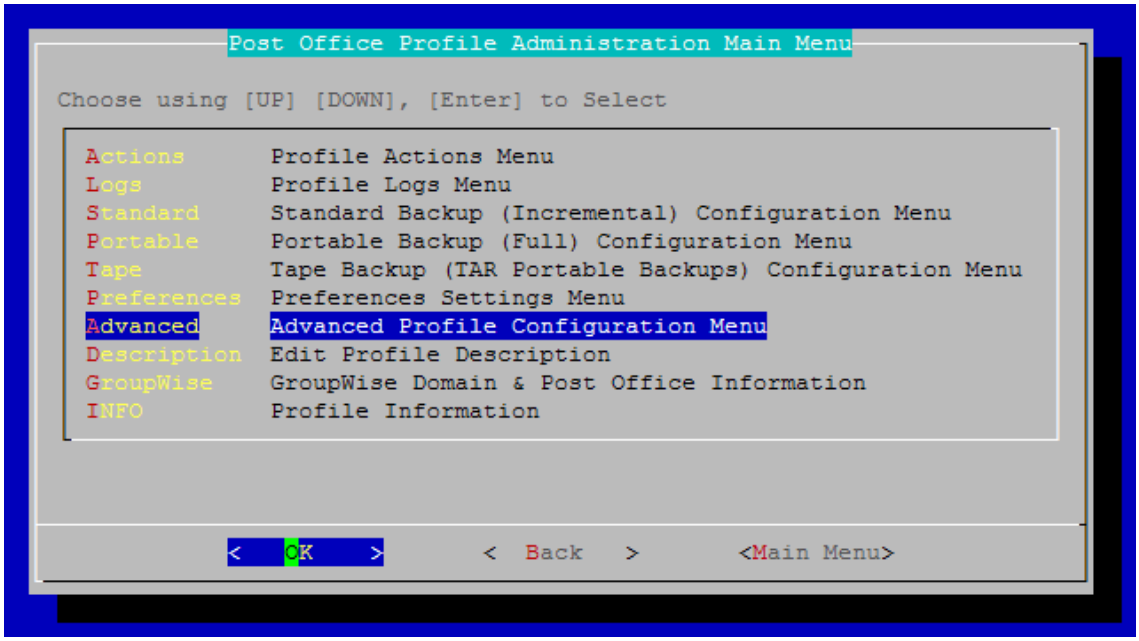
1. Start up Reload's Administration menu.



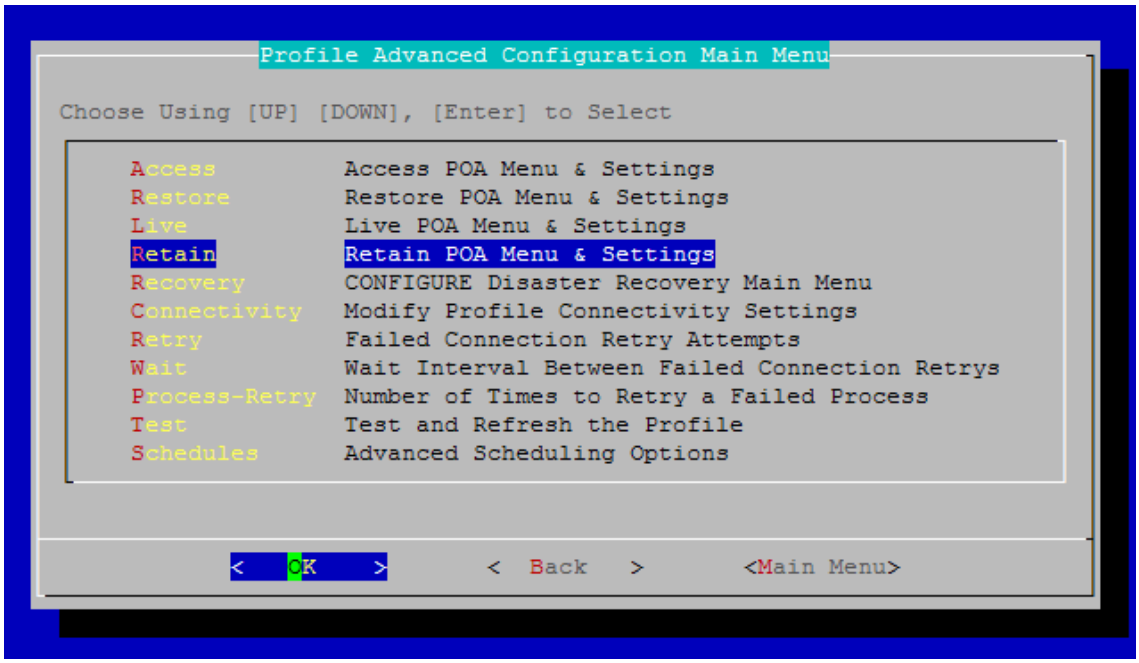
2. From the main menu, choose Profiles – Administer Profiles.

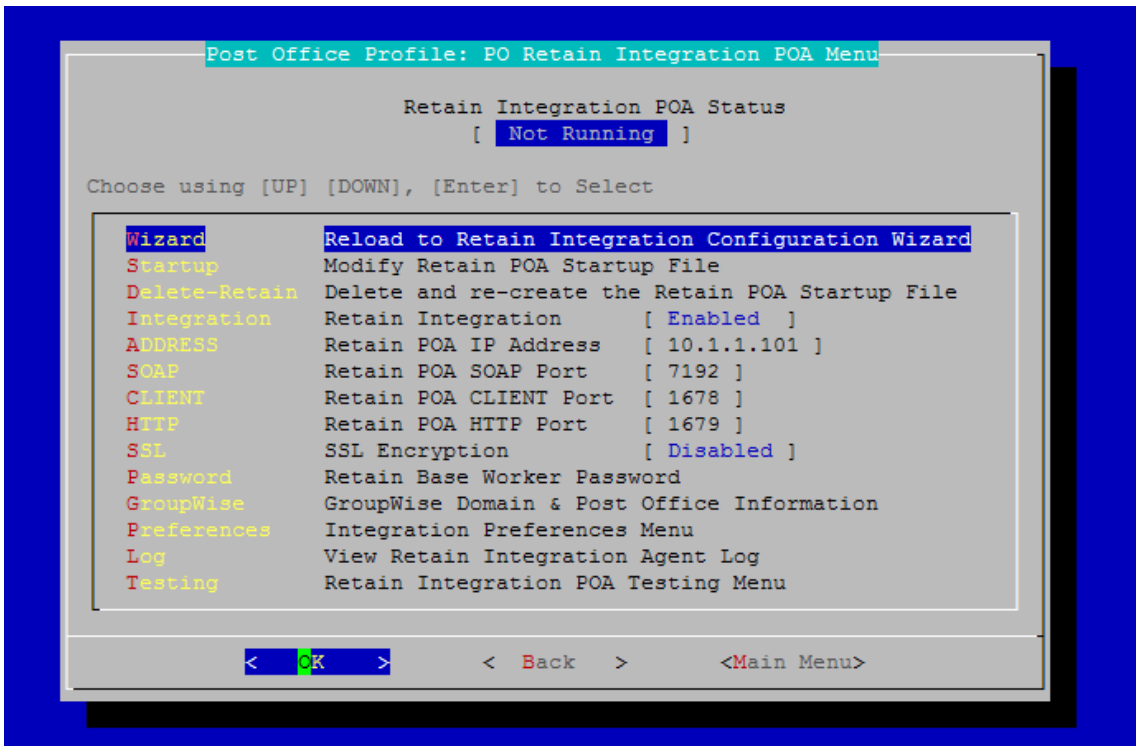


3. Choose Advanced Profile Configuration Menu



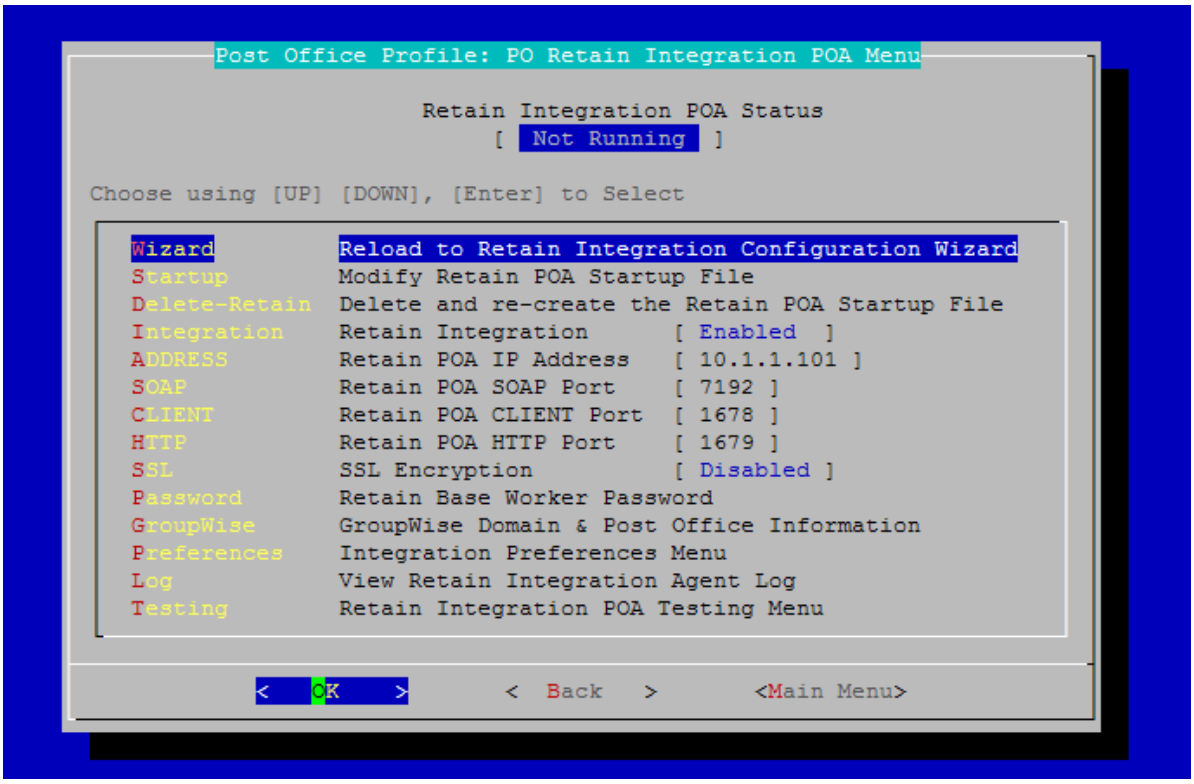
4. Choose Retain POA Menu & Settings





Now it's worthwhile examining this menu. It contains all the settings you will need to make the Retain integration work. This will be a new post office agent running and it will not interfere with the POA being used to access, backup or disaster recovery. Thus, the settings NEED to be different. The easiest way to start is to run the wizard.

Here is another shot of the configuration screen after running the wizard. You will see sample settings and the menu options will be described below.

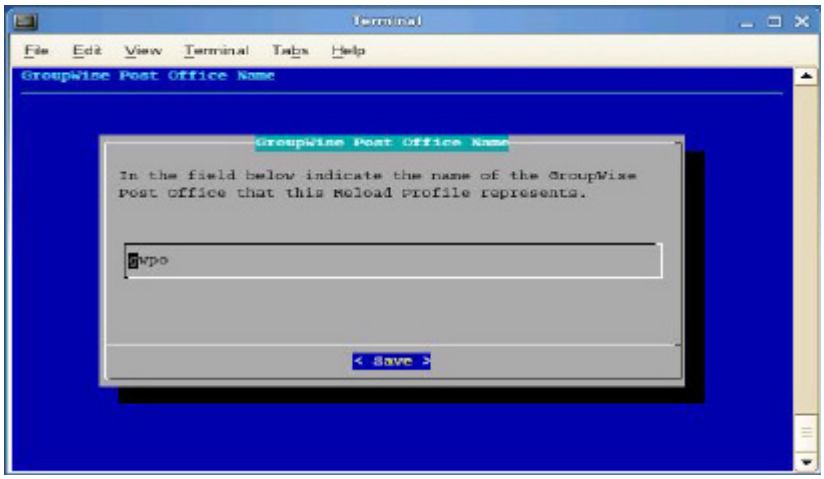


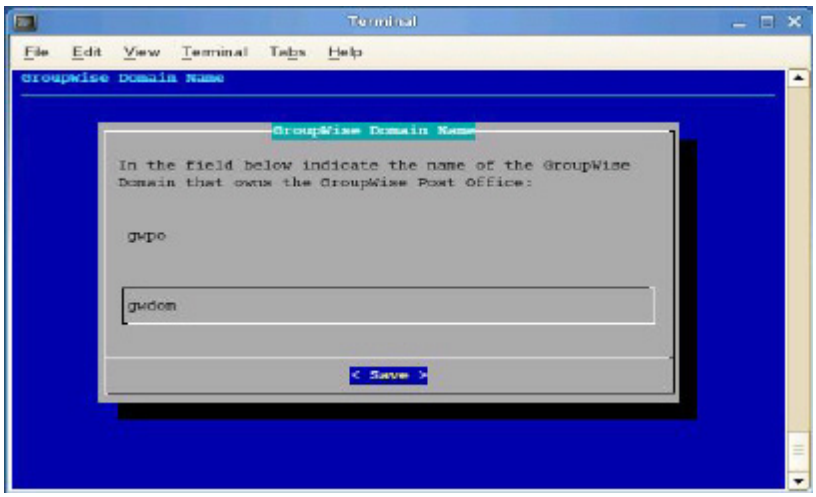
At the top, the status of the Retain Integration POA is displayed.

- Wizard: Run the configuration wizard.
- Startup: modify the startup file for the POA if you want to make specific changes to it.
- Delete-Retain: delete the startup file if you want to start fresh and configure from default.
- Integration: Enable or disable the Retain Integration
- Address: The IP address this POA will listen on.
- SOAP: The SOAP port this POA will use.
- CLIENT: The port that a GroupWise client may use to access this POA.
- HTTP: The HTTP port for this POA.
- SSL: Enable or disable SSL (Generally keep SSL Disabled)
- Key: A password Retain will use to access this POA.
- GroupWise: Specify the domain name and post office name for this POA.
- Log: View the Integration Agent Log.

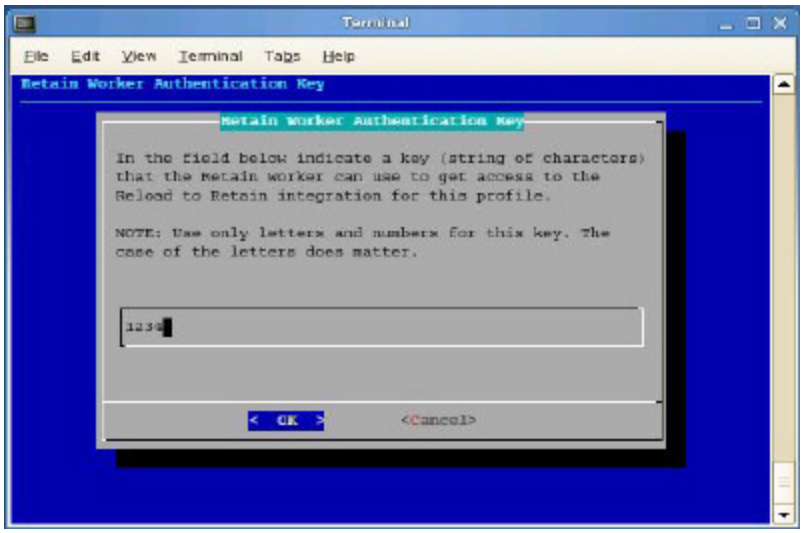
The wizard will be shown below.

1. Run the Wizard

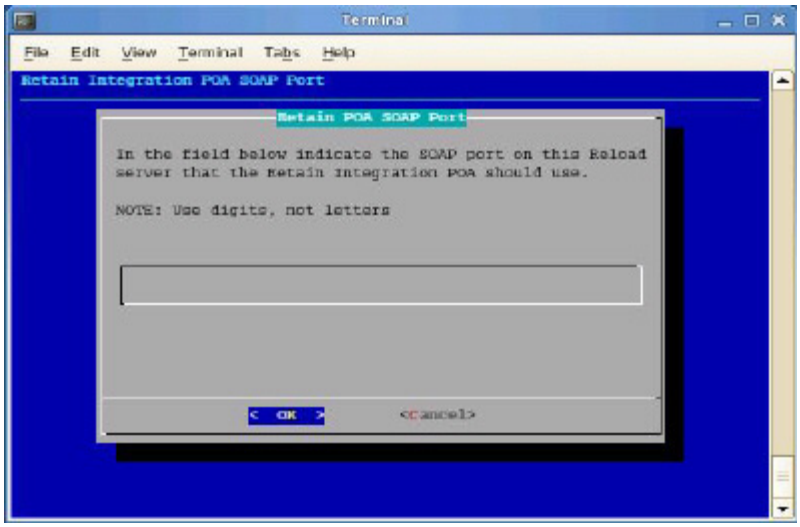




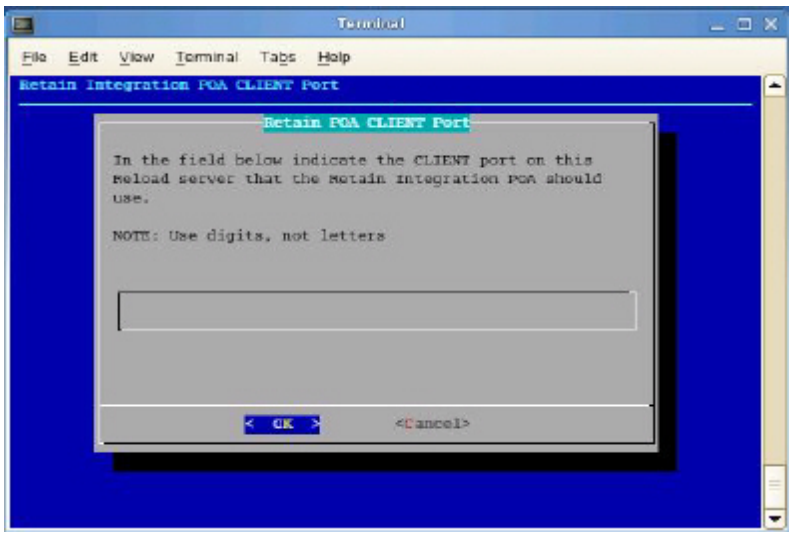
2. Enter the name of the post office and domain.



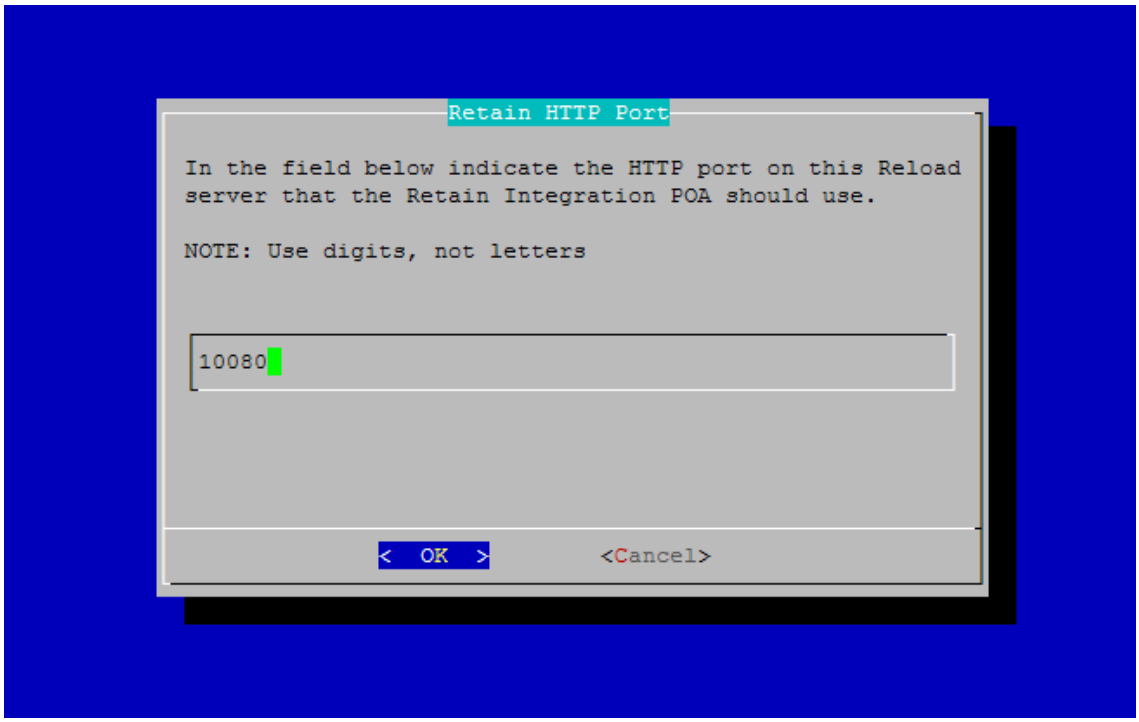
3. Choose an authentication key that Retain will use to access this POA. This must match the password you assigned to the Retain Worker.



4. Specify the IP address and SOAP port for this POA. Be sure it is unique. Some will choose one IP for the whole box with different client ports and SOAP ports for each POA. Others use the same ports but the IP addresses are different.



5. Choose the HTTP port for this POA.



Because Reload is creating a faux POA for Retain to archive mail from, the Reload POA must be on a dif-ferent or unique port, so there is no conflict with your original POA. If your Reload installation is on a sep-arate machine from your POA, any port will do, but if it is the same, pick a port that you know is open, different from the live system.

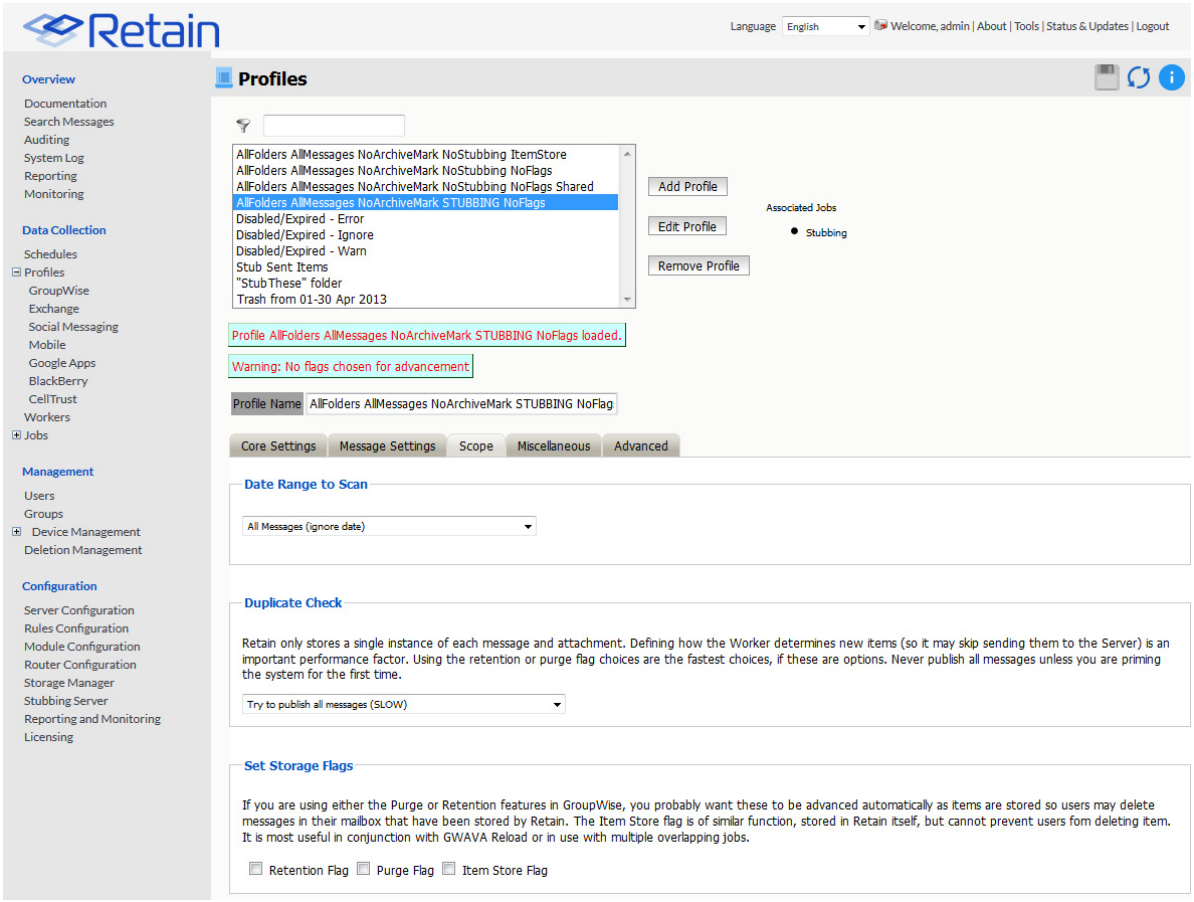
Retain will pull all necessary connection information from the Reload server. There is no need to enter these settings into the Retain Server.

Now that you have set up the basics, you may edit the POA startup file in case you wish you change any other settings, (Retain.poa), or you can re-run the wizard from step 1.

5.10.1 IMPORTANT Notes for the Integration

Retain

Because Reload essentially creates a snapshot of the Post Office, the duplicate checks that Retain can use are very limited. The retention flag and purge flag will not function as they are kept within GroupWise and would be changed back as soon as Reload creates a new backup. The Item Store Flag is the only duplicate check that is internal to Retain, and is the ONLY duplicate check ability that will work when Retain archives against a Reload system. Again, the retention and purge flags will not work but the item store flag will. Be sure your Retain Profile matches this setting.



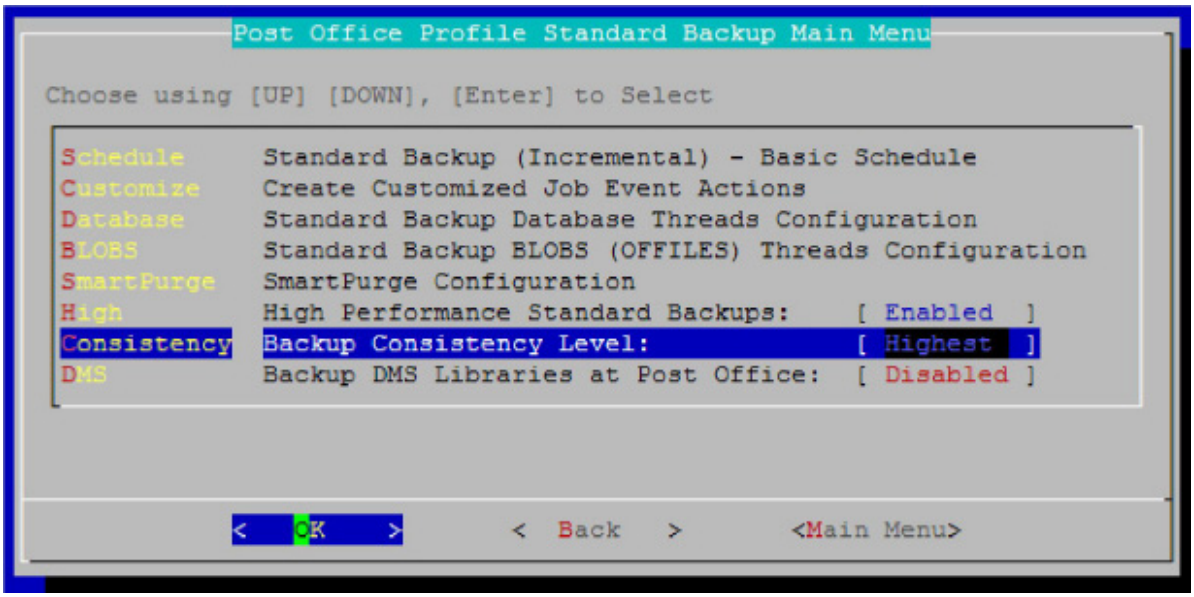
The item store flag is set in two places: Duplicate Check under the Scope tab and under Set Storage Flags under the Miscellaneous tab. The correct settings are shown.

Reload

To mitigate the chances of getting Retain Worker archive errors while working against a Reload POA, it is STRONGLY recommended that Reload is set to create highly consistent backups.

This setting is located at:

Main menu > Profiles (Administer Profiles) >Standard (Standard backup (incremental) Configuration Pro-file) > Consistency (Backup Consistency level): Set to highest.



This is enabled by default for new installs of Reload 2.5, but will have to be manually enabled on systems that are upgraded to Reload 2.5 You want a highly consistent backup, to make sure that you have all the blobs associated with the database. Database is picked-up first, so the blobs that are referenced in the database will be consistent with the current backup.

6 Management

6.1 Users

User and Rights Management in Retain include:

- Creating, deleting, and editing users
- Allowing new user accounts, and restricting specific ones from being created
- User expiration
- Assigning users to groups, to conveniently grant rights or set initial settings on a multiple user basis
- Granting access to mailboxes other than the user's personal mailbox
- Changing the specific functions the user can perform

To access User and Groups Management, the user logging in needs the “Manage users and groups” administrative right.

6.1.1 Creating Users

The primary purpose of a user account is to store their preferences, rights, mailboxes to which they have access, and authentication information.

Users come from one of two places

- They may be valid message system users logging in with their credentials
  - These users use SOAP authentication for GroupWise and Active Directory authentication for Exchange; Retain checks their login credentials with GroupWise or Exchange
  - These users initially belong to the group default. You may change this later
  - You may restrict users (prevent them from logging in)
- Users may be specially created in Retain independently of any message system
  - Users created in Retain do not need to have a message system account
  - Users who don’t exist in the message system will use the offline password

6.1.2 Offline Password

There are, however, occasions when you might want someone to search through the Retain archives, but who is not part of the mail system. Such a person might be an independent auditor, a lawyer, a user deleted from the live system, etc.

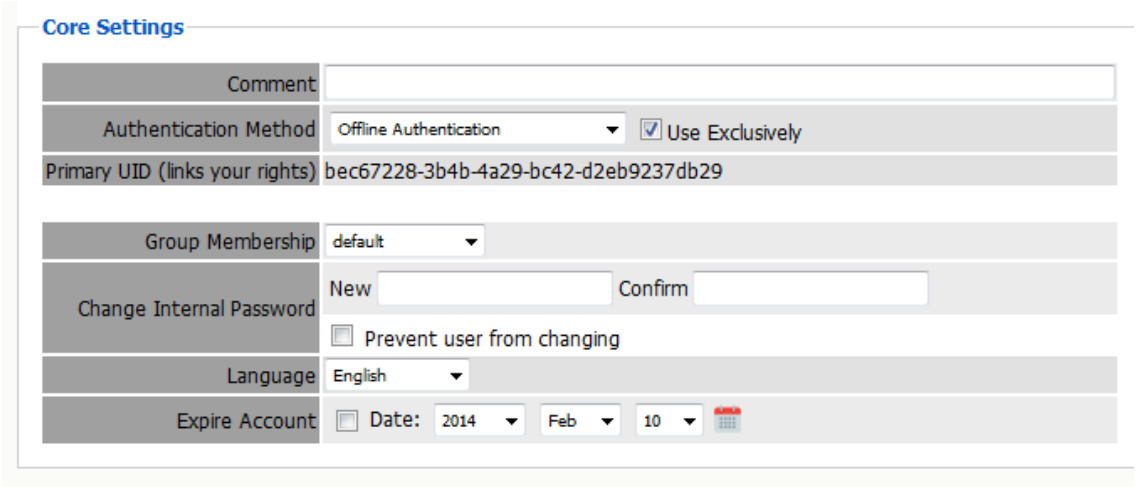
For this reason, Retain has an *offline password* system. These passwords are stored in Retain’s control database. Retain does not care how a user authenticates: whether offline, via SOAP for GroupWise, Exchange, or LDAP, the same rights can be assigned. An administrator who possesses the Manage Users and Groups' administrative right can assign all the rights they contain.

Users may also be assigned access to more than one mailbox. Offline users will need to be given access to at least one mailbox to perform searches. Users who are assigned “Search All Mailboxes” rights have access to all users’ mailboxes.

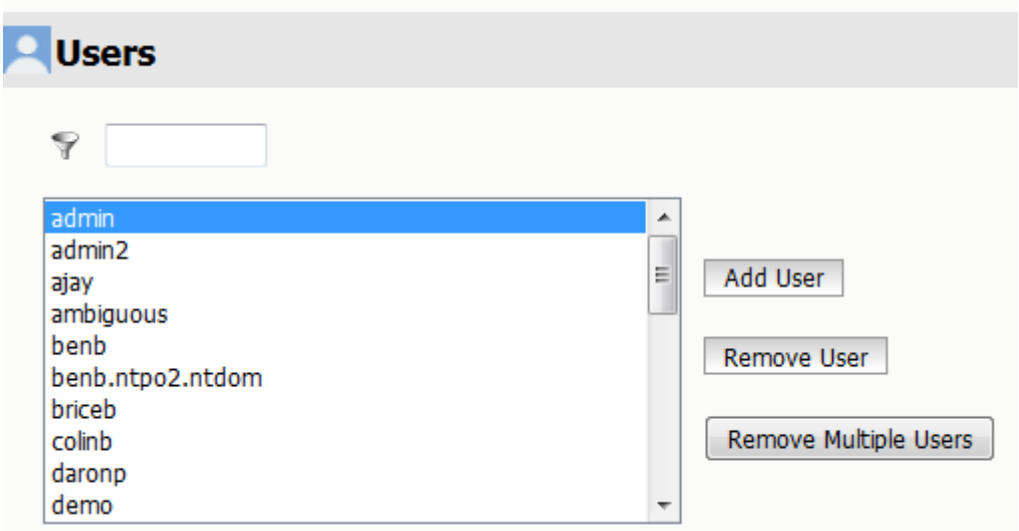
GroupWise Proxy support only works for users who authenticate via GroupWise SOAP protocol.

6.1.3 Creating a new user

1. Begin by clicking on the “Add User” button



2. Enter a new user name and then fill out the options under each tab
3. When you are done, click the save changes disk icon at the upper right



All previously created users are listed and can be edited or removed, by clicking on the user and editing the values, or through clicking on the Remove User button.

6.1.4 Core Settings

Authentication method

- SOAP (for GroupWise users)
  - oGroupWise users logging in are authenticated using SOAP
  - oThese users are automatically entered into Retain's user list
- Exchange (SOAP for Exchange). Exchange users are authenticated via SOAP for Exchange and users are added into Retain's user list
- Offline Password (credentials stored within Retain, any type of user)
- Google IMAP. Google users are authenticated through IMAP to the Google system.
- LDAP authentication – The LDAP Authentication must be setup under Server Configuration | accounts tab
- You may lock a user account so it can only use one type of authentication
  - oIf "exclusive" is not checked, it will try one, then the other

Primary UID

- Offline only accounts usually won't have one
- **offline:** The initial admin account is set to use offline exclusively, so it never has one

Group Membership

- Default is "default"
- Create groups under "Groups" and they will appear as choices here
- Users may belong to one and only one group. From an assigned group, additional rights, mailboxes, and initial settings may be inherited

Account Expiration

- Check this box if this user's account should not expire. Useful for administrator accounts
- By default, accounts never expire, (0=never), but can be changed in Server Configuration

Offline Password

- If you use this authentication method, store the password here
- May be changed as needed
- You can prevent the user from changing it themselves
- Passwords are always stored in an encrypted format – never in clear text

Default Language

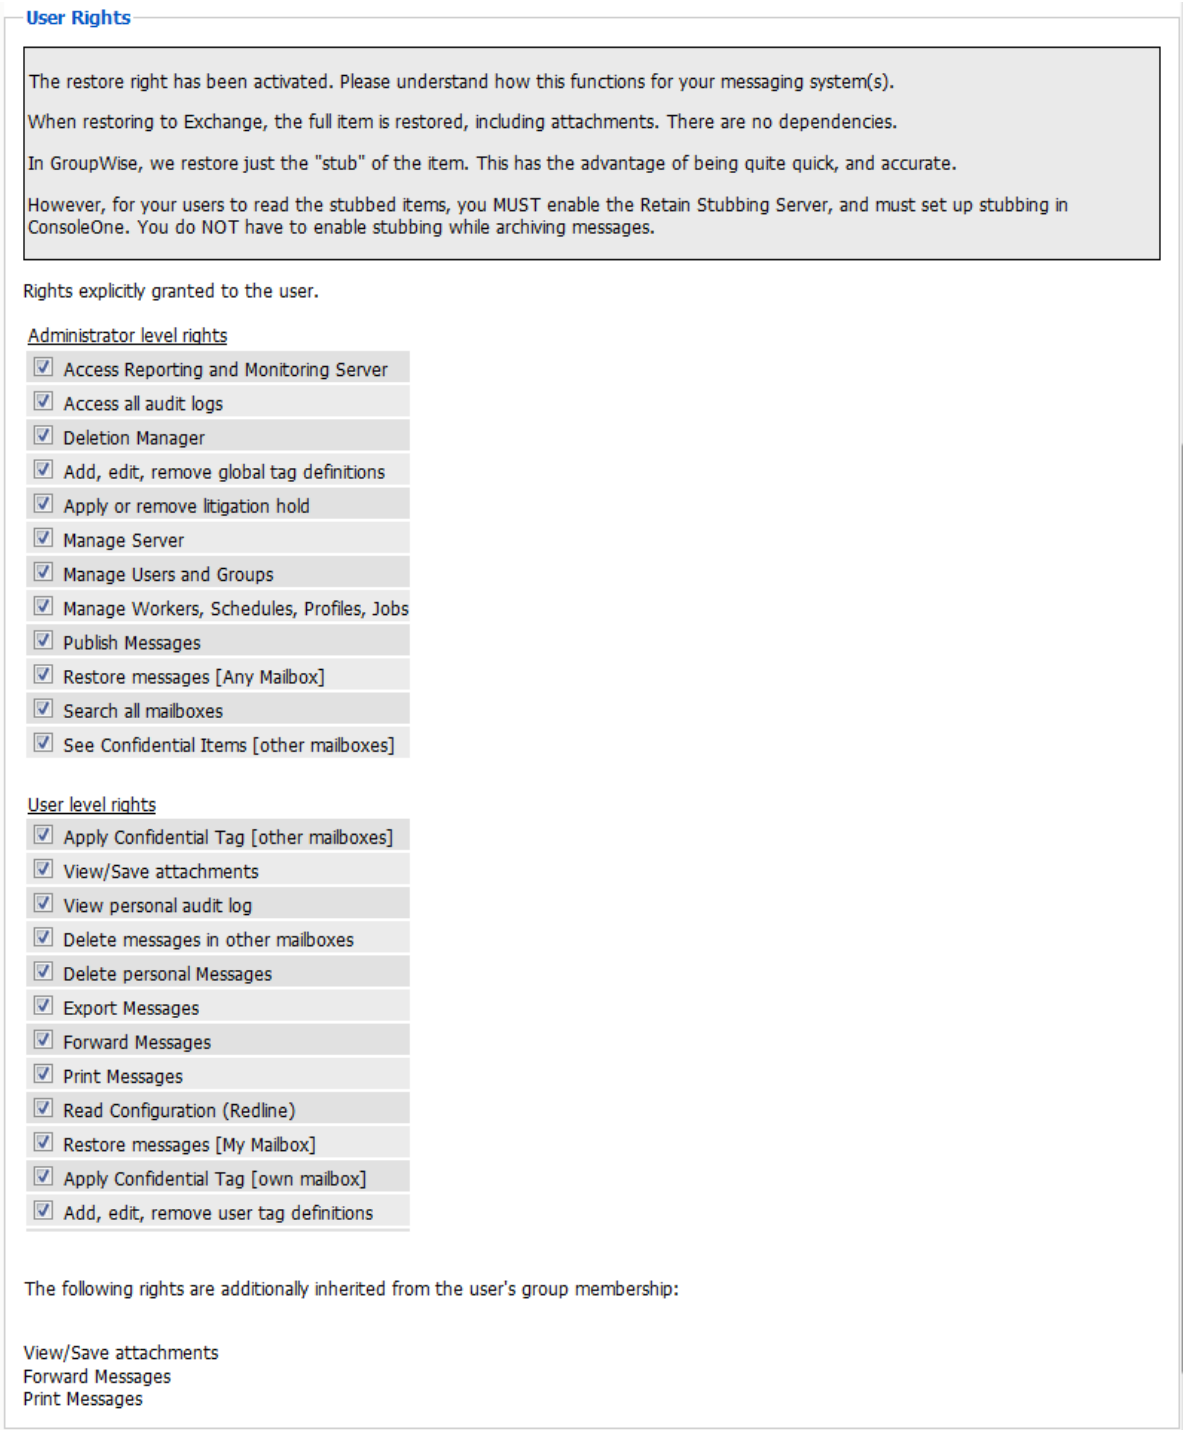
- Choose which language will be used in the Search Interface for this user

Disable account

- This allows the admin to pick a date when the account will no longer be allowed to login, but the account will not be deleted

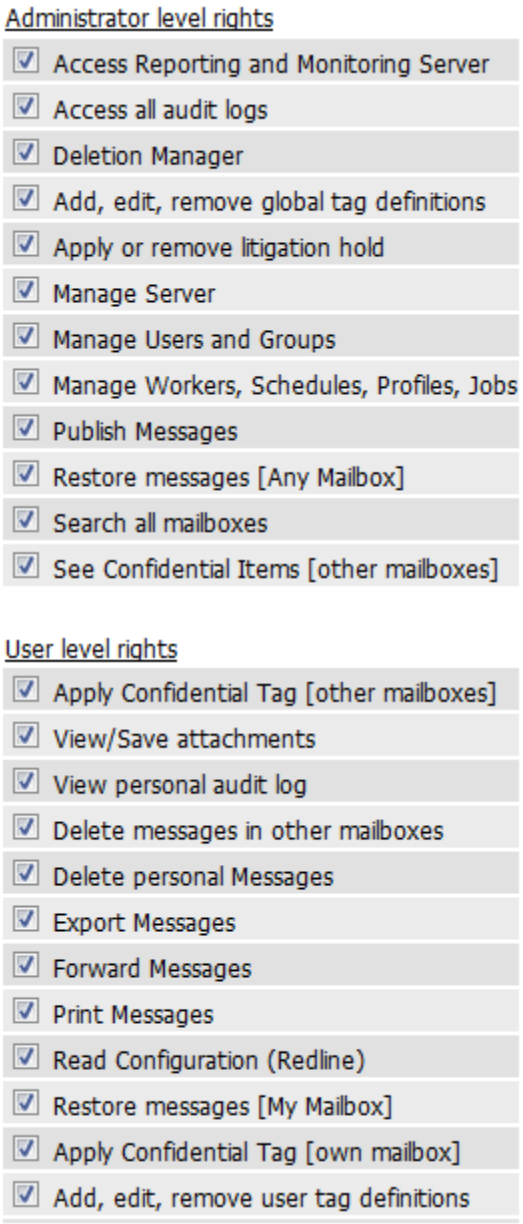
6.1.5 User Rights

When an administrator level right is granted a user, that user will see that right in the management con-sole when they log into Retain. If a right that the full Administrator can view is missing from the menu of that user, they are missing that right. To view and have access to that option, they must have the missing right granted to that user. If you have performed an upgrade and are missing options, check for a missing administrator right.



- Control what rights you grant to the user here. Check the box to enable the right

- These are extra rights
  - You don't need any of them for the user to access their mailboxes
  - You do need them to do “special things”. The first admin account gets them all
- Retain first checks your assigned group and you start with the Group Rights
- The rights you explicitly set here are added to the group rights for the user’s effective rights
- This way, you can control users as a group and give different rights to different groups
- If you don’t have rights to an administrative option, it won’t appear on the left



It should be clear from this screen that there is no such thing as an Administrator per se in Retain. Instead, some users simply have more rights to do more things than others. A distinction is made between Administrator level rights (which allow a user global system wide power) and User level

rights, but any user can have zero or more rights in either category. The Administrator you created in the setup wizard was simply a user account with all of the Administrator level rights granted by default.

6.1.6 User Rights Summary

Administrator-level rights

Access Audit Logs

- The Audit Logs become visible
  - Main menu audit log
  - Message properties audit log
- You can see user activity on two levels
  - Globally
  - On an individual message basis
- Deletion Manager
  - Allows management of Deletion jobs
- Add, Edit, Remove global tag definitions
  - Allows manipulation of global tags in the view messages interface
- Litigation hold
  - Allows excluding mail from Deletion
- Manage Server
  - Access Server Configuration and Licensing
  - Access diagnostic utilities
- Manage Users and Groups
  - Access Users and Groups
- Manage Workers, Profiles, Jobs, Schedules
  - Access data collection system
  - Login to Worker web interface
- Allows use of Publisher
- Allows messages to be restored to all mailboxes
- Search All Mailboxes
  - Grants access to all mailboxes in the system
  - Allows the user to search and browse them
- See Confidential Items [Other Mailboxes]
  - Allows users to view items which others have tagged as confidential

---

NOTE: Only users with administrative rights will see the administrator's screen on login. Non-admin users are simply forwarded to the Search Interface.

6.1.7 User-level rights

All user level rights are strictly optional, and add additional functionality.

- None are needed to access your own mailbox and other mailboxes assigned to you
- You may wish to grant Forwarding, View Attachment, and Printing rights
- The initial “default” group created upon installation does this

---

NOTE: There is no way to perfectly block printing in a web browser, so using this feature should not be taken as a 100% guarantee that users won’t be able to print. Nonetheless, for most users, it is effective

---

These rights are self-explanatory:

- View/Save Attachments
- View own audit log
- Forward Messages
- Restore Messages (My mailbox only)
- Print Messages

Other rights:

- Delete messages in other mailboxes
  - Allows access to delete items from mailboxes that the user is granted access to, but is not their personal account
- Delete personal Messages
  - Allows deletion of items from the user’s personal account
- Export Messages
  - Grants the right to use the export option which sends groups of messages or items as a single PDF document.
- Add, Edit, remove user tag definitions
  - This allows the user to create, remove, and edit the user’s tags in the search messages inter-face.

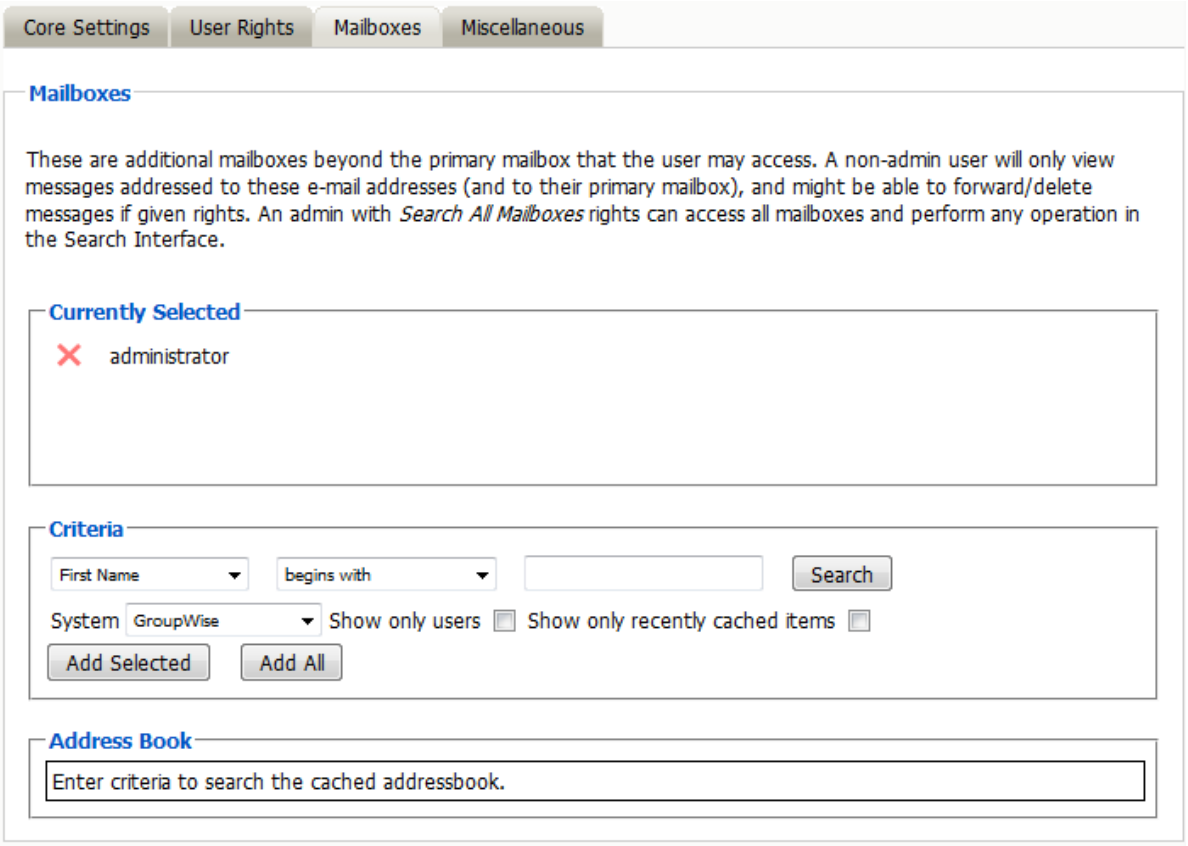
The Read Configuration right (Redline Integration)

If you are integrating with GWAVA’s Redline monitoring product, you will need to create a user account so that Redline can log in and retrieve monitoring information. We recommend the following settings:

- Account Never Expires
- Offline Password Authentication is required. (use exclusively) (be sure to set the password)
- Read Configuration (Redline) right.

Mailboxes

Select the mailboxes this user will be able to access in addition to their own.



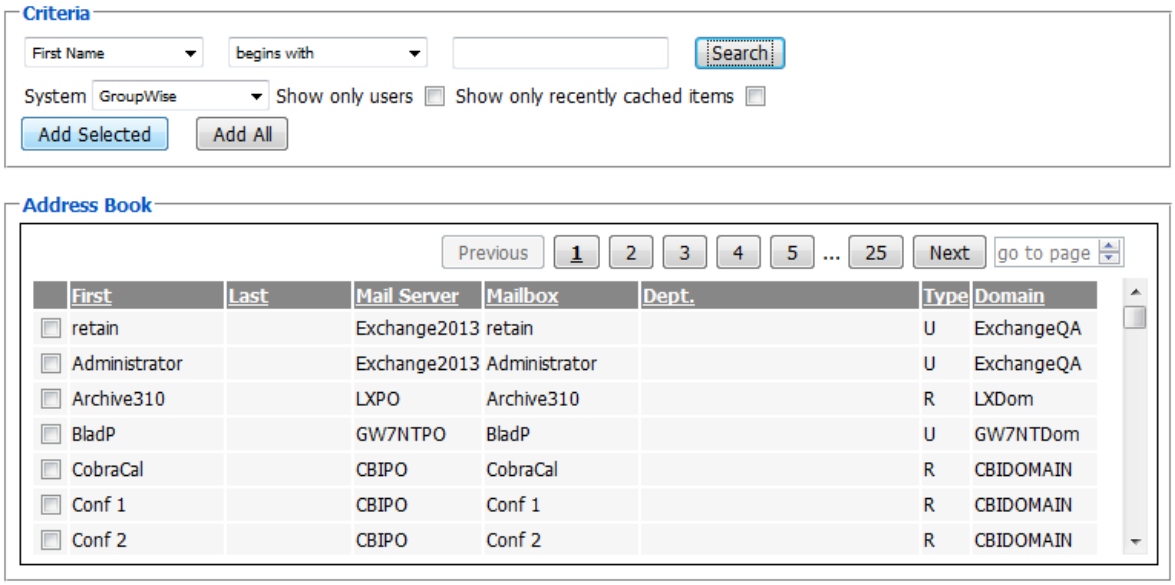
Often, you will want some users to be able to search through more than just their own mailbox. Adminis-trators can have “Search All Mailboxes” as a right which gives them access to everything. This section allows you to give a far more selective range of mailboxes to a user for searching.

In the example above, the user has explicit rights to two mailboxes. These mailboxes can be taken away from the user simply by clicking on the red ‘X’.

Adding users to the list is done using the Address Book selector. In the criteria section, you may enter information to search for a mailbox or a set of mailboxes. The search results will appear in the Address Book section. Each listed entry has a check box you can use to select that mailbox for addition to the list. Once you are done selecting, click Add Selected Items to add those mailboxes to your list of searchable mailboxes.

6.1.8 Address Book Selector

This interface is utilized in various other areas, but is described here.



It shows the currently selected items at the top, and lets you delete an item by clicking the red X.

(The New Mailbox selector in the Search Interface is an exception; just choose another item)

Adding Mailboxes

1. Choose between the configured module systems
2. Fill out basic criteria to narrow your search results (or no criteria for the first 100)
3. Click Search
4. The results up to a maximum of 100 are displayed
5. The user can then page back and forth among the first 5 pages of results
6. Choose which of the results you want to add to the selected list
7. Click Add Selected Items

NOTE: You can restrict to just Users (skipping Resources). You can show only recently cached items (last 10 days). The search is not case sensitive.

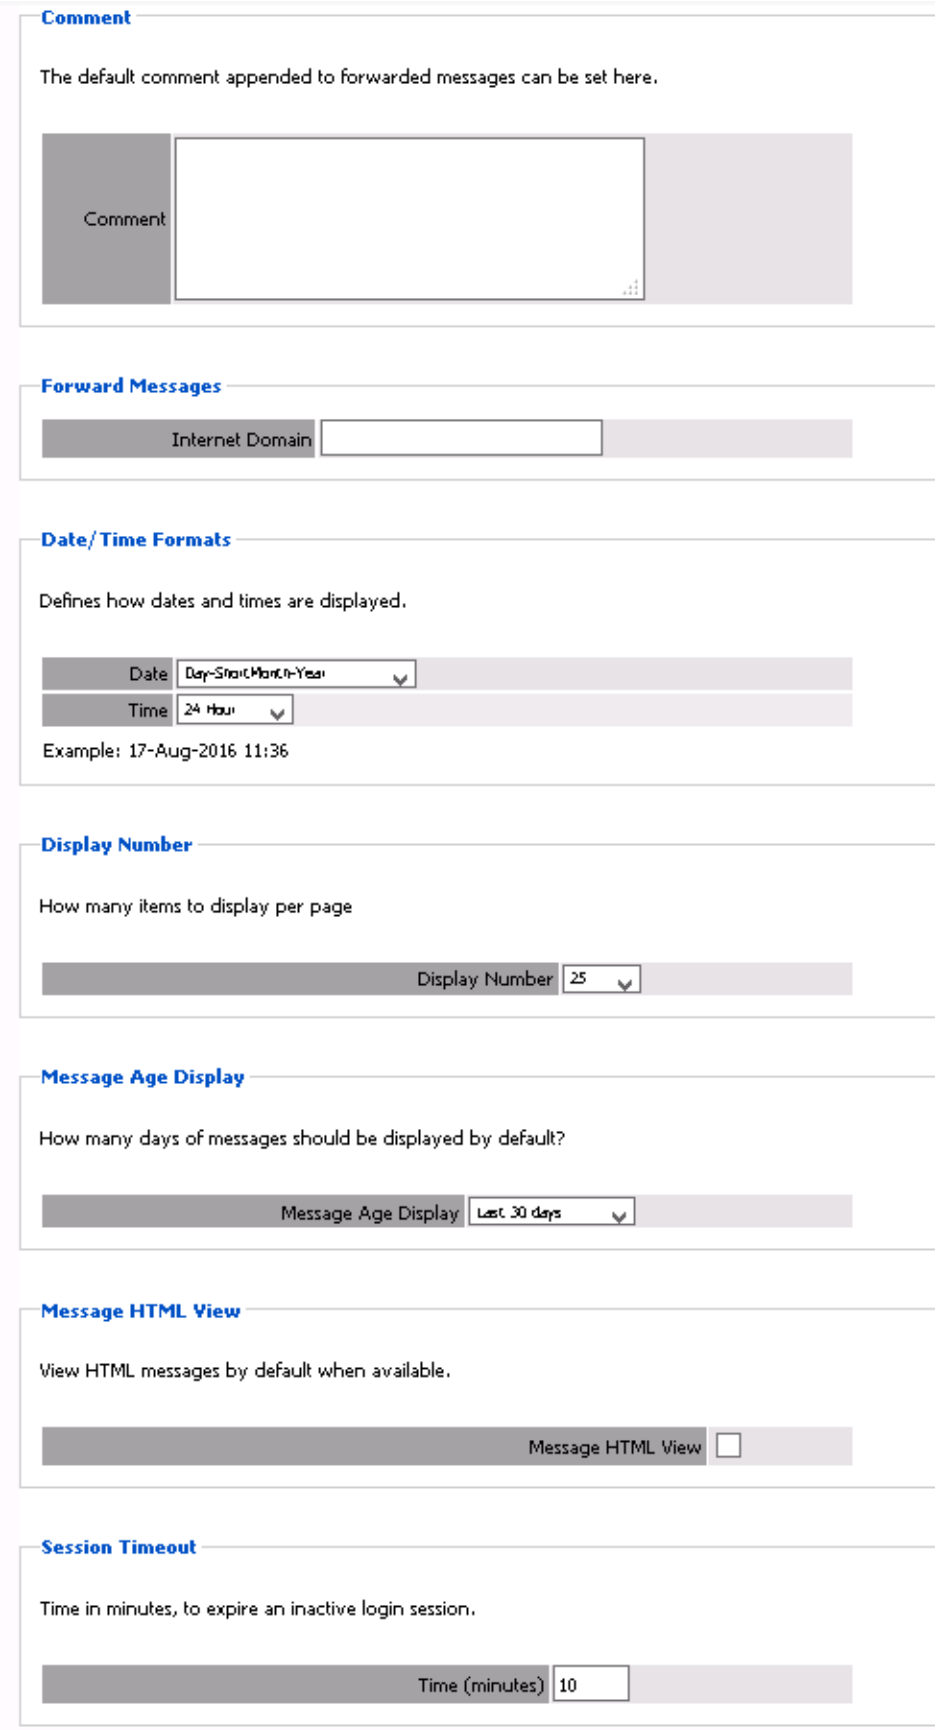
About “Show only recently cached items”

This option restricts the list of items shown in the selector to those with items stored within the last 10 days. In user/group management, it restricts the list to users who have logged in to the live Mail system within the last 10 days. The idea is to show only current items. If you DO want to see all items regardless of whether they’ve shown activity within the last 10 days, just uncheck this option.

Miscellaneous Tab

This tab contains settings that mainly govern the way the Search Interface works for the selected user.

Note that the user can change any of these settings by using the User Options tab in the Search interface.


Comment
The default comment appended to forwarded messages can be set here.
Comment
Forward Messages
Internet Domain
Date/Time Formats
Defines how dates and times are displayed.
Date
Time
24 Hour
Example: 17-Aug-2016 11:36
Display Number
How many items to display per page
Display Number
25
Message Age Display
How many days of messages should be displayed by default?
Message Age Display
Last 30 days
Message HTML View
View HTML messages by default when available.
Message HTML View
Session Timeout
Time in minutes, to expire an inactive login session.
Time (minutes)
10

- Comment: Default comment for forwarding messages
- Date/Time Format: How you want your dates and times to be displayed
- Display Number: How many items to display per page
- Forward Messages: Automatically append the specified address to forwarded messages
- Message Age Display: Default date filter for searching. May be changed on the fly
- Message HTML View: Have Retain display HTML messages by default, when available:
- Session Timeout: Inactive session timeout. Can be between 5 and 60 minutes

6.1.9 GroupWise Proxy Support

Retain supports the GroupWise proxy function. To enable it, check the box in the Module Configuration section. (NOTE: using proxy is useless if the user you wish to enable this function for is set to use offline authentication – found under the core settings of the user)

NOTE: The 'all user rights access' in GroupWise is not supported.

This function is used to enable a user to access the mailbox of another user. For example, if user B grants the right to user A to access their mailbox in the GroupWise client, then user A can "proxy" in to user B's mailbox.

Much the same way, if user A has proxy rights into user B's mailbox in GroupWise, and the function is enabled in Retain, then user A may select user B's mailbox for browsing or may search through user B's mailbox in the Search Screen.

In Retain, it is the MAIL READ right which grants access.

Retain uses the list of available mailboxes shown in the GroupWise client to determine which mailboxes will be made available to the logged in user (user A in our example). Thus, it is important that user A has logged into user B's mailbox as proxy using the GroupWise client before doing this in Retain. While user B might have granted the rights to user A, if user A has not yet logged in as proxy to user B's mailbox with GroupWise, then user B will not appear in user A's list of available accounts to proxy into.

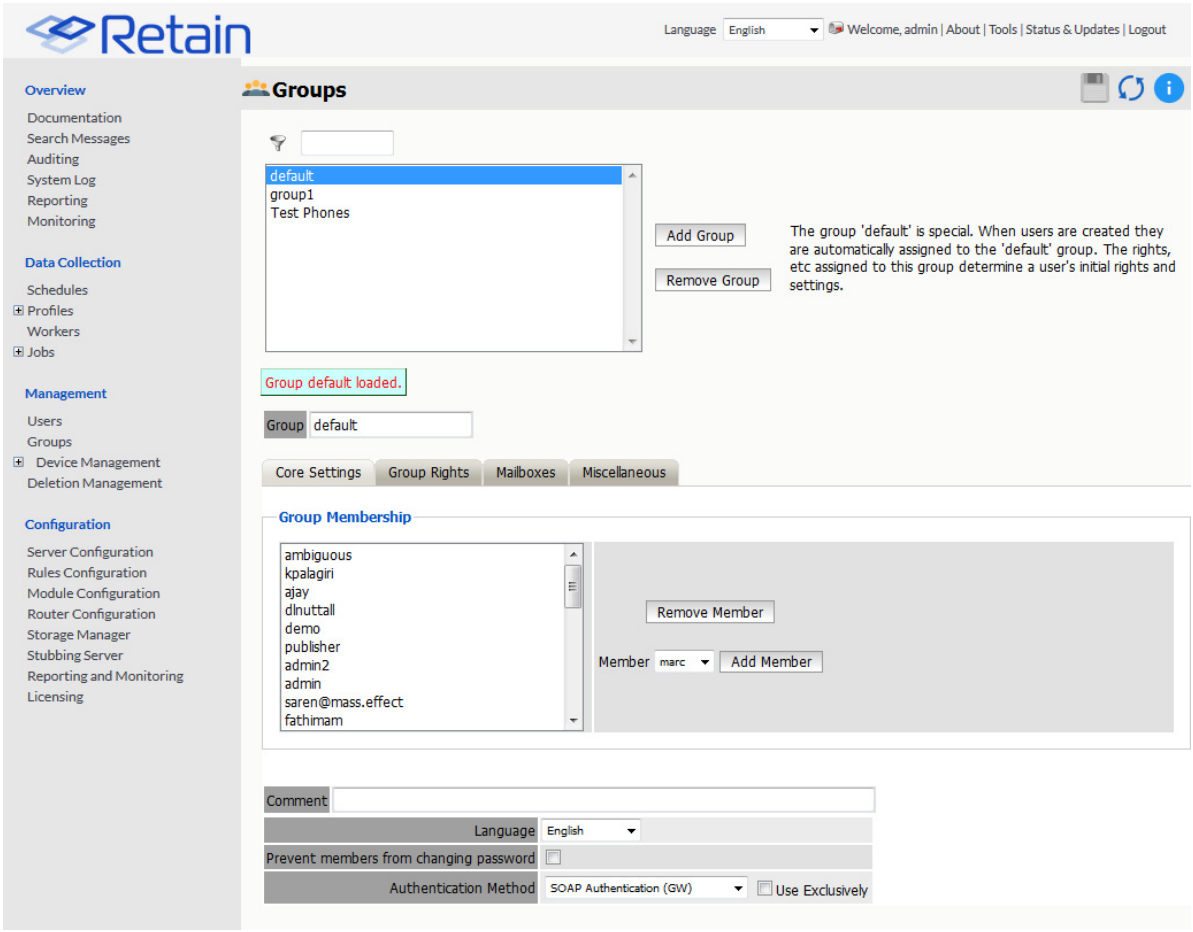
Retain checks these proxy rights the first time you access a proxy users mailbox, then caches the informa-tion for 7 days as configured in the server Configuration – Miscellaneous tab. (Default is 7 days.)

If you have access to another mailbox by virtue of GroupWise proxy, then you will see that mailbox appear in the mailbox selector in the search screen or you may search through that mailbox as well.

6.2 Groups

Once you have created users, you can group them and give them common rights. For instance, you can make a group called "Auditors" and this group will have rights to certain specific functions. If you then add users to this group, these users will get those specific rights in addition to the rights you have explic-itly given them.

This is a great way to save time in managing the rights of a large number of users.



This is the main Groups administration screen. It's laid out just like the Users screen and rights are assigned the same way, for the most part.

In the Groups screen, you will see all defined groups. The group default is created when you install Retain. Users are automatically assigned to this group. If you want users to be part of another group, you must first create that group and then add users to it.

You will notice the drop-down list that allows you to select a member. If it says "(no users)", it means that all users are part of the currently selected group and that there are no more users available to add.

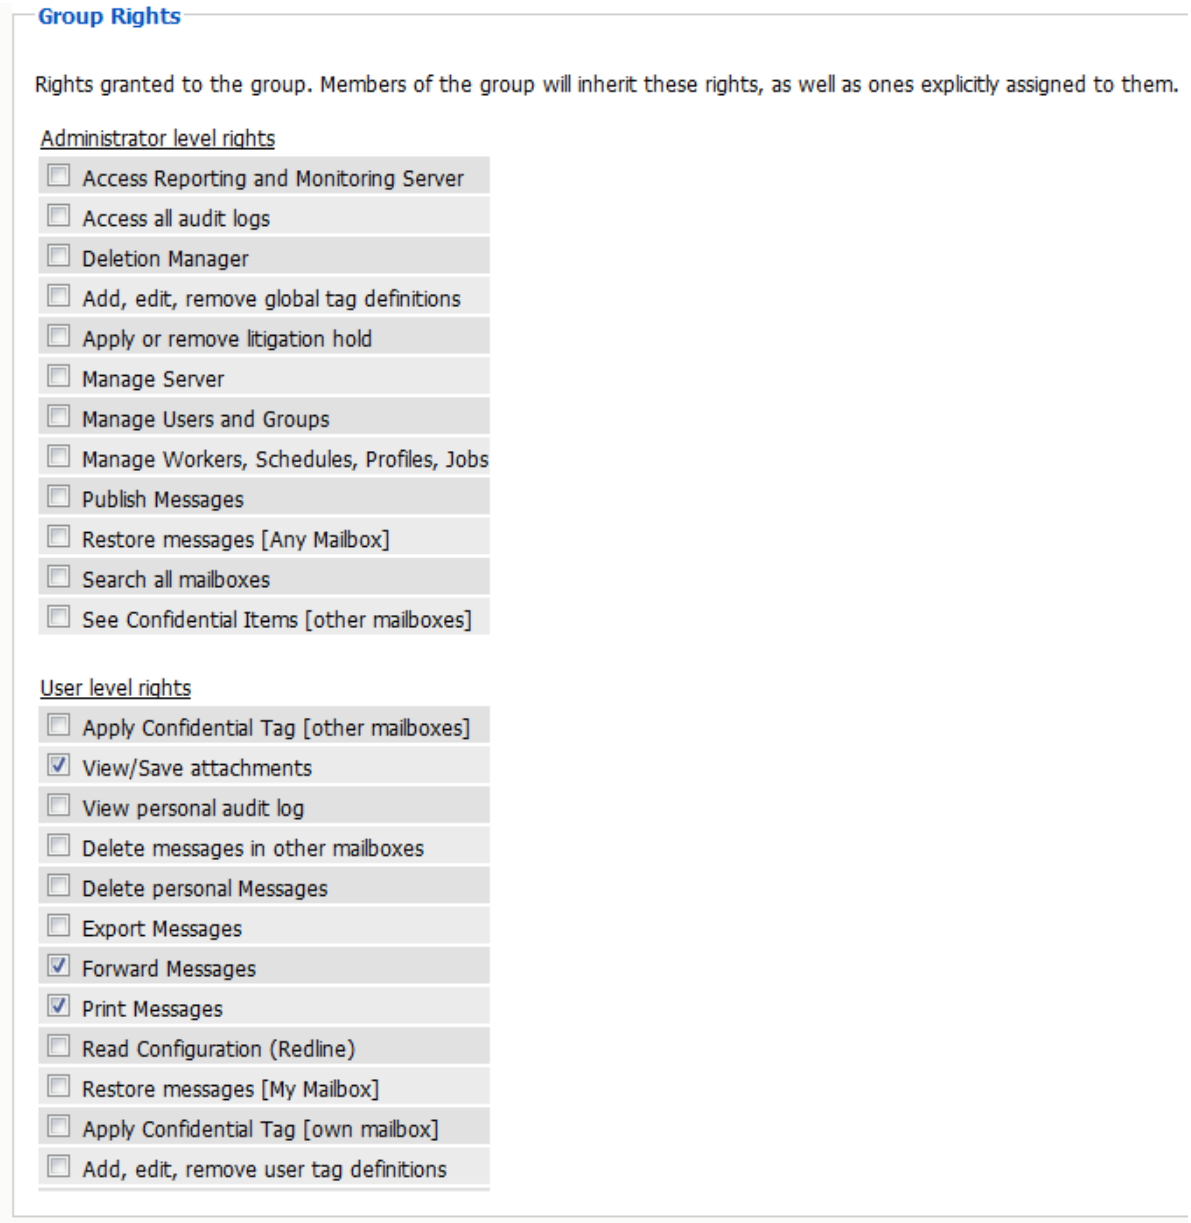
You may specify the default language used in the Search Interface for this group of users.

6.2.1 Creating a New Group

If you want to create a new group, simply

1. Click "Add Group"
2. Type the name in "Group"
3. Click "Save Changes"
4. Change the properties of the group in the tabs below

Group Rights



In this tab, you define all the rights that will be common to all members of the group. These rights are ADDITIVE and will be in ADDITION to the rights you have explicitly given to the individual user.

For this reason, if you want to manage users as a group, you would typically not assign any individual rights. Rather, you would assign rights to their group. These rights have all the same meanings and func-tion as the users' rights.

To log into and manage or monitor Workers, a User or Group must have the manage workers, Schedules, Profiles, Jobs right. To set or work with the Deletion Manager, the user must likewise have the Deletion Manager right, or they will not be able to modify those settings.

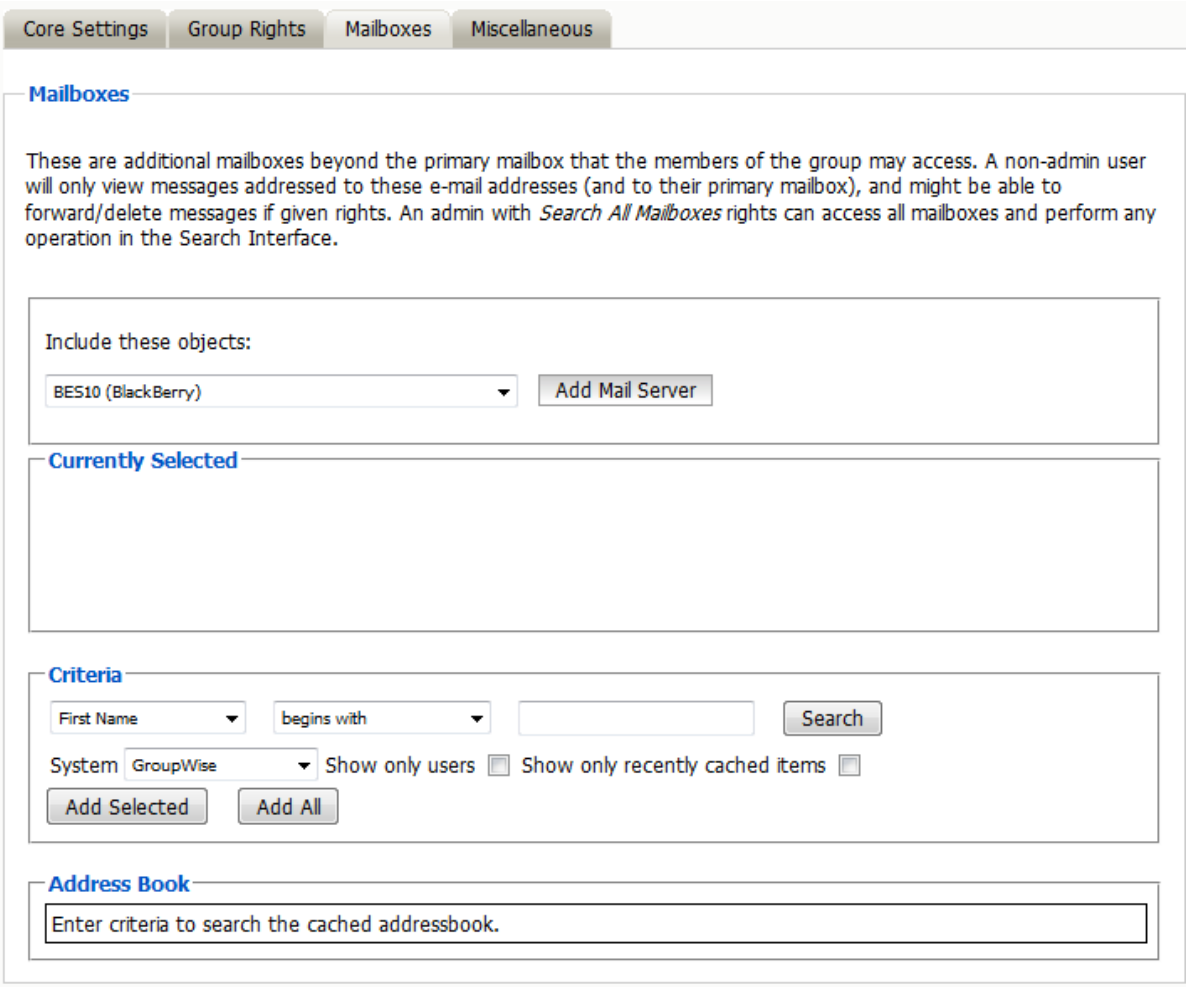
6.2.2 TIPS

Typically, you will not want to grant Redline rights to a group, but in a situation where you have more than one Redline control center and you want to see which one is monitoring, group-level Redline rights would be appropriate.

Typically, you will want to make the default group's rights rather restrictive, granting very limited access by default. If you wanted a user to have more rights, you would simply assign them to another group or you could assign additional explicit rights.

By using groups, you can have groups of individuals with rights to totally different sets of archives. Users only get rights to the mailboxes you allow them.

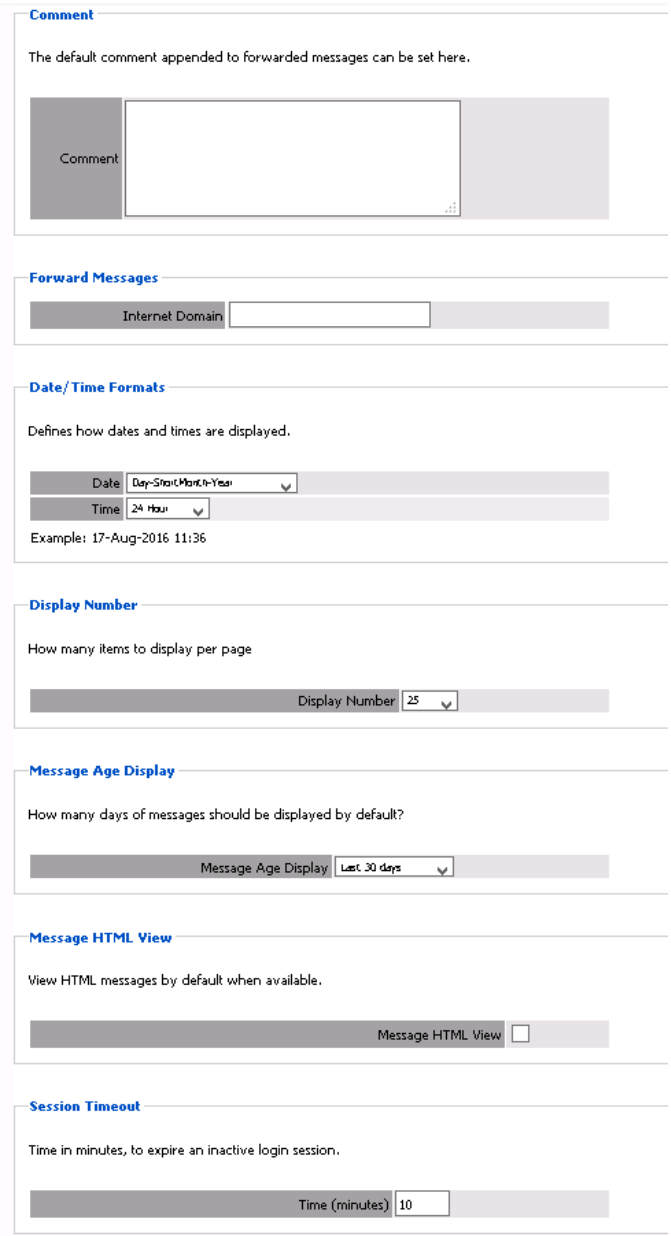
Mailboxes



- This screen works exactly like the user's access to mailboxes. Please see the user's mailbox sec-tion for details on how to select which mailbox or mailboxes to assign to the group.
- You use the address book selector to choose which mailbox or mailboxes to give the group access to.

- By default, groups have access to NO mailboxes; users may only access their own mailboxes. If the user exists only in Retain (no GroupWise account) and their group has access to no mailboxes, then they will have access to no mailbox either.
- Typically, you'll give group access to mailboxes only when you want all members of the group to be able to access a particular mailbox. Some examples could include a general sales account or accounts being audited by a group of auditors, a workgroup needing to access each other's archives, and so on.

6.2.3 Groups – Miscellaneous

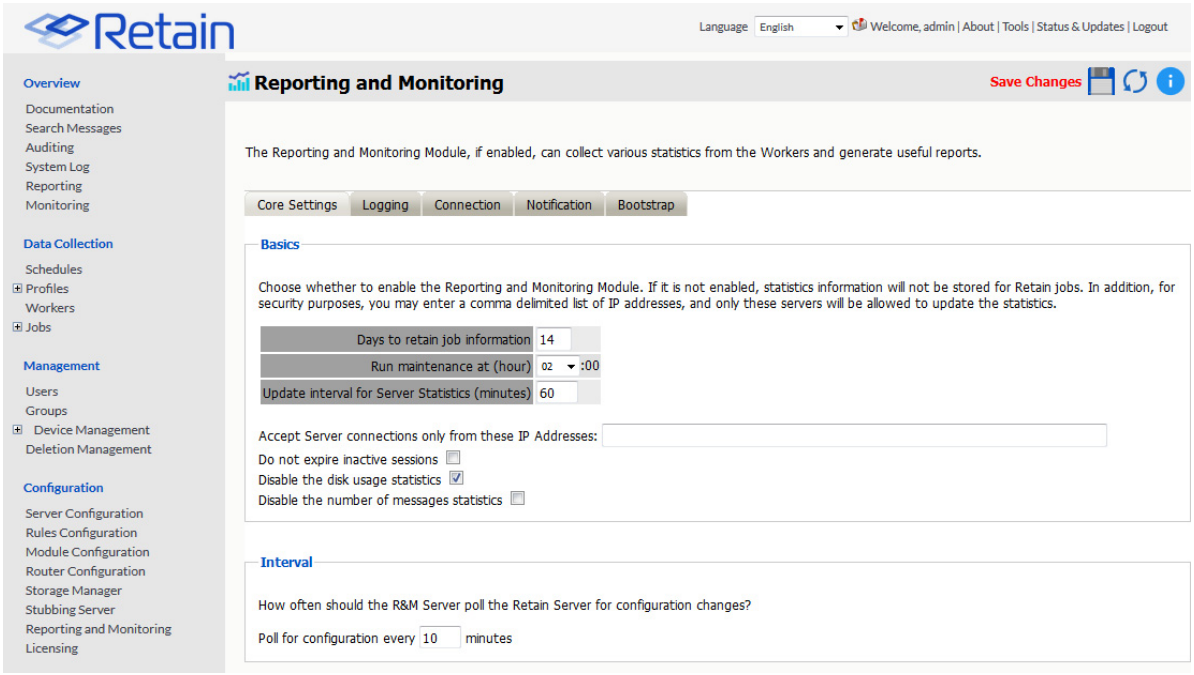


This is exactly like the users' miscellaneous configuration screen. It applies to all users in the group and users' explicit settings will override the group settings.

These settings govern the way the search interface works for the selected group. See more information about these settings in the users' miscellaneous section.

6.3 Reporting and Monitoring Configuration

Retain comes equipped with a separate Reporting and monitoring server. Reporting and Monitoring may be installed on the same or a different physical or virtual machine as the Retain Server.



Reporting and Monitoring functions much like the Retain Worker in that it is installed as a separate entity and requires connection information from the Retain Server. Additional security is provided through the option to specify exactly which IP address(es) the Retain Server will allow a connection from. Separate with commas the IP addresses allowed if more than one. No mail data is transferred through this connec-tion, only server statistic information.

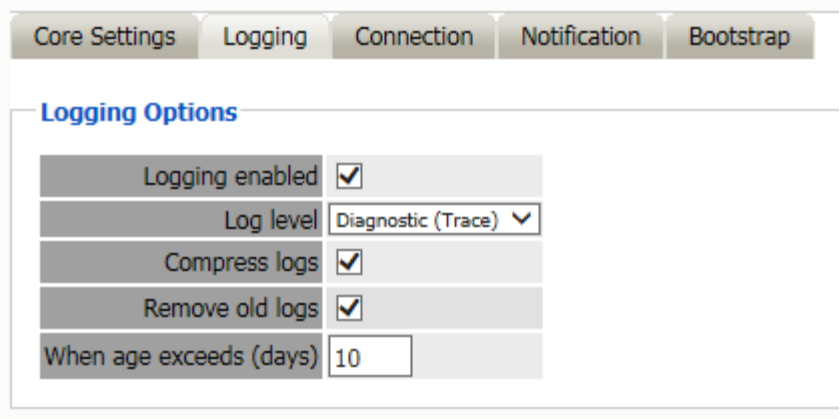
6.3.1 Setup

The Core Settings, Logging, Connection and Notification sections of Reporting and Monitoring configura-tion work identically to the Worker connection settings, and should be filled-out accordingly. Most server implementations will require little to no modification to these settings save the notification tab. Limiting the server connections to specified IP Addresses limits which machines have access to the Reporting and Monitoring module. Separate multiple addresses with a comma.

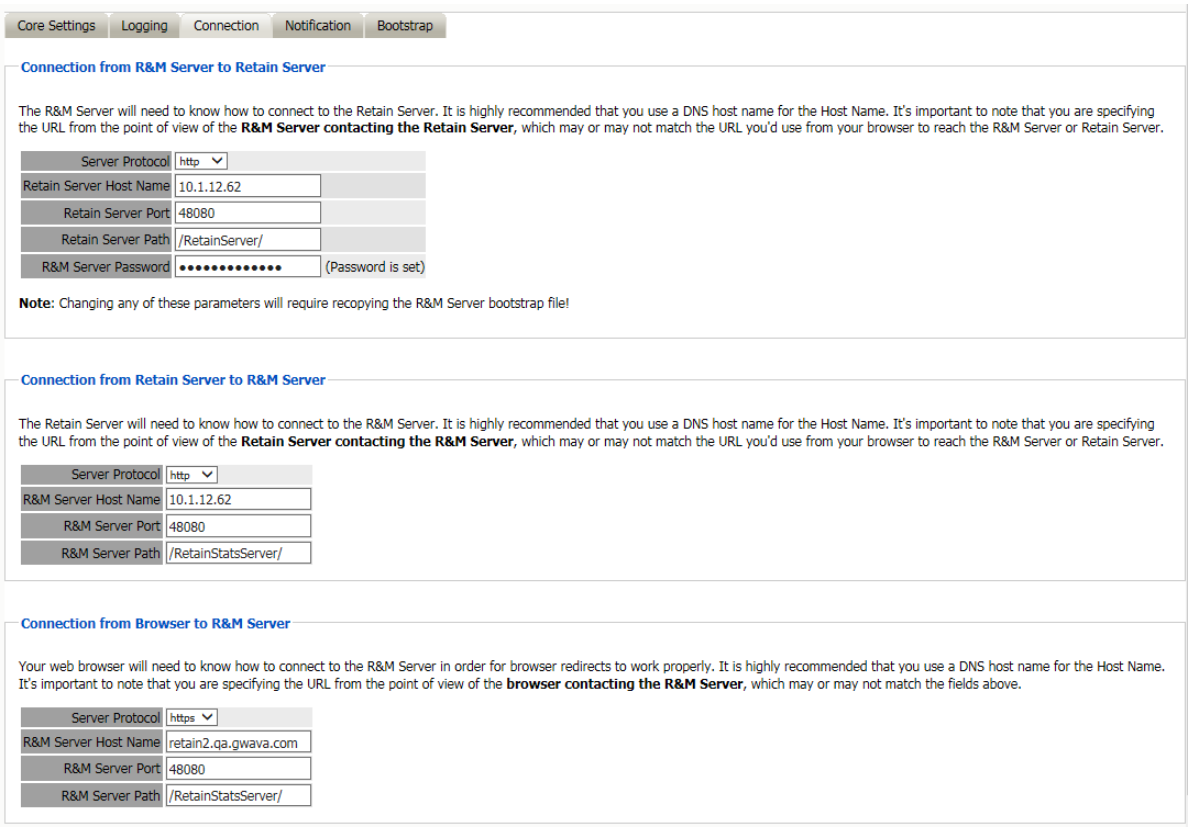
Putting a checkmark in the 'Do not expire inactive sessions' will allow sessions with the R&M module to never expire. This is useful in conjunction with the 'monitoring' page.

Selecting the 'Disable the disk usage statistics' box will disables the 'archive size' from server utilization report. Selecting 'Disable the number of messages statistics' disables the 'total messages stored' on the Reporting and Monitoring dashboard. These options are for systems which encounter a bottle-neck which impacts system performance with these checks.

Set the logging level to the desired setting. Leaving the logging as default 'normal' is recommended unless instructed to increase the logging level by Support.



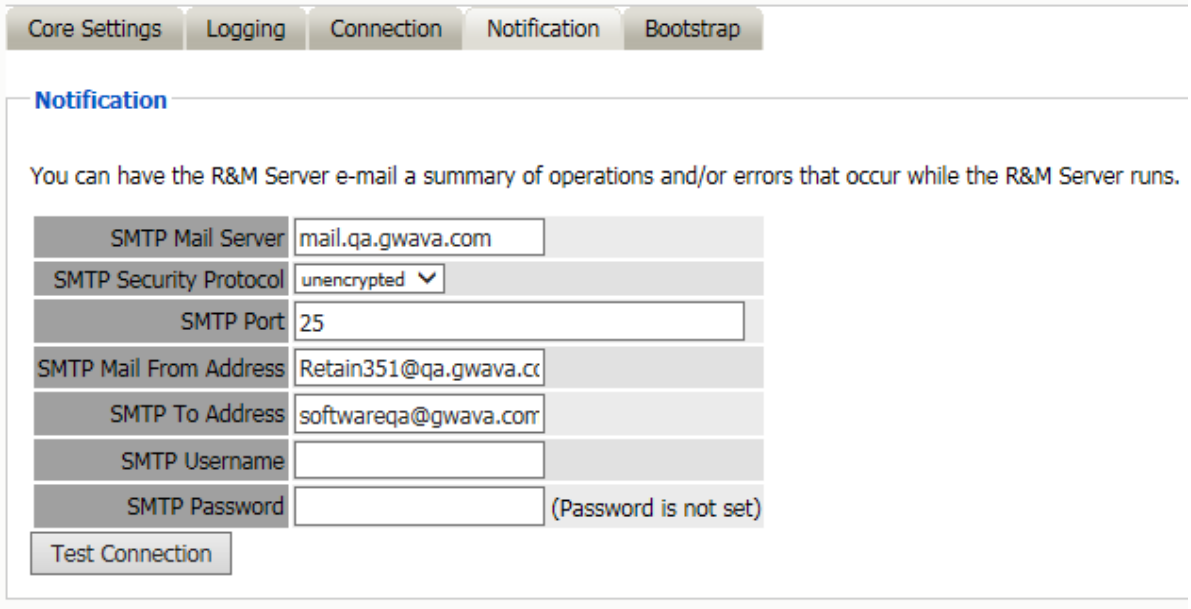
The Connection between Reporting and Monitoring, and the Retain Server should not be changed unless the Retain Server requires a modified connection path, (such as a custom port.)



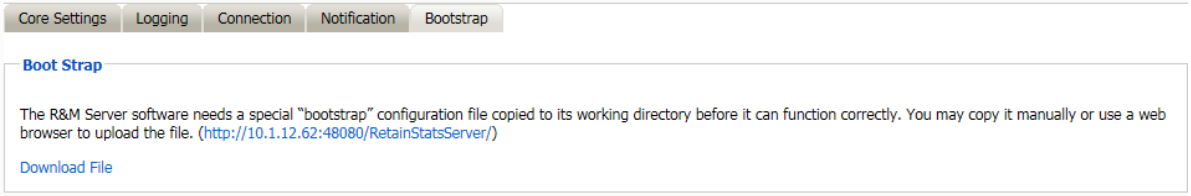
The R&M server is different from the Retain Worker in that the Retain Server is able to contact the R&M server as well. The connection address to R&M should not be modified unless custom connection settings have been created.

The connection from the Browser to the R&M Server is how the browser is going to connect to the R&M server when redirected from links in the Server. This connection information needs to be accurate for where the R&M Server is installed and listening.

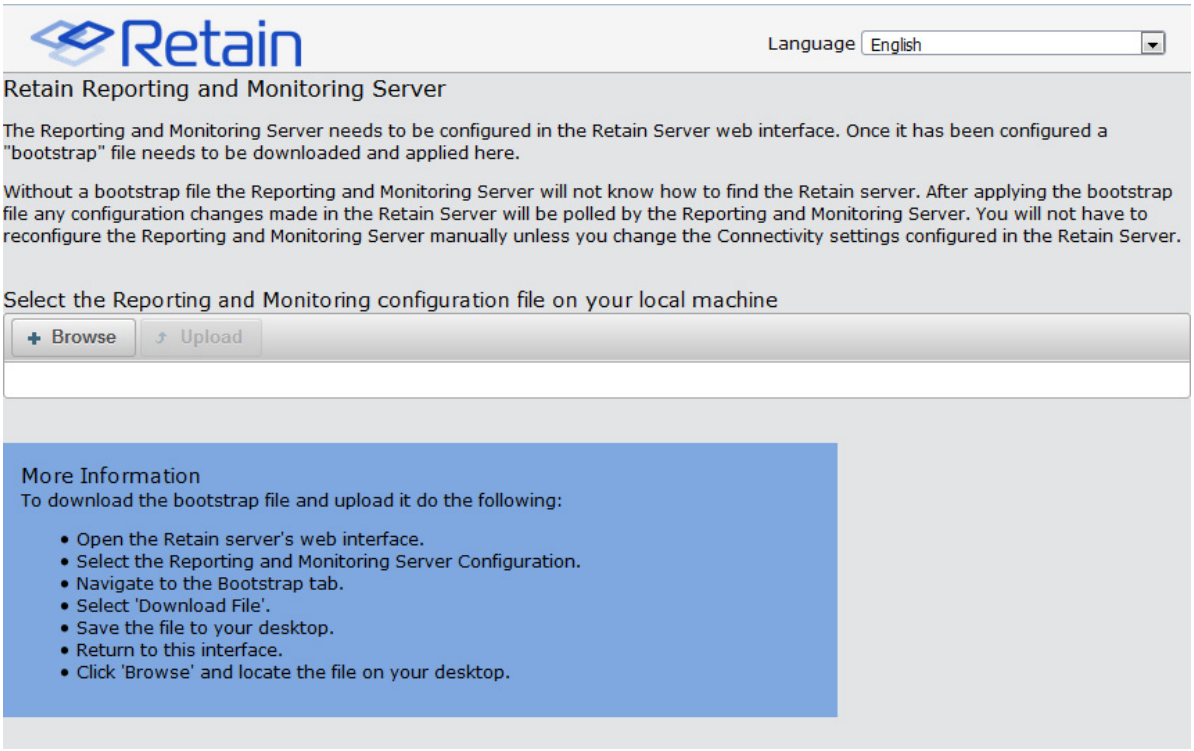
When configured to use HTTPS, the web server must be configured to use a full TLS certificate chain signed by a trusted third party Certificate Authority, not self-signed, and configured with both server and intermediate certificates.



Notification defaults to the standard notification address specified in the Retain Server activation and ini-tial configuration. If the notification is desired to be different, modify the settings here. Make sure to test the connection to ensure proper function.



Like the Retain Worker, the R&M server must be uploaded a saved configuration or 'bootstrap' file. To access this tab, the configuration must be saved first. Then the bootstrap tab will become available. Download the file to a known location and then browse to the R&M server webpage. The default connec-tion address is shown on the connection tab.



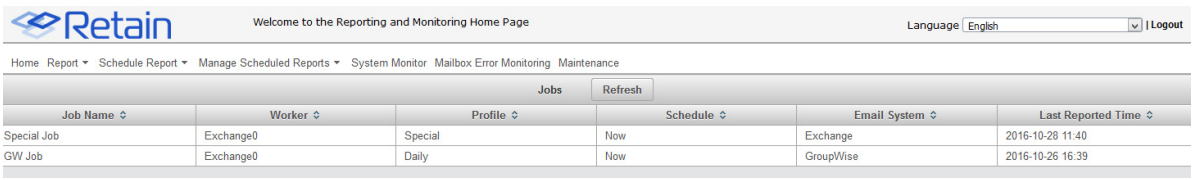
From the R&M initial webpage, click on the “Browse” button and browse to the saved bootstrap file and select it.

Click ‘Continue’ to upload the bootstrap file. As the bootstrap file is uploaded, R&M functions will become available. Only users with the administrative “Access Reporting and Monitoring Server” right will be able to login to the stat server.

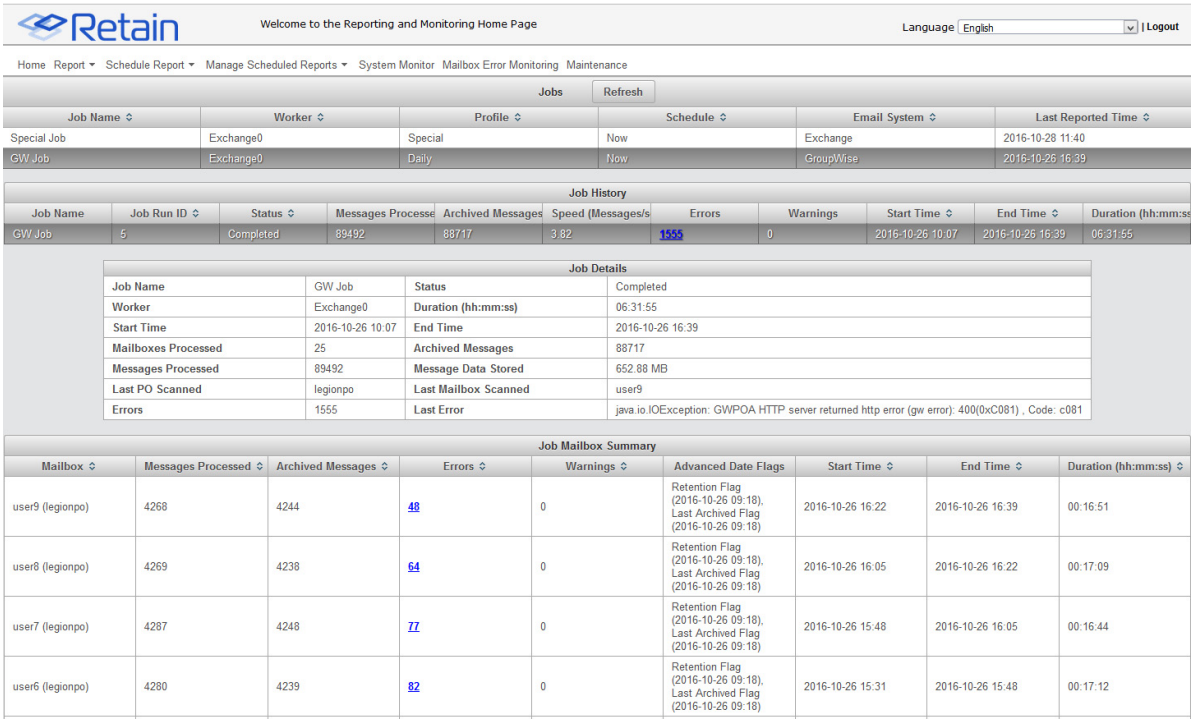
To log into the R&M server, select the ‘Reporting’ or ‘Monitoring’ links from the Retain Server left menu to spawn a new tab or page to the R&M module.

6.4 Reporting and Monitoring Home Page

On login, Reporting and Monitoring opens to the Home page. The Home page is also accessible through the ‘Home’ button along the top bar. The R&M server has different page links and menus across the top of the menu bar, accessing different functionality. The R&M server has many tables and columns display-ing a large amount of information. To make this information more accessible, every column has been made sortable. To sort the columns, simply click on them. They may be sorted to show the highest num-ber first, lowest number first, or not sorted at all.



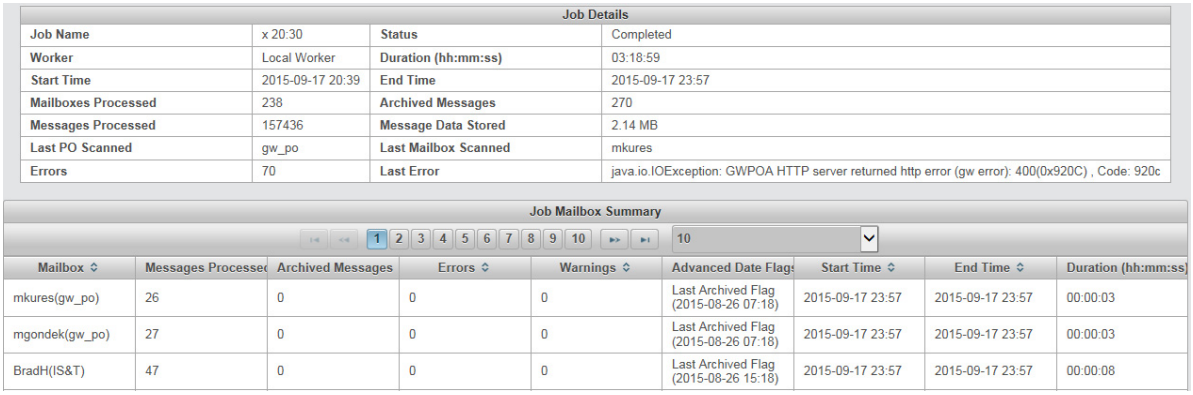
The Home page list the different jobs, associated workers, profiles, schedules, and last reported run times configured for the system. These jobs can be selected to display the job history and detailed infor-mation for each instance of when that job was run.



Further selection of a listed job instance will display details for that selected run of the selected job. If an error exists on any job, it will show up as a linked number under the error column. Errors reported here are linked to information on the error which has occurred.

A selected error will spawn a new window which displays the error and any applicable information, such as any associated files or error code numbers, a quick description (if a Retain error), the severity level, how often it has occurred, and in how many different mailboxes it has occurred.

The last error encountered in a job is also displayed in the Job details window. Select a column header to sort it as desired. The Errors column shown below has been sorted to show the highest number of errors first.



6.5 Reports

6.5.1 Report and Schedule Report

Report and Schedule Report functions are accessed through drop-down menus in the interface. To access the different reports, select the ‘Report’ or ‘Schedule Report’ menus from the top menu and then click on the desired report. Most reports which are accessible through the ‘Report’ menu are also available as a ‘Scheduled Report’, except for ‘Job Run Mailbox Summary’. All reports under the ‘Report’ menu except ‘Job Run Mailbox Summary’, can be scheduled and have a ‘Schedule’ button at the bottom of the deliv-ered HTML report. Selecting this button copies the settings used to generate the visible report to the Schedule Report interface. This convenience is designed to allow rapid fine-tuning of reports to deliver the desired information. First select the report you wish to run, and then fine-tune the report until it reaches the desired effect. Once the desired settings have been achieved, select the ‘Schedule’ button from the bottom of the report to carry all settings into the corresponding ‘Schedule Report’ interface, and schedule the report.

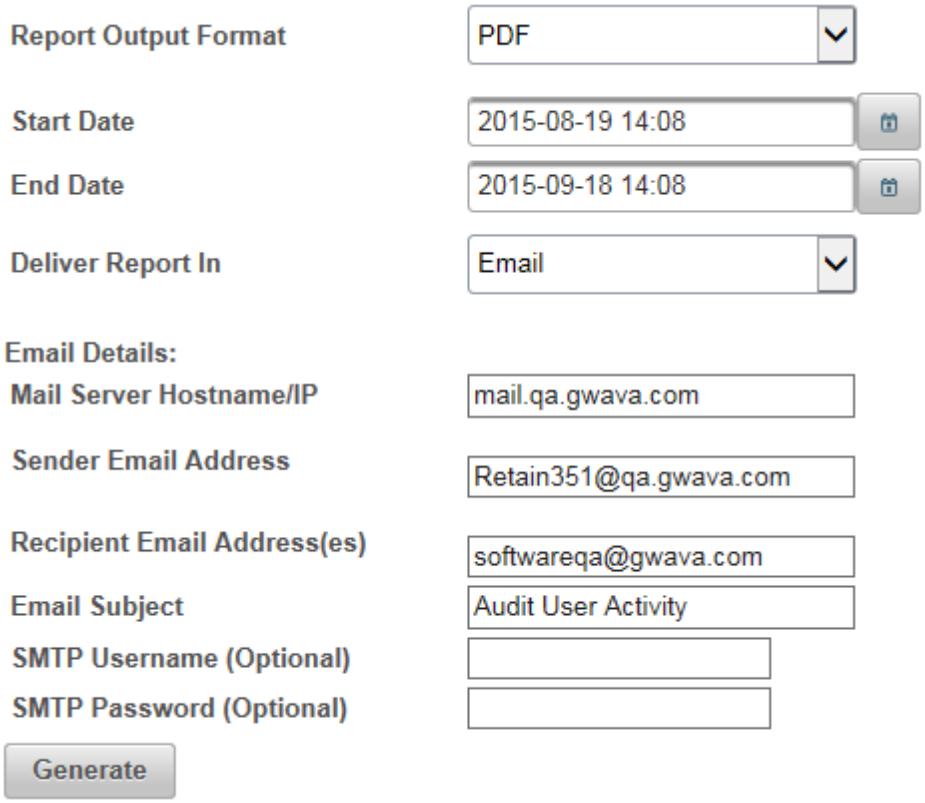


Reports differ from Scheduled reports in more than immediate delivery, reports support HTML in addi-tion to PDF and XLS delivery formats. However, both report structures may be emailed to a desired recip-ient. When a report is created and emailed it is formatted as selected and then attached to a message which is sent to the desired recipient.

When the Email delivery option is selected, additional configuration is required. The R&M module needs to know not only what address to send the report to and what mail server to use, but also allows the con-figuration of the email subject and sender address. If an SMTP username and password are required for the desired host system, they should be provided.

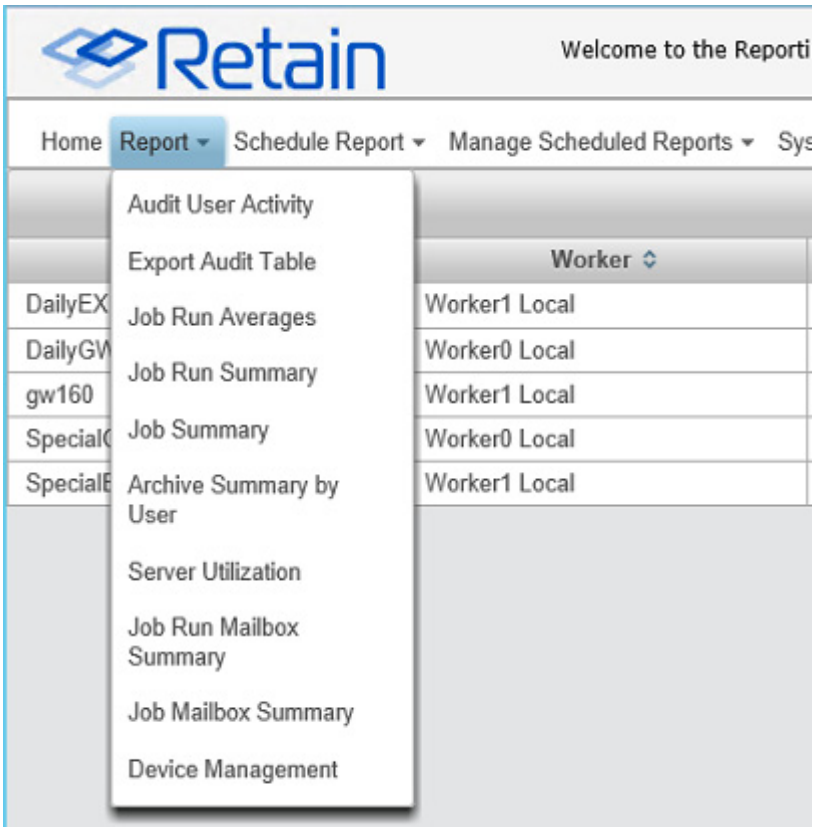
The R&M module will automatically copy the SMTP configuration from the Retain Server. If this informa-tion is not correct, it should be configured here before selecting the ‘Generate’ button.

This configuration is the same interface for all emailed reports, immediate and scheduled.



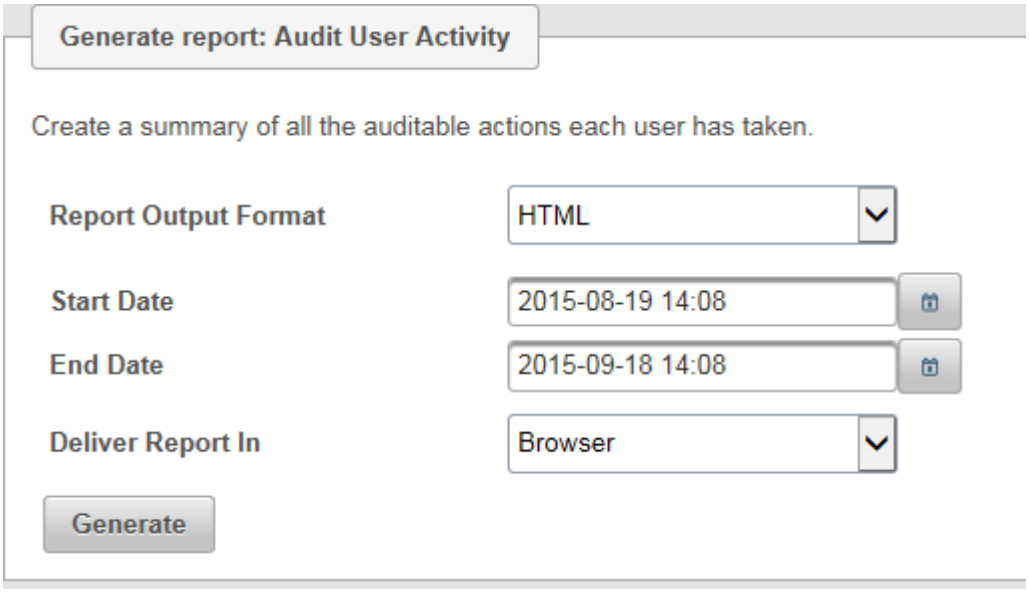
6.5.2 Reports

Reports may be created in HTML, PDF, or XLS, and may be delivered straight to the Browser or in a speci-fied email. Different reports have different settings, however the basics are the same; select or specify the desired criteria, specify the format and delivery method, then click on the “Generate” button to cre-ate the report.

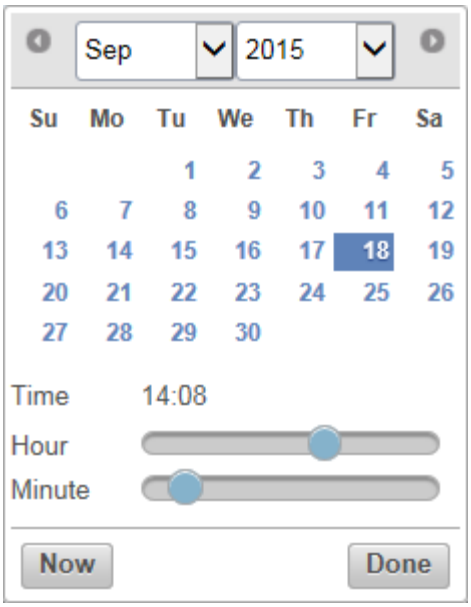


6.5.3 Audit User Activity

The Audit User Activity report details the different actions performed in the Retain system by different users. This report is a detailed list of all auditable actions taken and which users took them in the Retain system.



While this report lists all actions by all users, it may still be restricted by a date range. The date ranges are specified through a calendar interface.



The date and year are selected through the top bar, while the hour and minute, if desired, are selected through sliders. The date can be immediately set to the current moment by selecting the 'Now' button. When configuration has been completed, select the 'Done' button to set the date. Both the Start and End Date have identical date tools.

When the settings are as desired select the 'Generate' button to create the report.

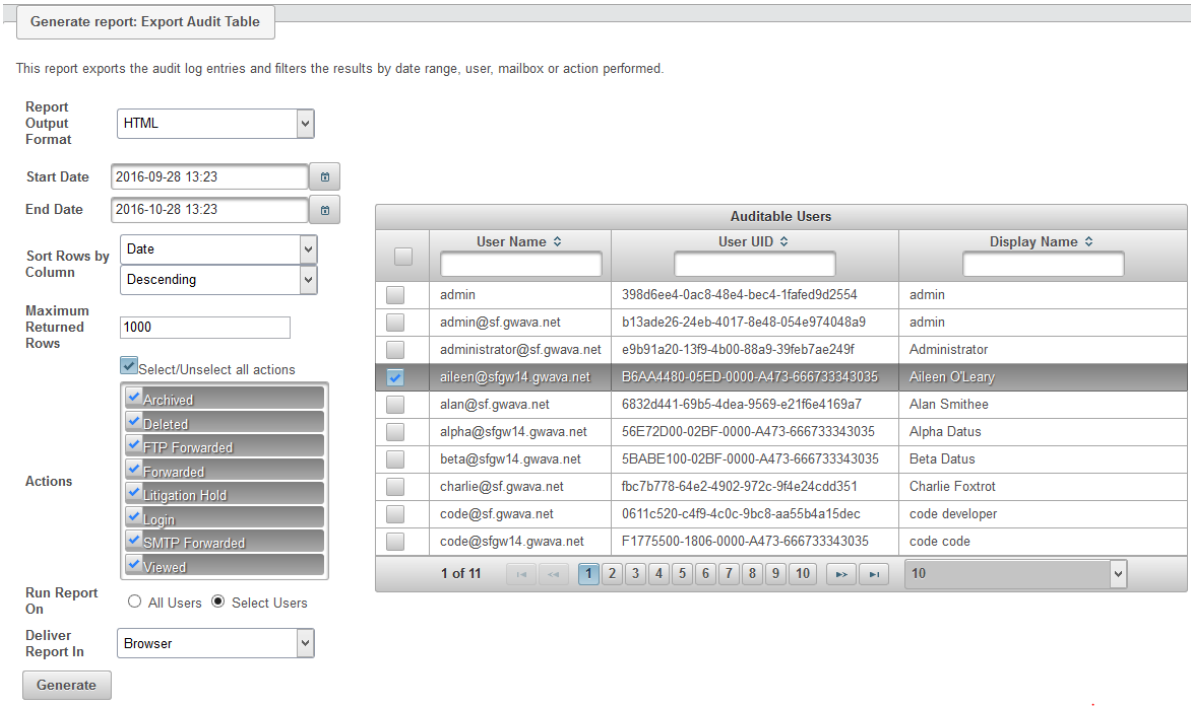
The report lists all users, a quick description of the action taken and how many times it was performed, and the date of the last time it was performed. Actions such as viewing an item, searching the archive, changing settings on either the server or in the interface, forwarding items, applying litigation holds, pub-lishing items, etc. To take the same settings used when creating this report, select the 'Schedule' button to enter the Schedule Report interface with the same settings.



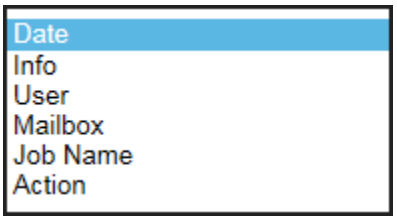
User Name	Count	Action	Last Action
admin	138	Login	2015-09-18 14:01
	36	Changed User or Group	2015-09-17 09:47
	89	Switched	2015-09-17 11:04
	2	Litigation Hold	2015-09-17 09:38
	187	Viewed	2015-09-18 00:03
	88	Searched	2015-09-18 00:02
admin2	4	Switched	2015-09-10 11:59
	3	Changed User or Group	2015-09-17 14:02
	54	Login	2015-09-18 13:32
	4	Changed Server	2015-09-15 15:53
	39	Searched	2015-09-11 15:19
	14	Viewed	2015-09-10 11:57
ajay	0	No action performed	None
ambiguous	0	No action performed	None

6.5.4 Export Audit Table

The Export Audit Table sorts the entire Audit log with the settings desired and creates a report for specific users or all users, with specific events or all events. This report can be limited in size and date.

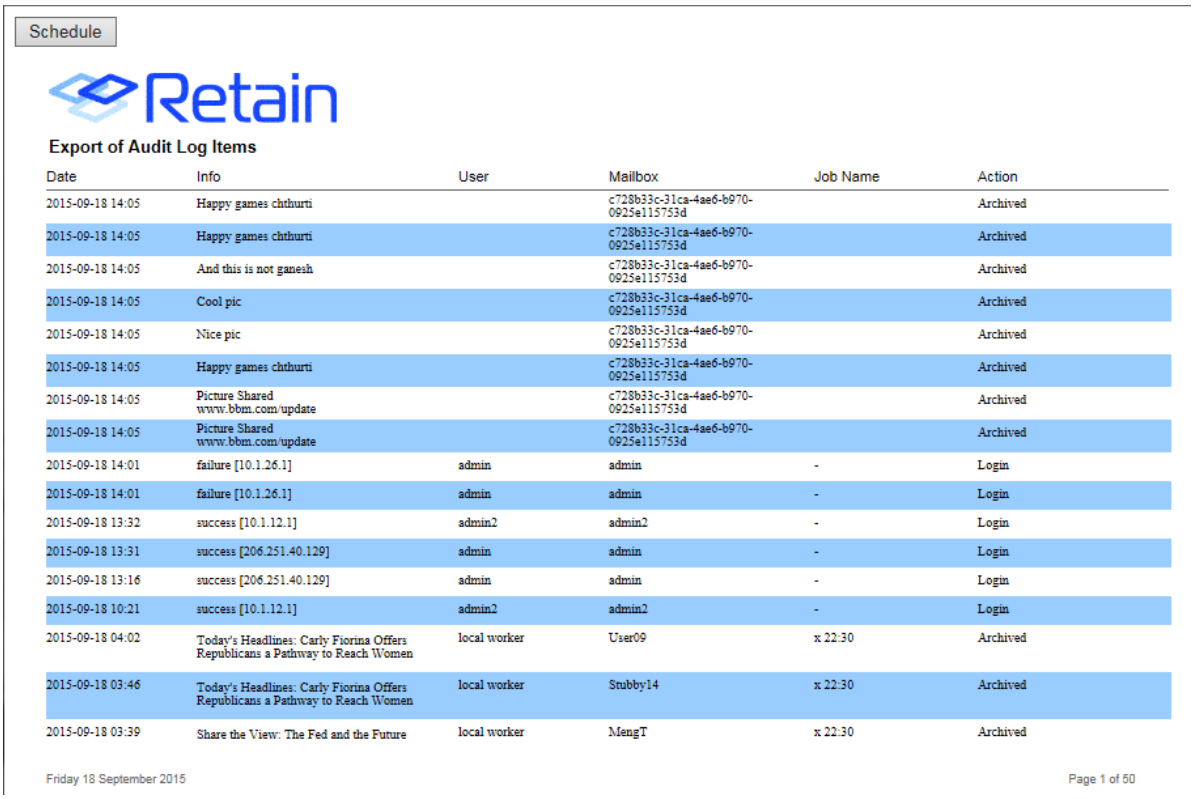


If restricting the report to a specific user or users, a menu displaying ‘auditable users’ is displayed. The different fields are searchable. The option to select all, or none, is also present. Users may be searched for via User Name or email address, User UID, or the user Display Name.



The rows and columns are not sortable in the generated report, but they can be sorted or organized from the report generation interface. The rows may be organized by the columns listed.

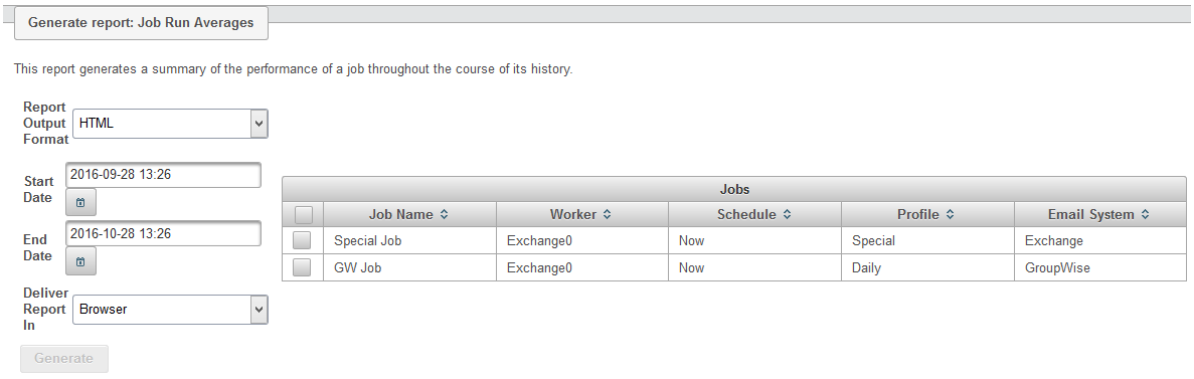
This has the exact same effect as organizing columns for the jobs on the reporting and monitoring home page. Configure as desired and then select ‘Generate’ to create the report.



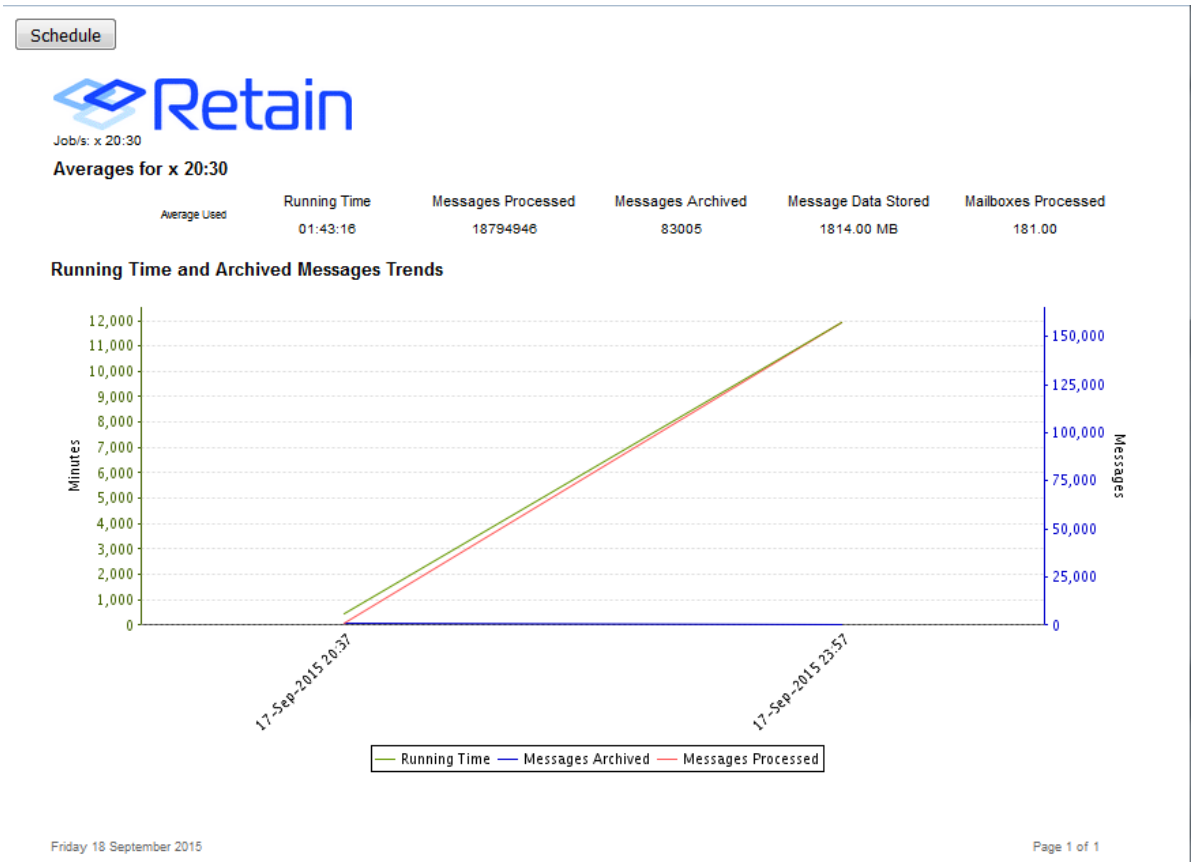
To Schedule the report, select the ‘Schedule’ button from the top of the report to keep the settings.

6.5.5 Job Run Averages

The job run averages report takes the performance of any selected particular job and creates a chart to display the trends. The effectiveness, overall load, and performance can be rapidly displayed. The Job Run Averages report requires a single job to be selected. The ‘Jobs’ selection window has columns which are sortable for convenience.



Once the date period has been configured, select ‘Generate’ to create the report.



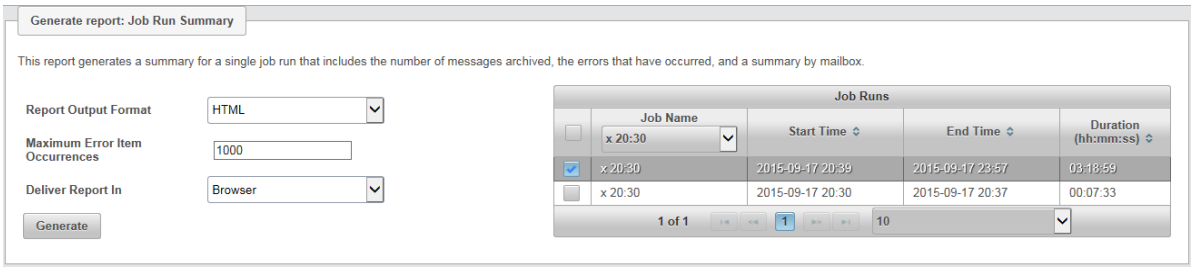
Performance trends are immediately apparent as run time is compared with message counts over a time period. The averages used to create the graph are listed above the graph, as well as a legend below.

As with the rest of the reports which exist in the ‘Schedule Report’ interface, to schedule this report with the settings which created it, select the ‘Schedule’ button from the top of the report to use the same set-tings and be linked directly to the ‘Schedule Reports’ interface.

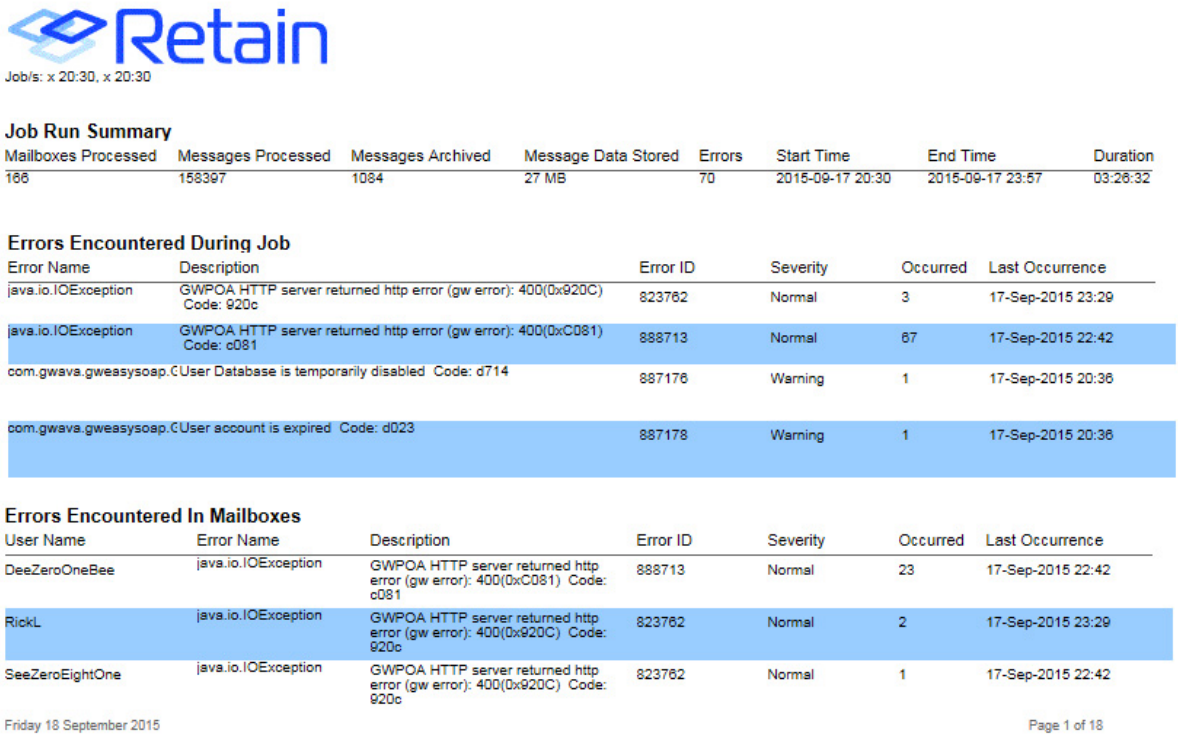
6.5.6 Job Run Summary

The Job Run Summary report gives a complete report of the performance of a single job. Whereas the Job Run Averages report charts raw performance averages, the Job Run Summary report goes into greater detail on the total number of archived messages, the errors, and mailbox summary for a single instance of a job.

Job Run Summary and Job Summary reports have the option to limit the number of reported errors dis-played.



Select the desired single instance of a job and then select ‘Generate’.



Retain

Job/s: x 20:30, x 20:30

Job Run Summary

Mailboxes Processed	Messages Processed	Messages Archived	Message Data Stored	Errors	Start Time	End Time	Duration
166	158397	1084	27 MB	70	2015-09-17 20:30	2015-09-17 23:57	03:26:32

Errors Encountered During Job

Error Name	Description	Error ID	Severity	Occurred	Last Occurrence
java.io.IOException	GWPOA HTTP server returned http error (gw error): 400(0x920C) Code: 920c	823762	Normal	3	17-Sep-2015 23:29
java.io.IOException	GWPOA HTTP server returned http error (gw error): 400(0xC081) Code: c081	888713	Normal	67	17-Sep-2015 22:42
com.gwava.gweasysoap.C	User Database is temporarily disabled Code: d714	887176	Warning	1	17-Sep-2015 20:36
com.gwava.gweasysoap.C	User account is expired Code: d023	887178	Warning	1	17-Sep-2015 20:36

Errors Encountered In Mailboxes

User Name	Error Name	Description	Error ID	Severity	Occurred	Last Occurrence
DeeZeroOneBee	java.io.IOException	GWPOA HTTP server returned http error (gw error): 400(0xC081) Code: c081	888713	Normal	23	17-Sep-2015 22:42
RickL	java.io.IOException	GWPOA HTTP server returned http error (gw error): 400(0x920C) Code: 920c	823762	Normal	2	17-Sep-2015 23:29
SeeZeroEightOne	java.io.IOException	GWPOA HTTP server returned http error (gw error): 400(0x920C) Code: 920c	823762	Normal	1	17-Sep-2015 22:42

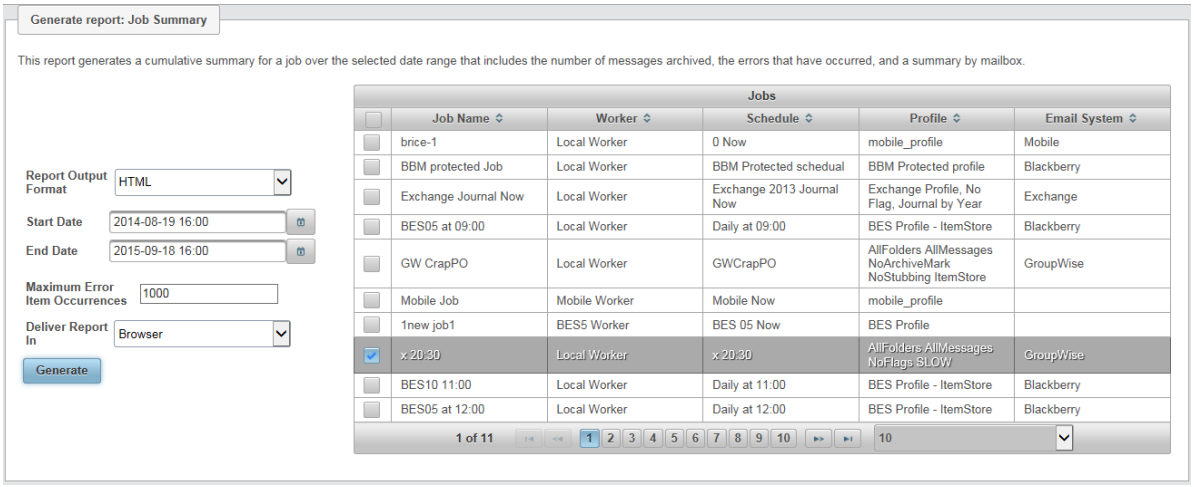
Friday 18 September 2015 Page 1 of 18

The summary displays the total mailboxes, messages, data store, errors, and running time of the job. Afterwards, this information is broken down by each mailbox the job covers.

6.5.7 Job Summary

The Job Summary report is much the same to the Job Run Summary report, except that this report covers the entire history of a job and not just a single run. To run the report, select the desired job and any time period, if desired, the desired output and delivery method, and click ‘Generate’.

Job Summary contains the option to limit the number of reported errors displayed.

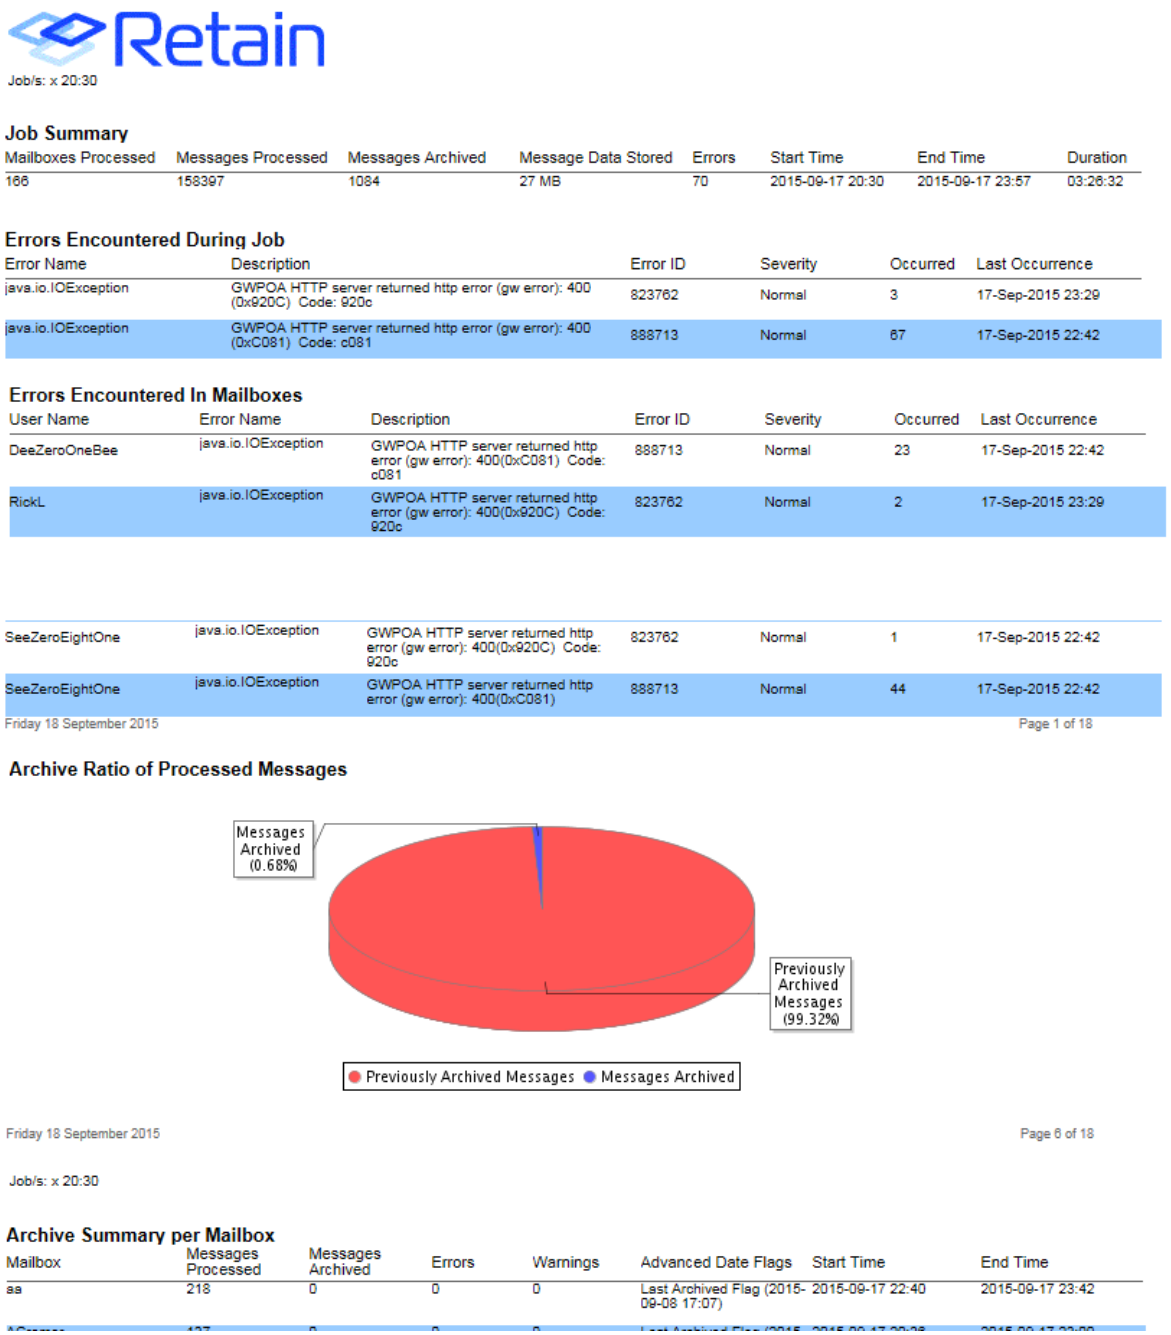


The report generated displays the number of mailboxes that the job processes, the total messages the job has processed, how many messages are actually archived, the message data stored, the total errors encountered, when the job first started, and the total time the job has been active.

Immediately below the initial information, a pie chart displays the amount of messages actually archived in comparison to the number of messages the job processed.

Finally, below the pie chart is a summary, by individual mailbox, of job actions. Each mailbox is displayed with the pertinent information, including the messages archived, messages processed, critical errors, a record of the advancement of message store flags, and the duration of the job in that mailbox.

To schedule this job with the settings desired, select the ‘Schedule’ button from the top of the report. All settings used to create the report will automatically be copied into the ‘Schedule Report’ interface.



Retain

Job/s: x 20:30

Job Summary

Mailboxes Processed	Messages Processed	Messages Archived	Message Data Stored	Errors	Start Time	End Time	Duration
166	158397	1084	27 MB	70	2015-09-17 20:30	2015-09-17 23:57	03:26:32

Errors Encountered During Job

Error Name	Description	Error ID	Severity	Occurred	Last Occurrence
java.io.IOException	GWPOA HTTP server returned http error (gw error): 400 (0x920C) Code: 920c	823762	Normal	3	17-Sep-2015 23:29
java.io.IOException	GWPOA HTTP server returned http error (gw error): 400 (0xC081) Code: c081	888713	Normal	67	17-Sep-2015 22:42

Errors Encountered In Mailboxes

User Name	Error Name	Description	Error ID	Severity	Occurred	Last Occurrence
DeeZeroOneBee	java.io.IOException	GWPOA HTTP server returned http error (gw error): 400(0xC081) Code: c081	888713	Normal	23	17-Sep-2015 22:42
RickL	java.io.IOException	GWPOA HTTP server returned http error (gw error): 400(0x920C) Code: 920c	823762	Normal	2	17-Sep-2015 23:29
SeeZeroEightOne	java.io.IOException	GWPOA HTTP server returned http error (gw error): 400(0x920C) Code: 920c	823762	Normal	1	17-Sep-2015 22:42
SeeZeroEightOne	java.io.IOException	GWPOA HTTP server returned http error (gw error): 400(0xC081)	888713	Normal	44	17-Sep-2015 22:42

Friday 18 September 2015 Page 1 of 18

Archive Ratio of Processed Messages

Friday 18 September 2015 Page 6 of 18

Job/s: x 20:30

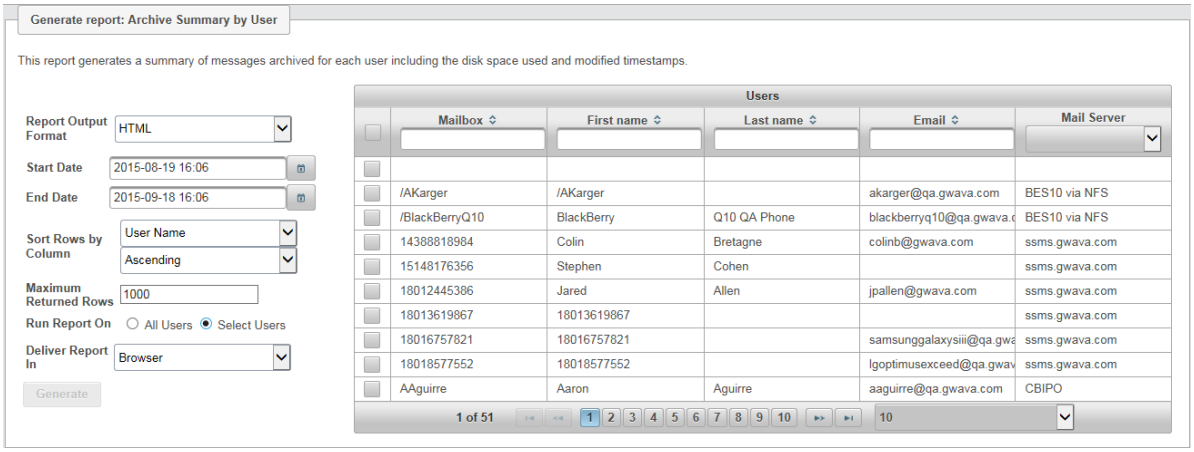
Archive Summary per Mailbox

Mailbox	Messages Processed	Messages Archived	Errors	Warnings	Advanced Date Flags	Start Time	End Time
aa	218	0	0	0	Last Archived Flag (2015-09-08 17:07)	2015-09-17 22:40	2015-09-17 23:42

6.5.8 Archive Summary by User

The Archive Summary by User report creates a summary of all users', or a specific user's, archive history. The summary includes all pertinent information for each mailbox's archive over the last 3 months of jobs. To run the report, either all or one mailbox must be selected. If all mailboxes are selected, the maximum returned rows may be a very important setting to keep the report from becoming immense.

While the report cannot be sorted after it is generated, it may be sorted at generation through the 'Sort Rows by Column' feature. The columns may be sorted by either descending or ascending order. Select 'Generate' to create the report.



The report displays the total information for every mailbox selected; User name, the time stamps and their advanced dates, and the message store for the associated user.

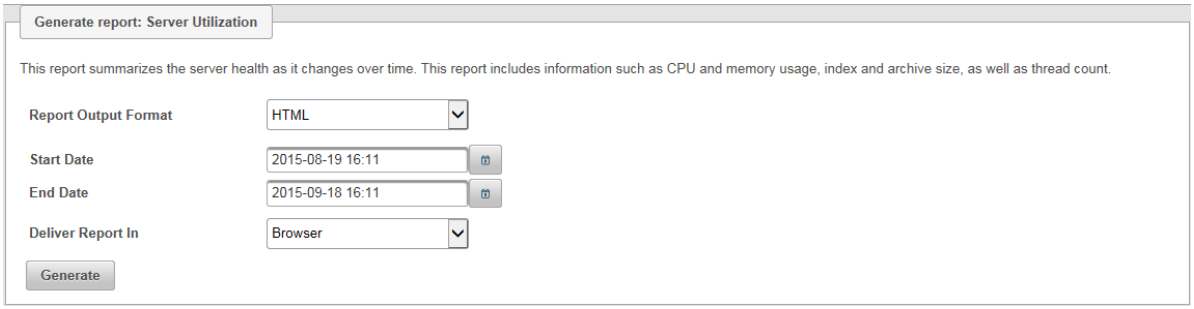


The most valuable part of this report, and what will make this report a staple for administrators, is the ability to monitor and get a concise report for the disk space used for each individual user along with the messages processed and stored.

If the report is to be scheduled, select the ‘schedule’ button from the top of the report to use the same settings which created the report.

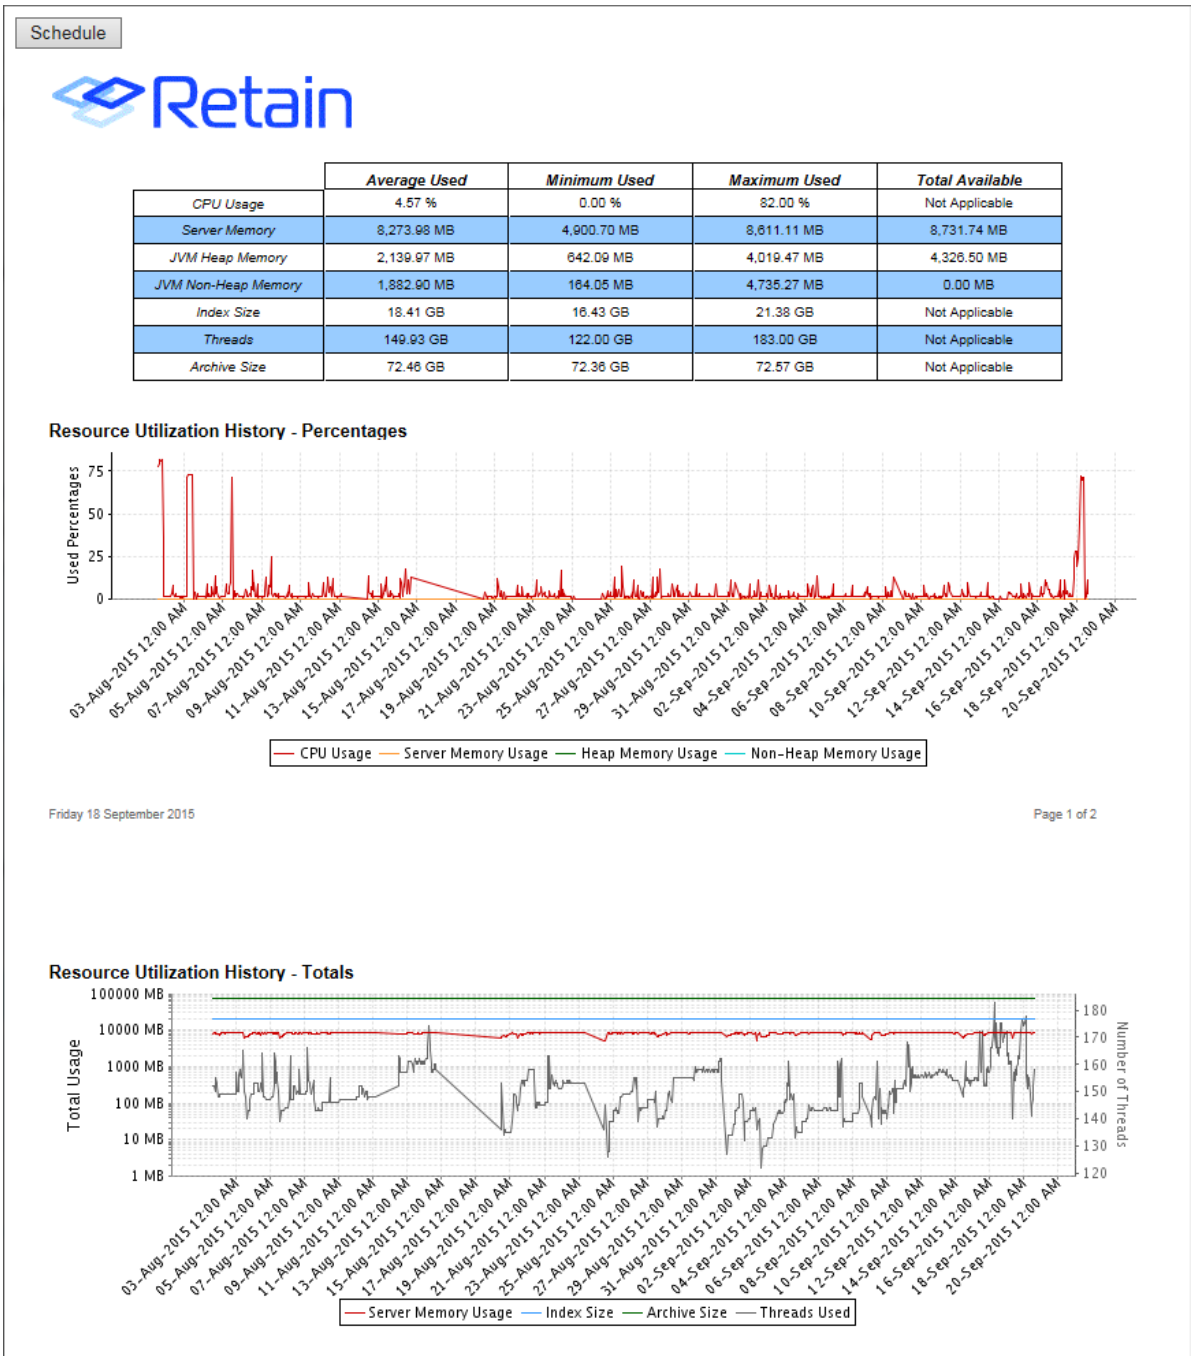
6.5.9 Server Utilization

The Server Utilization report is a direct report on the load that the Retain Server is under. The only option is to limit the date range for which the report is active. Select the report period and then select ‘Gener-ate’ to produce the report.



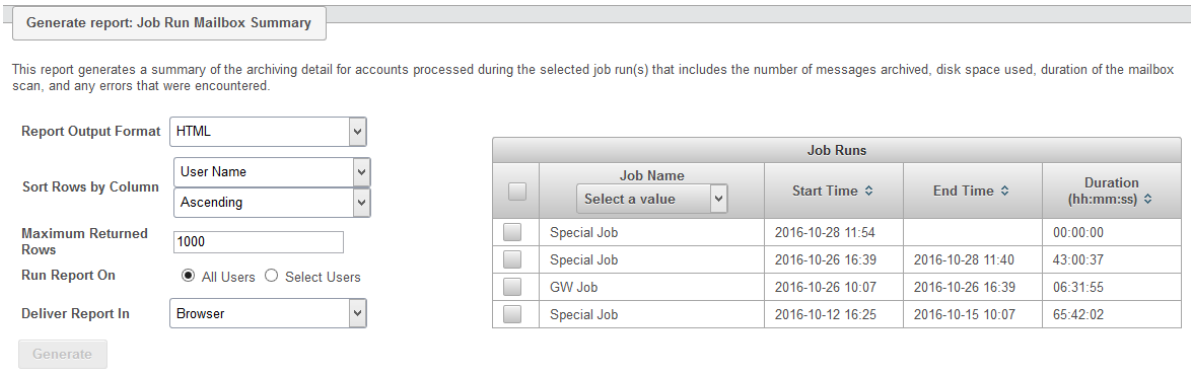
The Server utilization report details the memory, CPU utilization, Archive size, Thread usage, Index size, heap memory, and non-heap memory usage all in quick to digest charts.

To schedule the report, select the ‘Schedule’ button from the top of the report to utilize the same settings.



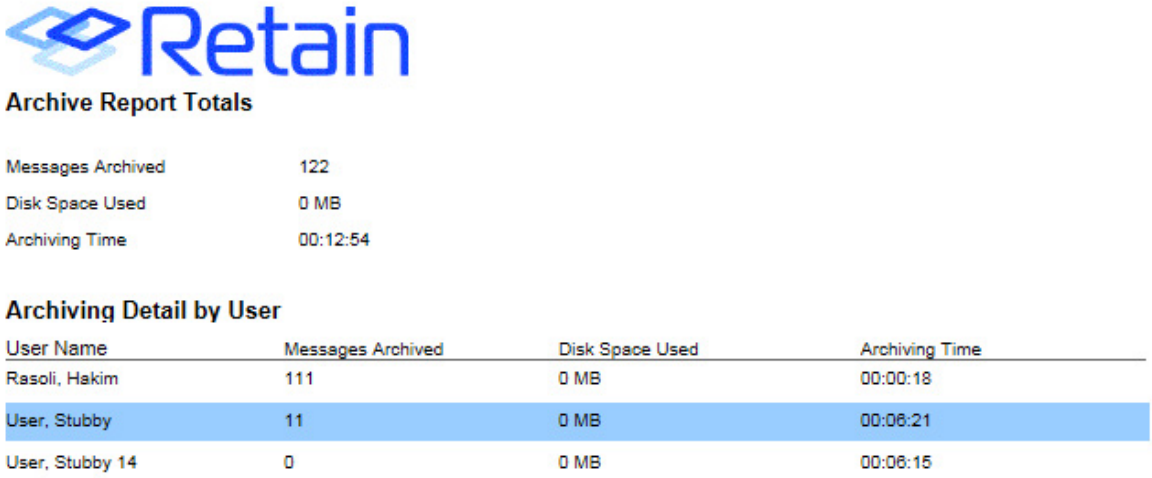
6.5.10 Job Run Mailbox Summary

The Job Run Mailbox Summary report is much like the Job run Summary except it is dedicated to specific mail-boxes and specific jobs. This job requires both a job, and a user or users to be selected.



Select the desired job and either all users, or only select users. If desired, the columns may be sorted by a selected value through the ‘Sort Rows by Column’ drop-down menu. Once set as desired, select the ‘Gen-erate’ button.

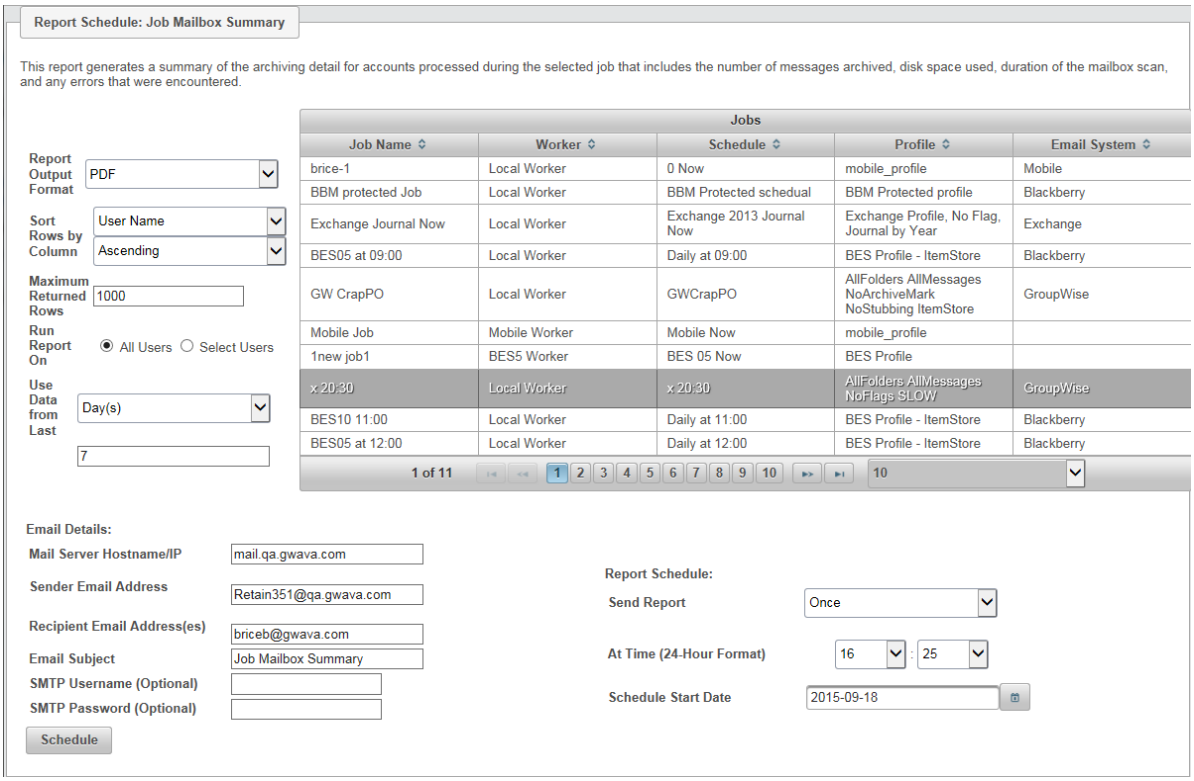
The report shows the details for the selected job and mailboxes. The user name, messages archived, disk space used, and archiving time is all displayed for each mailbox as well as for the total job.



This report cannot be scheduled.

6.5.11 Job Mailbox Summary

The Job Mailbox Summary report is a scheduled report only, that displays the number of messages archived, the disk space used, and the time the job took, separated by separate mailboxes. This report may be restricted to a date range, a limited amount of rows, and may be performed on all, or only select users. This report may also sort the data according to a selected column, in either ascending or descending value.



Select the ‘Schedule’ button to create the report.



Archive Report Totals

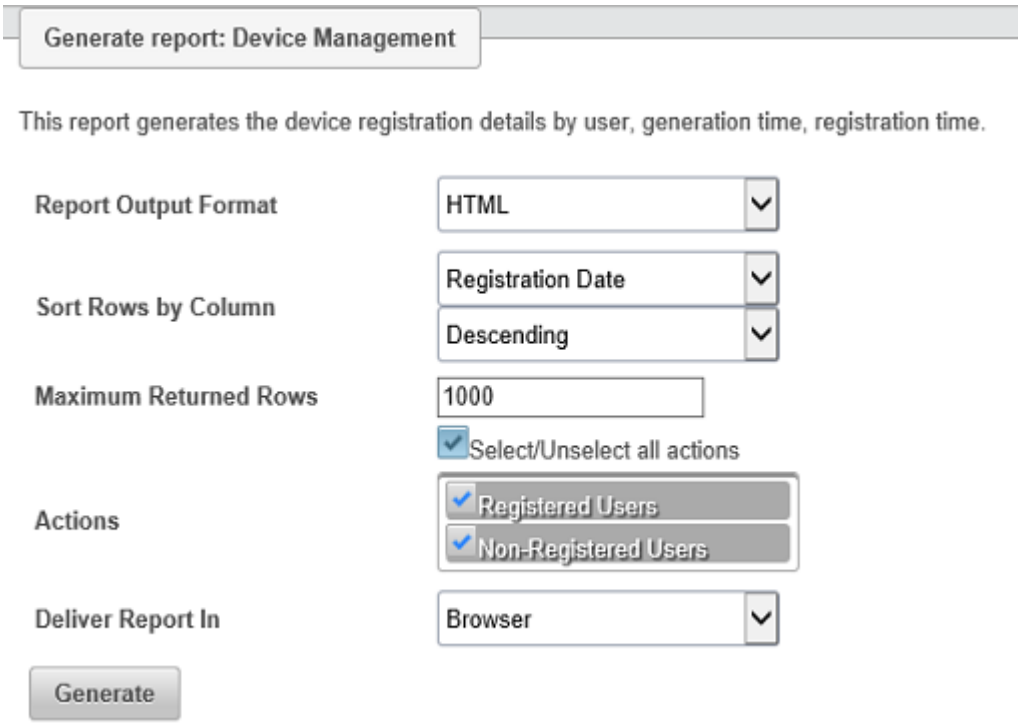
Messages Archived	1084
Disk Space Used	27 MB
Archiving Time	03:19:18

Archiving Detail by User

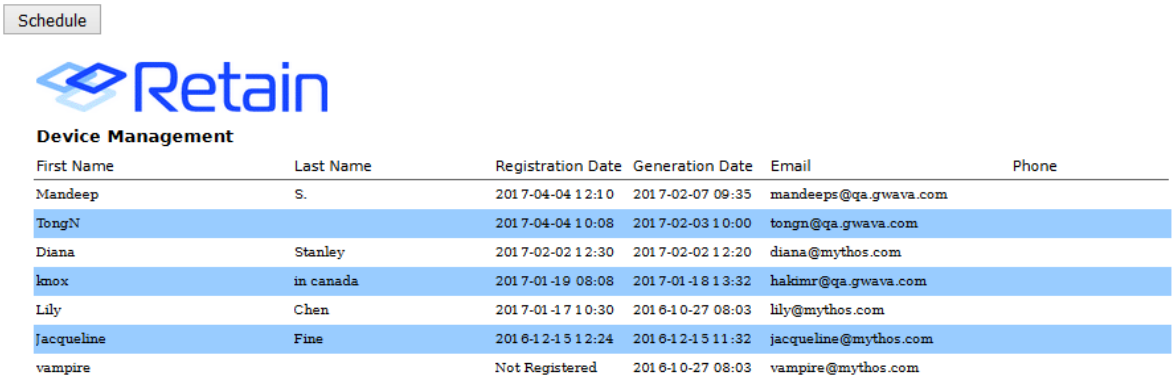
User Name	Messages Archived	Disk Space Used	Archiving Time
aa	0	0 MB	00:00:18
Administrator	0	0 MB	00:00:04
Allen, Jared	4	0 MB	00:06:59
Alungh, Suman	160	4 MB	00:00:41

Device Management

The device management report generates the device registration details by user, generation time, registration time.



The resulting report will show Non-registered Users with a Registration date of Not Registered until the registration process is complete on their device. The Generation Date shows when the registration process was begun.



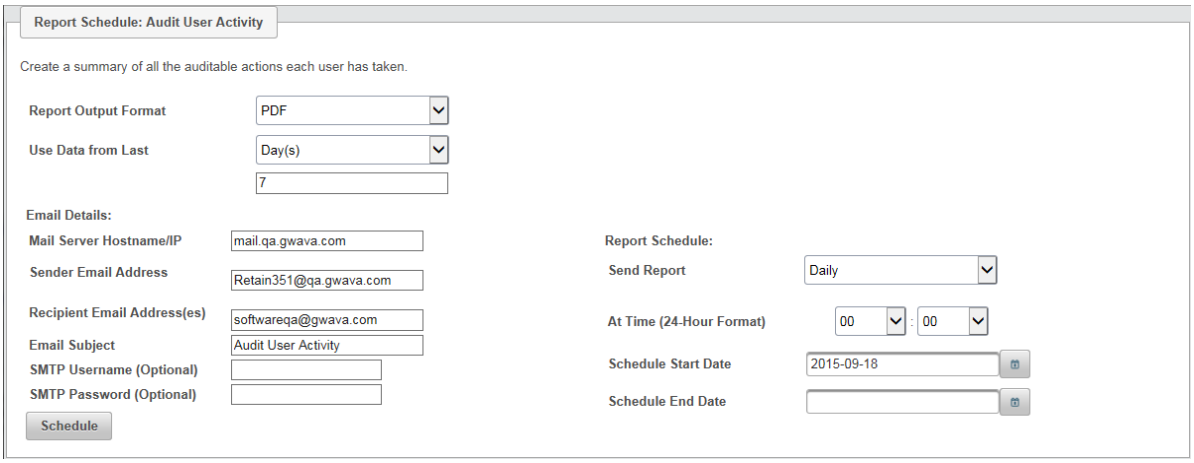
First Name	Last Name	Registration Date	Generation Date	Email	Phone
Mandeep	S.	2017-04-04 12:10	2017-02-07 09:35	mandeeps@qa.gwava.com	
TongN		2017-04-04 10:08	2017-02-03 10:00	tongn@qa.gwava.com	
Diana	Stanley	2017-02-02 12:30	2017-02-02 12:20	diana@mythos.com	
knox	in canada	2017-01-19 08:08	2017-01-18 13:32	hakimr@qa.gwava.com	
Lily	Chen	2017-01-17 10:30	2016-10-27 08:03	lily@mythos.com	
Jacqueline	Fine	2016-12-15 12:24	2016-12-15 11:32	jacqueline@mythos.com	
vampire		Not Registered	2016-10-27 08:03	vampire@mythos.com	

6.5.12 Schedule Report

Scheduling reports is accomplished through the Schedule Report interface. Every report which can be scheduled is also available as an instant report with a ‘schedule report’ button located at the top of the HTML report. Because these instant reports are easier to fine-tune, it is recommended to first configure a scheduled report through the ‘Report’ interface first, to allow immediate fine tuning and

viewing of the results. Once achieved, selecting the ‘Schedule’ button on the report will take all settings used to gener-ate that report and input them into the schedule Report interface. To schedule any report, a destination email address, SMTP information, and a set schedule are required.

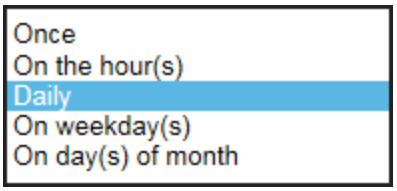
Scheduled reports differ from normal reports in that, due to the scheduled nature, they must be given a time to generate the report and be given an email address to deliver the report to. Reports are delivered as an attachment in an email sent to a specified user. As such, scheduled reports are only offered in PDF and XLS formats.



Reports which have a time frame attached to them as part of the criteria will have a floating time window which they may take information from. The configuration allows reports to the use data created in the past number of hours, days, weeks, or months. The number may be set by the user.



Also, the frequency of the report may be configured to be either a single instance or a reoccurring report. The frequency allows for hourly, multiple times on specified hours every day, daily, weekly or on any specified number of days in the week, or at any specified day or days in the month.



In addition, a scheduled start and end date allow the administrator to control exactly when the start and stop dates for any specific report can be. The end date is not required, but is optional.

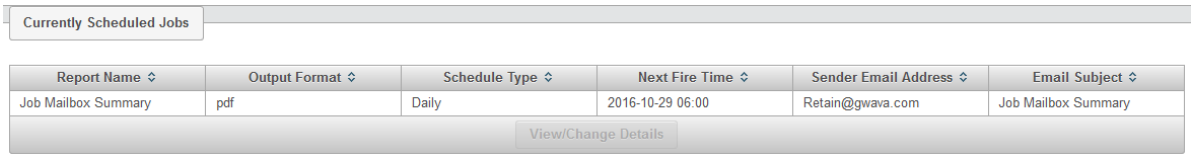
Manage Scheduled Reports

Once reports have been scheduled, it may be necessary to either cancel or modify the frequency or set-tings of the report. The Manage Scheduled Reports interface allows jobs to be immediately run, modi-fied, or removed.



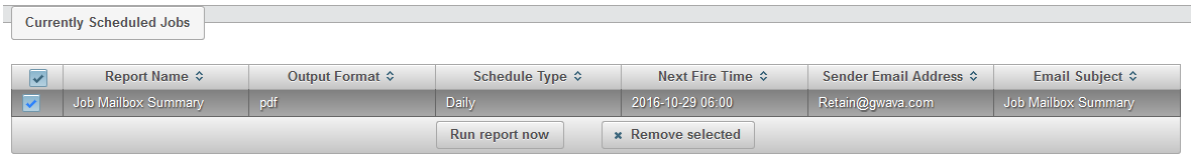
As with all other columns in the R&M interface, the columns are sortable and even searchable to allow quick access to any desired Report through any of the scheduled criteria. For instance, a user which doesn't want to receive a report anymore, but doesn't know what report is being delivered to their email, only needs to know the email subject, or the next delivery time, to immediately identify the report. Unique sender addresses or subjects will readily allow any report to be identified.

To modify or view a report's settings select the 'View/Change details' window.



Once the desired report has been located, select the report and then select the 'View/Change details' button at the bottom to be taken to the Report's configuration. Make any desired changes and select the 'Update' button at the bottom to save the changes.

To remove or immediately run a scheduled report, select the 'Run and Delete reports' window.



Any report or reports may be immediately run or removed from the system through this window. All scheduled reports will be displayed. A select all/none box is at the top left of the table, and allows for a quick selection or clearing of all previous selections which have been made.

To run a report or reports, select the desired report by clicking on the associated checkbox and then select the 'Run report now' button at the bottom. The report will be immediately triggered, and emailed to the specified recipient.

To remove a report or reports, select the desired report(s) by placing a check in the associated checkbox. Once the reports are selected, click the 'Remove selected' button at the bottom of the table.

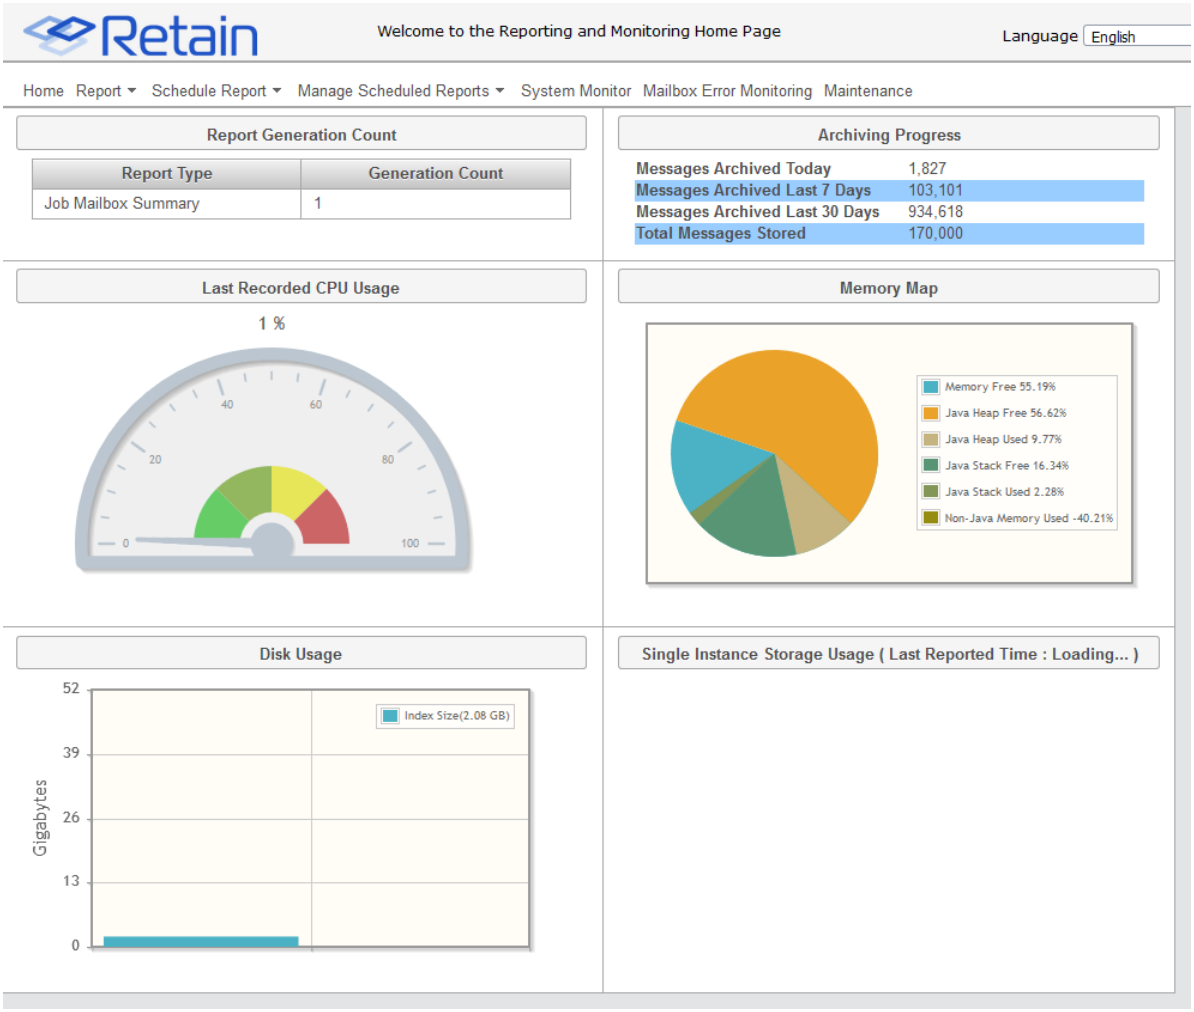
6.5.13 System Monitor

The System Monitor page is a quick, graphed page which displays the essential stats of the Retain Server system. The R&M module statistics are displayed, with how many reports have been created and of what type. The basic archive progress is also displayed showing the daily, weekly, monthly, and total amount of messages which have been archived in the system.

The system CPU load, memory usage map, and Disk Usage are displayed. The dashboard also contains statistics for disk space saved through single storage instances. These numbers are for the actual host server, and not just the Retain system, with exception of the memory map, which has

both. The Memory map is broken out to total system memory free and used, but also includes Java memory statistics – which will be Retain specific unless additional Java applications are running on the Retain Server. This page will not time-out or expire, to allow quick monitoring of the system.

NOTE: A problem has recently been identified with gathering some system-level variables for the System Monitor page. Affected variables include CPU, RAM, and system memory usage. To avoid additional problems these variables have been set to '0' until the issue is resolved.

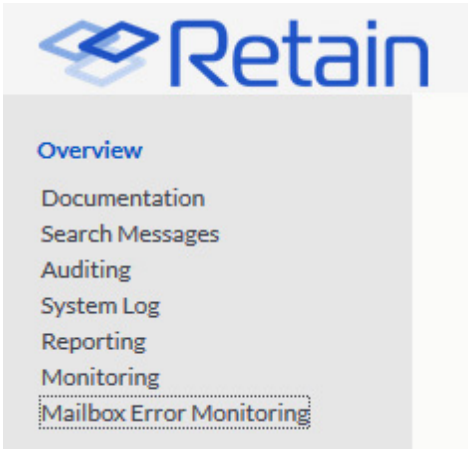


6.5.14 Mailbox Error Monitoring

Mailbox Error Monitoring is a function that allows the Administrator to deal with message errors. The email system reports errors to Retain when there is an issue accessing a message. Sometimes these errors are transient and other times they are permanent. This function will help the Administrator determine the next action.

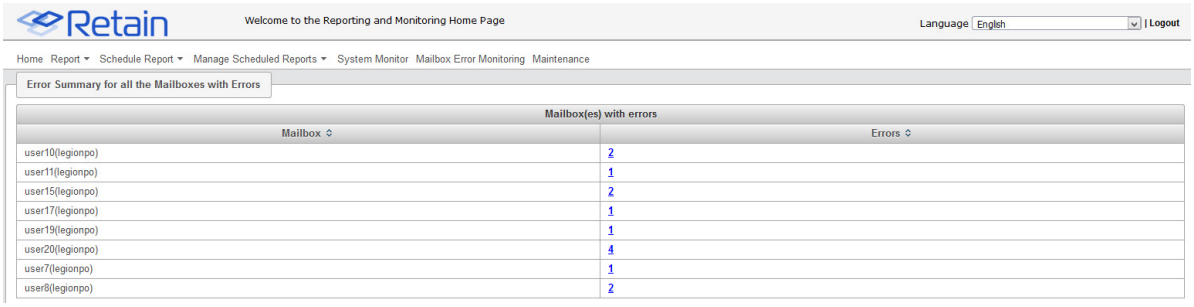
This function requires the Reporting and Monitoring server to be installed.

Find "Mailbox Error Monitoring" under the Overview section of the left navigation bar.

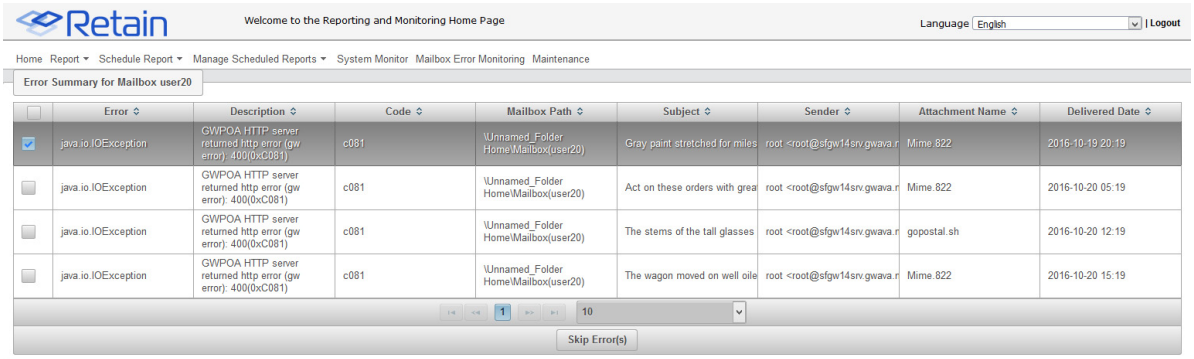


Clicking on "Mailbox Error Monitoring" will open a new window or tab. This will show the Mailbox Error Monitoring page of the Reporting and Monitoring server.

"Error Summary for all the Mailboxes with Errors" shows a list of mailboxes with errors.



To access the errors click on the error number, which is a link that opens a new window or tab. This window will show the "Error Summary for Mailbox [mailboxName]" list.



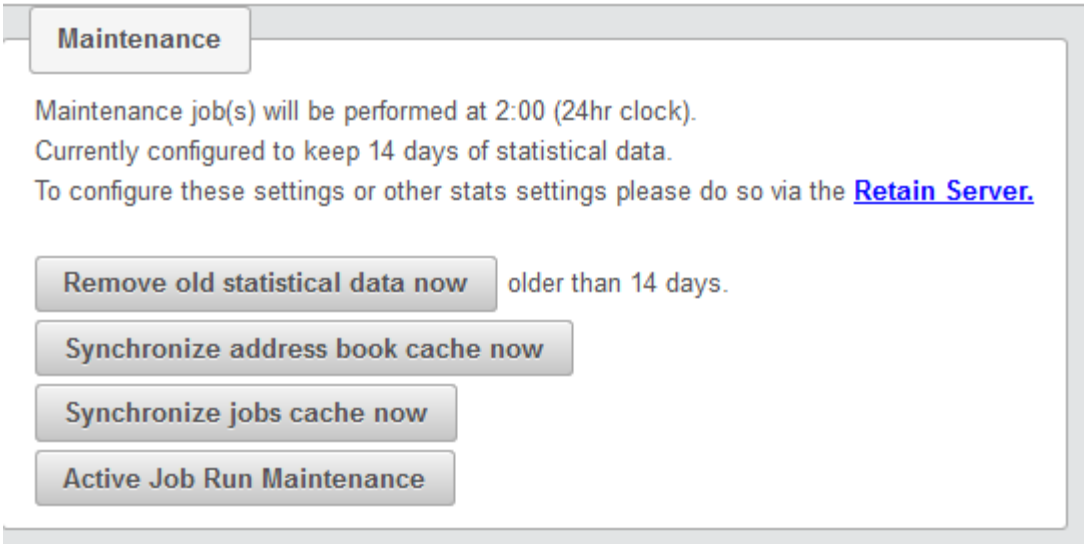
Each line will include: the error, a description of the error, the error code, the mailbox path to the folder the message was found in, the subject, the sender, the attachment that the error occurred on, and the delivered date.

If the message is accessible from the user's client, save the message from the client. Import the message into Retain using the PST Importer or GroupWise Personal Archive Migrator.

To skip a message, select a message by clicking the checkbox by a message, or use the title bar checkbox to select all. Click on the "Skip Error(s)" button to have Retain skip those errors the next time the job runs. This will remove the message from the list.

6.5.15 Maintenance

The Maintenance page contains the basic maintenance for the R&M module. This is not maintenance for the entire Retain Server. Because the R&M module maintains individual statistics, it needs to be refreshed and purged at regular intervals. The defaults are shown.



Remove old statistical data now: By default, the statistics server only maintains a record of statistics for 14 days, and all data older than 14 days will be automatically expired and the data deleted. If the statistic data is required for longer or shorter, configure and save the setting. It is not recommended to maintain statistical data older than 30 days.

Synchronize address book cache now

The R&M Module maintains a separate record of the address book and the job cache. The system auto-matically synchronizes this information with the Retain Server every 24 hours. Still, the different records may become out of sync. If a synchronize job is required before regular maintenance, a manual sync may be individually triggered.

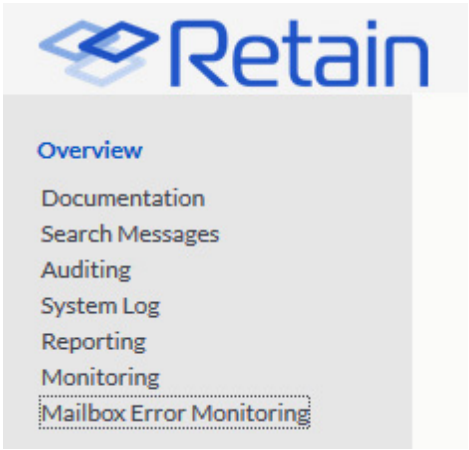
Active job run maintenance: If a job has completed but R&M thinks it is still running, press this button to clear.

6.6 Mailbox Error Monitoring

Mailbox Error Monitoring is a function that allows the Administrator to deal with message errors. The email system reports errors to Retain when there is an issue accessing a message. Sometimes these errors are transient and other times they are permanent. This function will help the Administrator determine the next action.

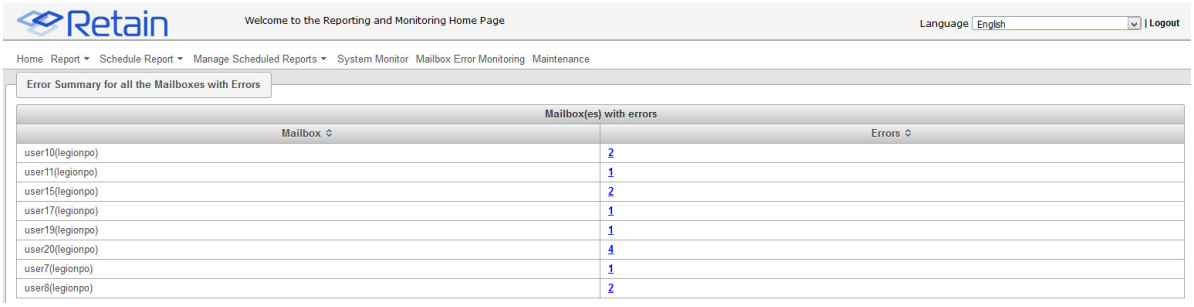
This function requires the Reporting and Monitoring server to be installed.

Find "Mailbox Error Monitoring" under the Overview section of the left navigation bar.

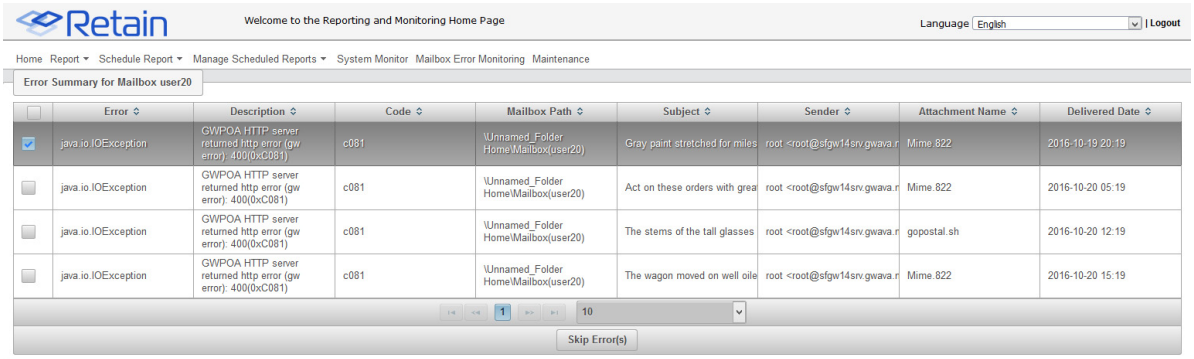


Clicking on "Mailbox Error Monitoring" will open a new window or tab. This will show the Mailbox Error Monitoring page of the Reporting and Monitoring server.

"Error Summary for all the Mailboxes with Errors" shows a list of mailboxes with errors.



To access the errors click on the error number, which is a link that opens a new window or tab. This window will show the "Error Summary for Mailbox [mailboxName]" list.



Each line will include: the error, a description of the error, the error code, the mailbox path to the folder the message was found in, the subject, the sender, the attachment that the error occurred on, and the delivered date.

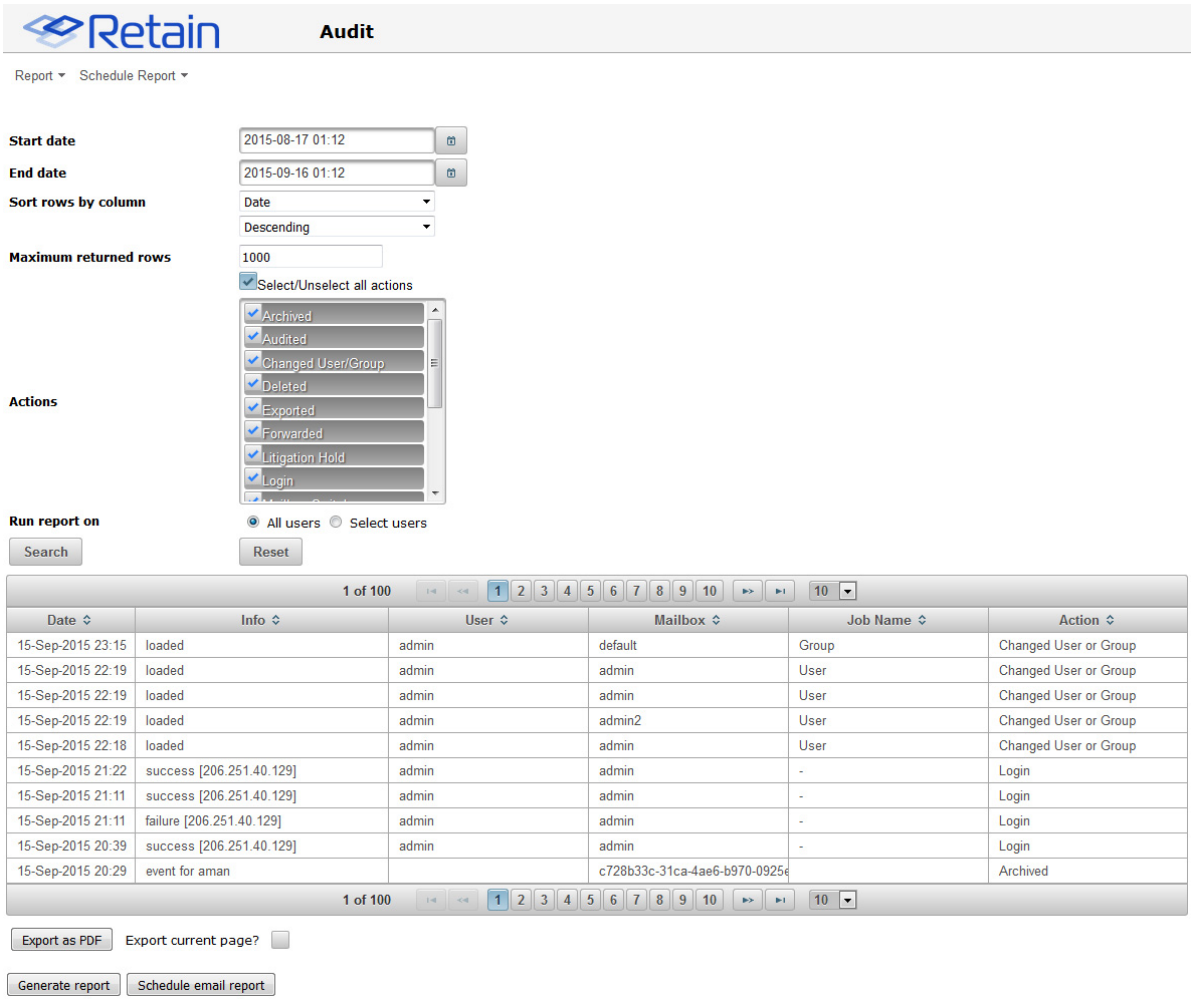
If the message is accessible from the user's client, save the message from the client. Import the message into Retain using the PST Importer or GroupWise Personal Archive Migrator.

To skip a message, select a message by clicking the checkbox by a message, or use the title bar checkbox to select all. Click on the "Skip Error(s)" button to have Retain skip those errors the next time the job runs. This will remove the message from the list.

6.7 Audit and System Logs

6.7.1 Audit Log

All pertinent actions performed by users in Retain are logged and can be discovered and printed. If Reporting and Monitoring is installed, four buttons are added to the following page, as shown below.



- Accessing the audit log requires the “Access all audit logs” administrative right.
- Ordinary users cannot access their audit logs unless you grant the “View Personal Audit Log” right, in which case they can find the audit log on a per-message basis under the Properties tab.

To view actions, select the desired actions, one or all, a date range, and a user, one or all, and then select ‘Search’. The results will be displayed in a table below. All columns in this table are sortable.

The results appear below. This may be printed by selecting the ‘Export as PDF’ button at the bottom, and printing the resulting PDF. The PDF export button will only display what is shown in the action table, not the criteria or the interface.

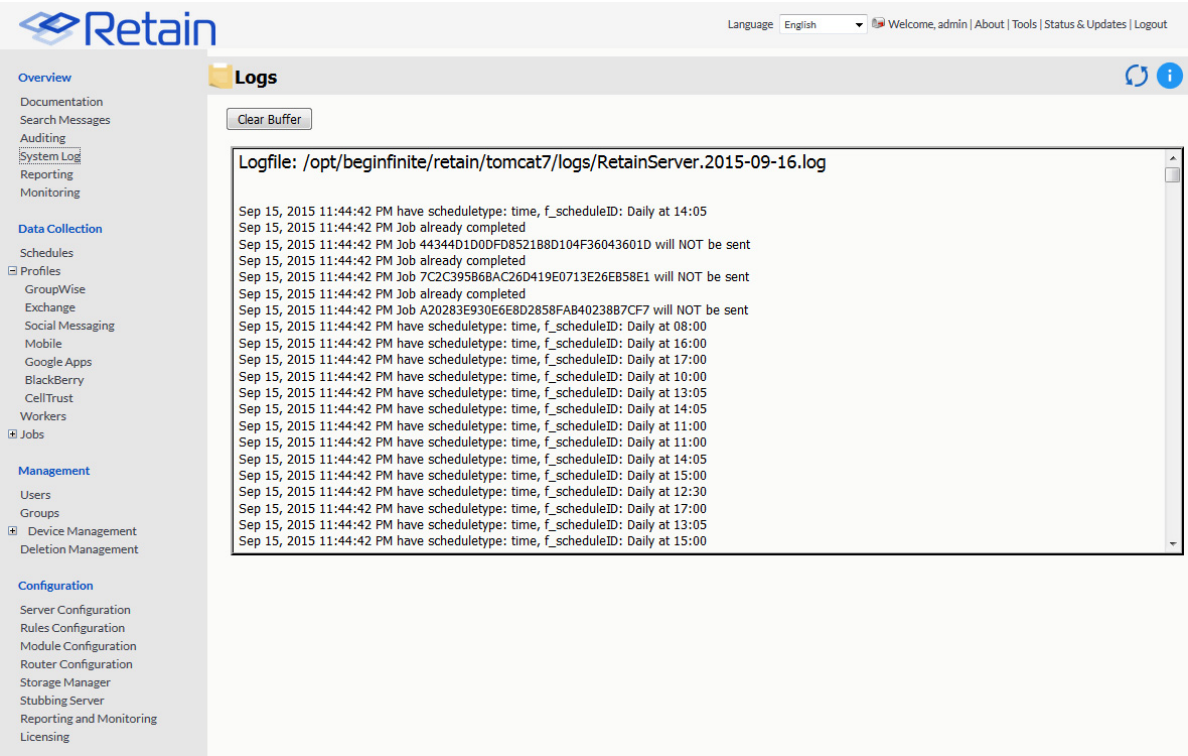
If you have rights, you may click on the Action link to open the message. An auditor does not automati-cally gain the rights to READ messages in mailboxes that are not assigned to them.

If the Reporting and Monitoring module is installed, the ‘Report’ and ‘Schedule Report’ menu options on the top bar, and the ‘Generate Report’ and ‘Schedule email report’ buttons along the bottom will appear. These buttons link to the R&M interface to extend R&M ability to generate and schedule reports based on the selected criteria. If the Generate report or Schedule email report buttons are pressed, the criteria used in the Audit window will be automatically propagated to the Reporting interface.

6.7.2 System Log

This option lets you read and parse through the latest system log file. The level of detail is controlled in system configuration.

The actual filename and path on the server is also listed. You may clear the view buffer at any time in case you want to monitor a particular operation for troubleshooting purposes.



6.7.3 Location of log files

At the end of each day the logs are closed and compressed to save space. After 10 days, by default, the logs are removed from the system.

The full length logs can be found in, be default:

- **Linux:** /opt/beginfinite/retain/tomcat8/logs
- **Windows:** C:\Program Files\Beginfinite\Retain\Tomcat 8.0\logs

There are several log files that are stored here, the ones of primary interest are:

- **RetainServer.<date>.log:** Contains server transactions.
- **RetainWorkerN.<date>.log:** Contains worker transactions. N being the worker number.
- **Indexer.<date>.log:** Contains indexer transactions.

6.8 Data Removal

At some point the data in Retain will reach the end of its lifespan or otherwise need to be removed from the system. Formerly known as deletion management.

- Mailbox Deletion: Users leave the company and eventually their mailboxes may be removed. Retain's Data Removal section is able to accomplish these tasks.
- Item Deletion: The data retention period has been met and you are free to dispose of the excess data en masse, which will free up considerable disk space.
- Delete Messages: Individual emails may also be removed from Retain.

6.8.1 Mailbox Deletion

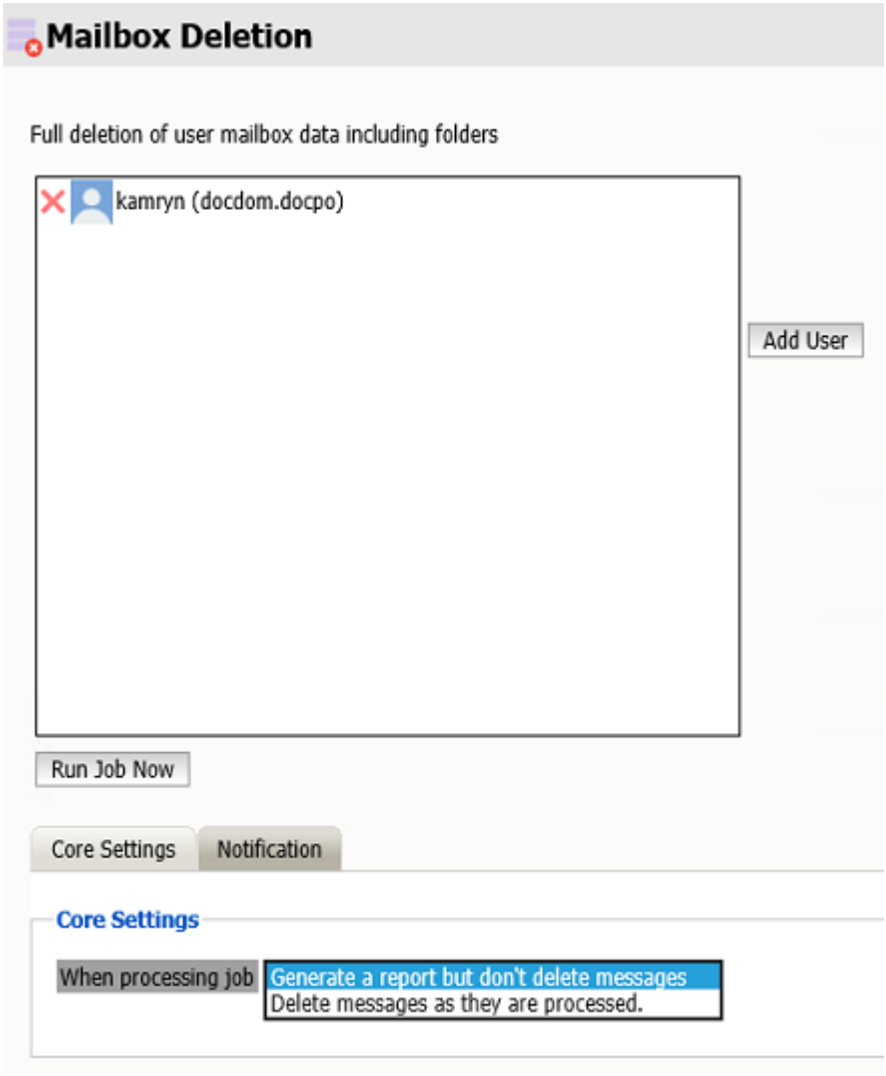
Mailbox Deletion will delete one or more entire mailboxes from Retain.

Add User

The primary action is to press the *Add User* button to select users to be deleted from Retain.

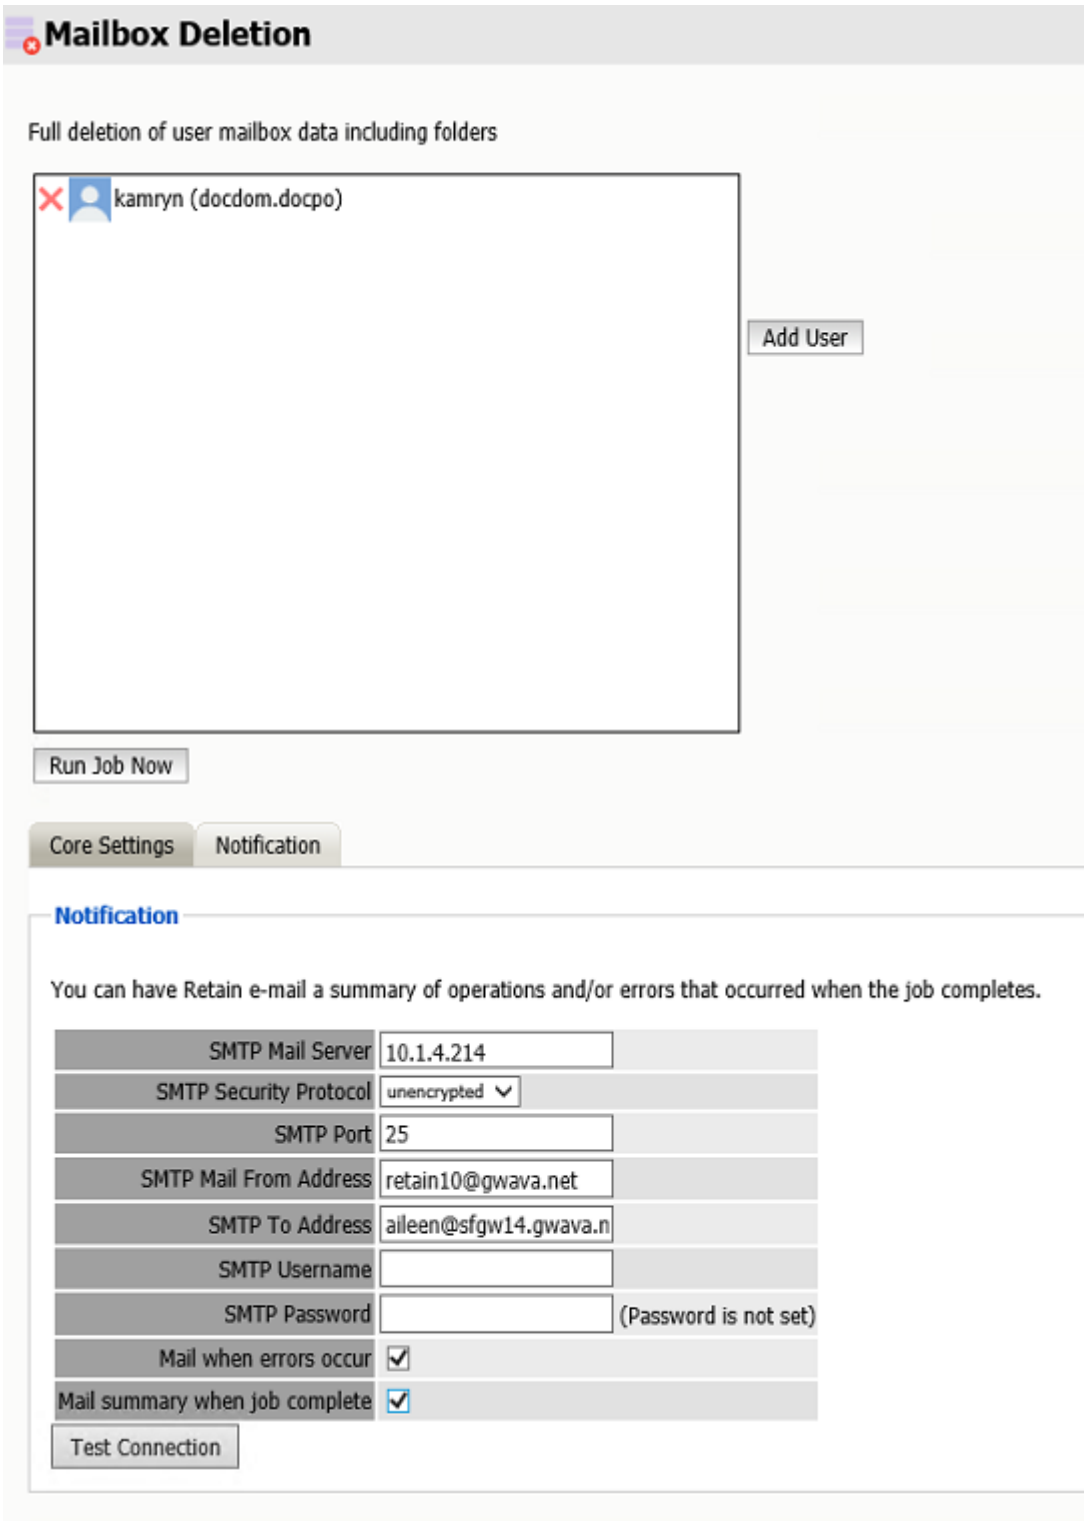
Core Settings

Core Settings allows you to choose between *Generate a report but don't delete message Delete messages as they are processed* and *Delete messages as they are processed* when processing the job.



Notifications

To be notify about errors and/or when the job is complete, fill out this tab. By default it is filled in with the notification data from Server configuration.



Press the *Run Job Now* button to begin the process.

6.8.2 Item Deletion

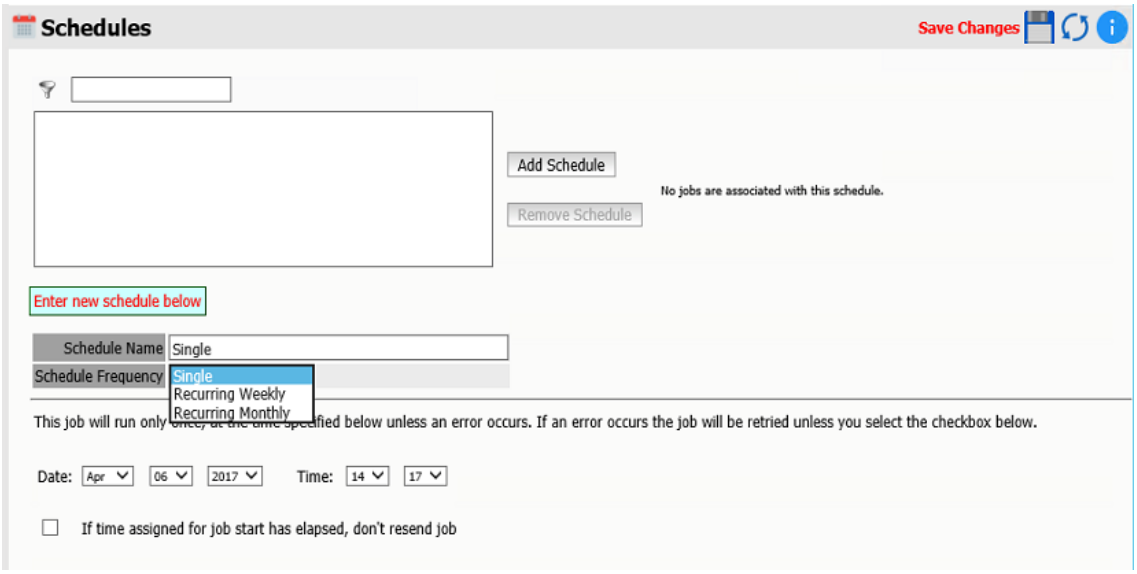
When items pass out of their retention period it is time to remove them from the system.

An item deletion job is set up much like an archive job with a schedule, profile and job, but there is no need for a worker because Retain can connect to itself just fine.

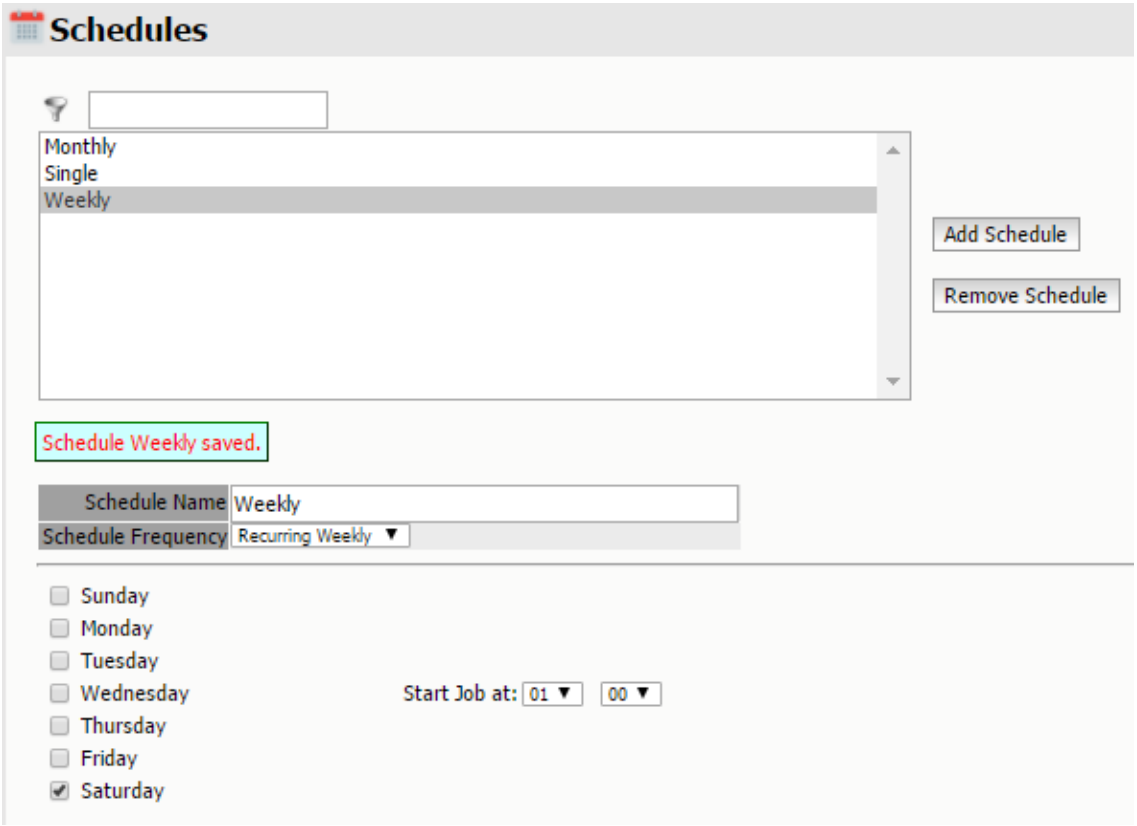
Schedule

There are three types of schedules:

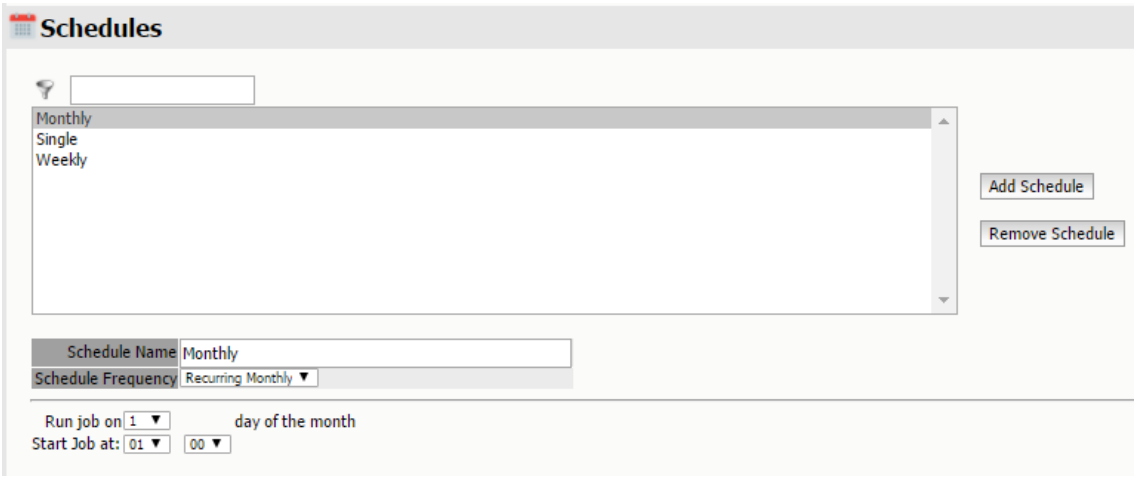
- **Single** , the job will run only once at the assigned date and time



- **Recurring Weekly:** , the job will run each week on the assigned day and time



- **Recurring Monthly:** , the job will run each month on the assigned day and time

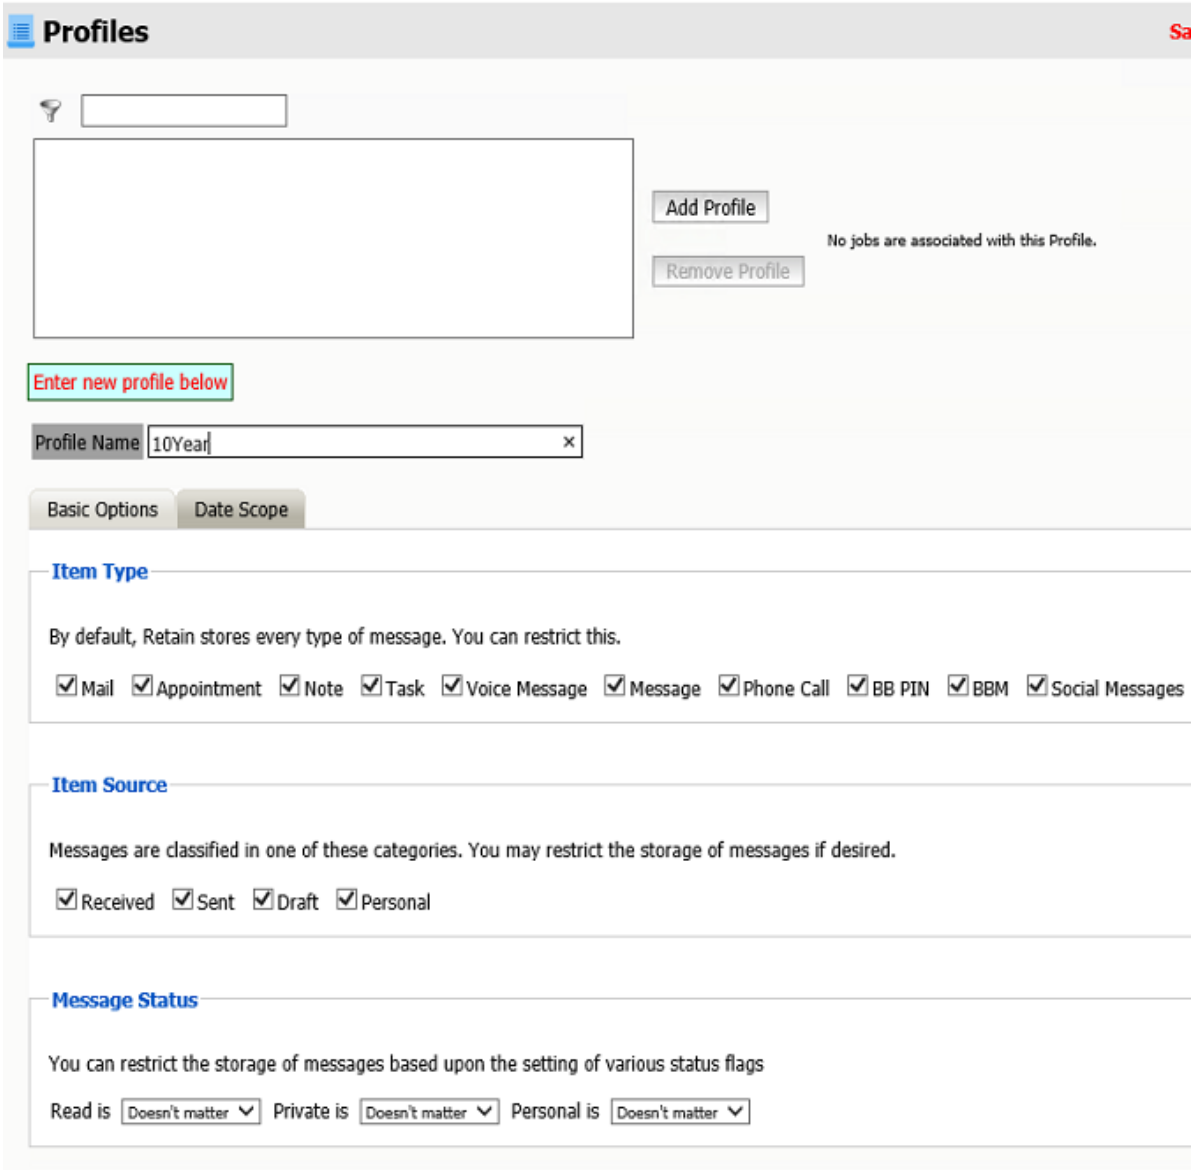


6.8.3 Profiles

The profile sets the type and age of the items to be deleted.

Basic Options

Sets the limits on the Type, Source and Status of the items to be deleted



- **Item Types:** include: Mail, Appointment, Note, Task, Voice Message, Message, Phone call, BB PIN, BBM, and Social Messages. Social messages include wall posts, chats, likes, comment, group posts, media, ratings, requests, events, status update, search, forward, and webmail.
- **Item Sources:** include: Received, sent, draft and personal.
- **Message Status:** can include Read, private or personal, depending on Doesn't matter, True or False.

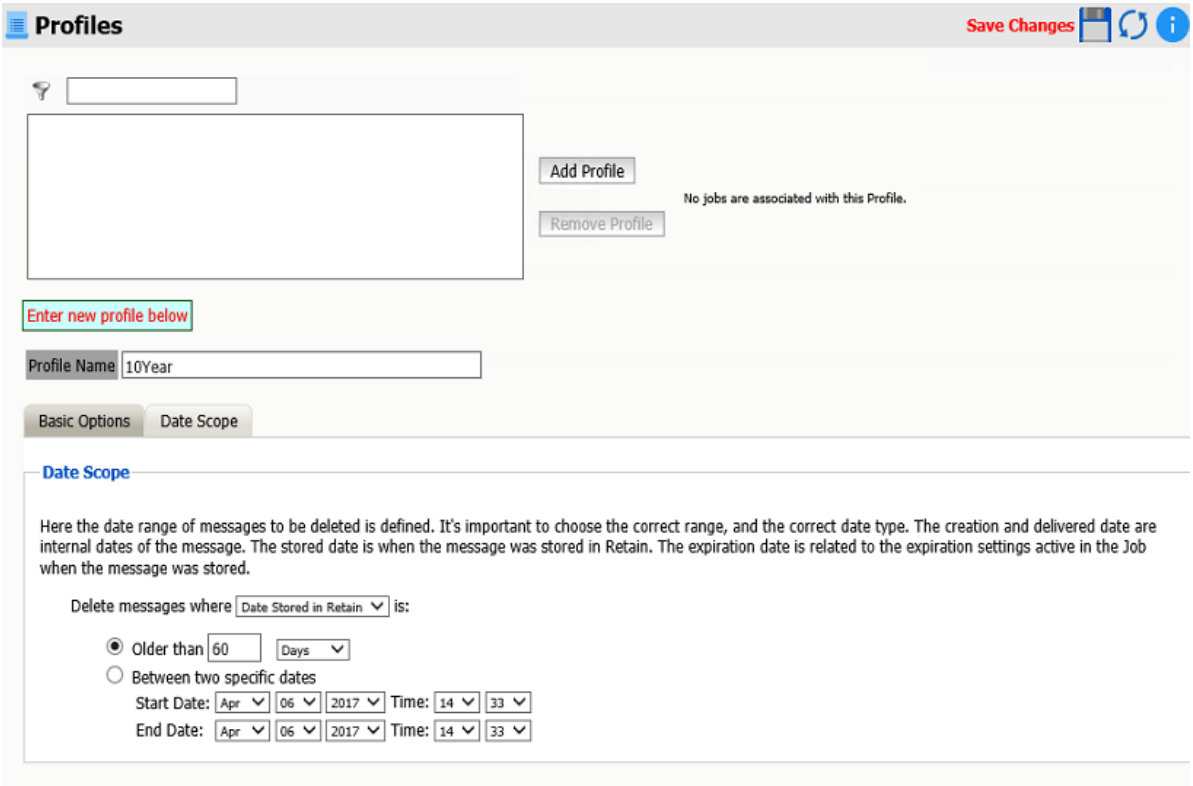
Date Scope

Items can be deleted by the

- **Date Stored in Retain:** , if you have been using Retain longer then the retention period this is good.
- **Delivered Date:** , this is usually the best choice as the receiving message system marks this and would be most accurate

- **Expiration Date:** , this is set by Retain during the archive job, if not set won't be much use.
- **Created Date:** , this cannot be trusted as it is set by the sender and spammers often set this to the far future.

And if the items are older than a specified number of Days, Months or Years or between two specific dates.



6.8.4 Job

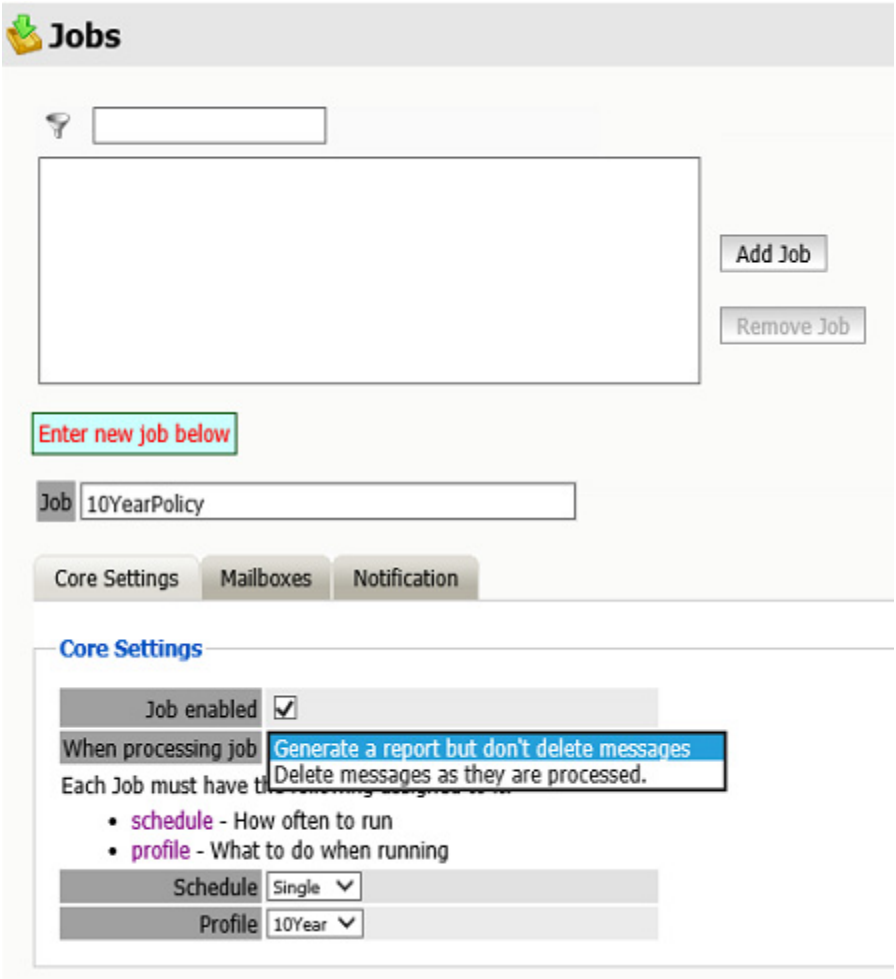
Creating a deletion job is just like creating an archive job select a schedule and profile to prepare the job.

Core Settings

A job can be enabled or disabled.

Core Settings allows you to choose between *Generate a report but don't delete message Delete messages as they are processed* or *Delete messages as they are processed* when processing the job.

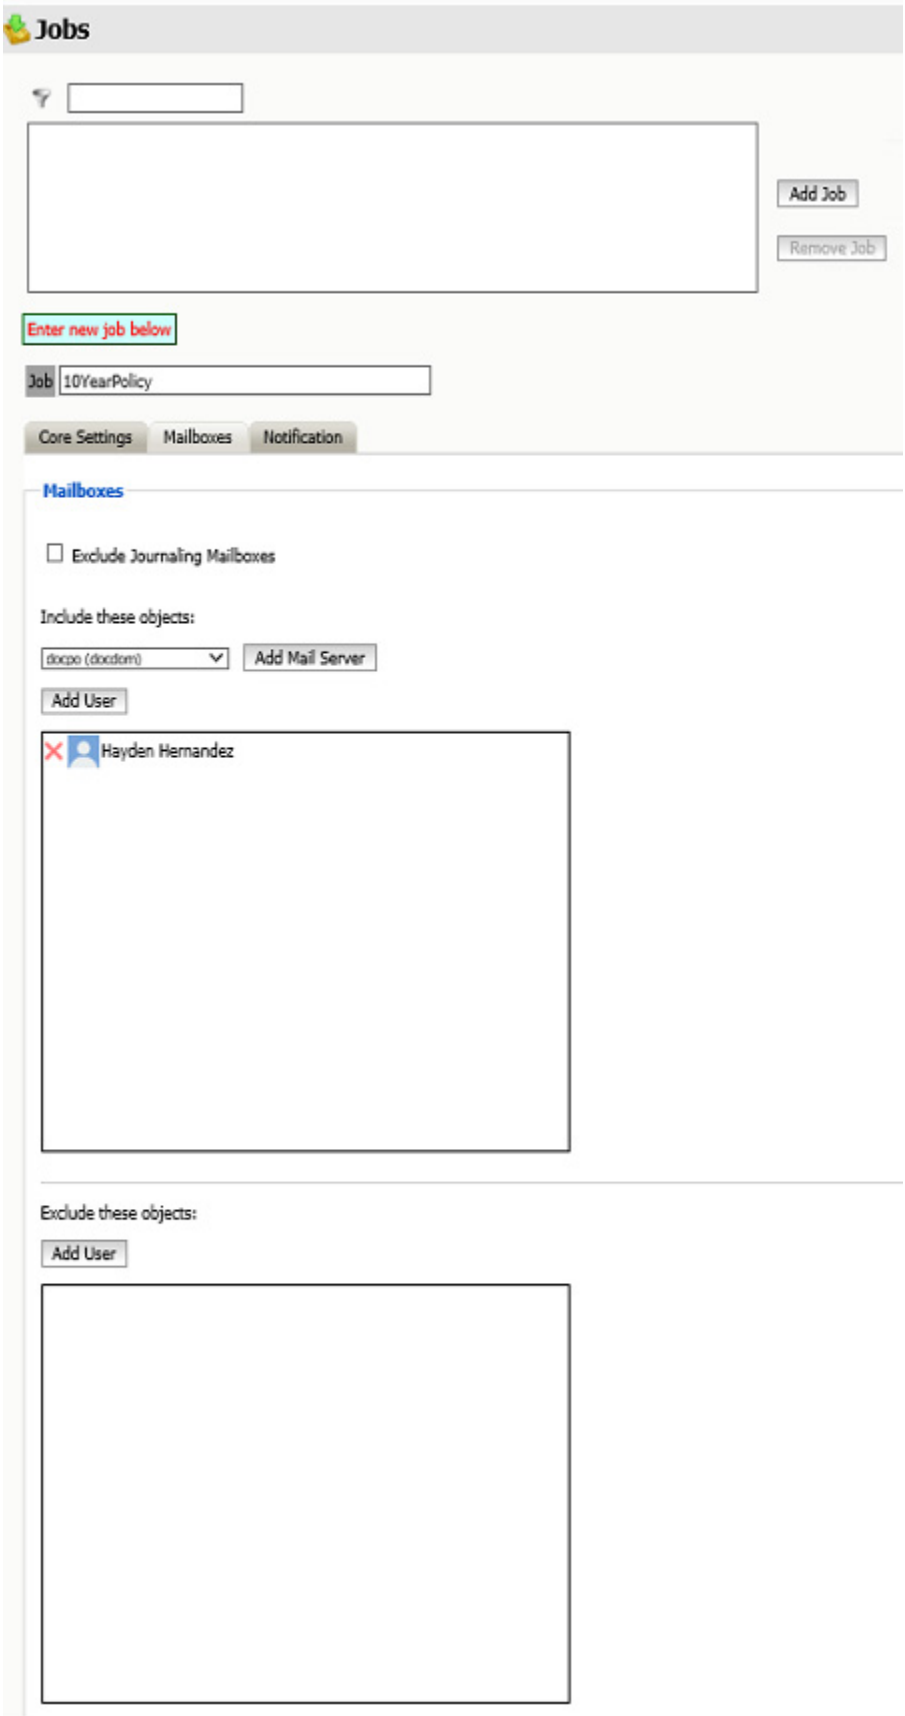
A Schedule and a Profile must be chosen.



Mailboxes

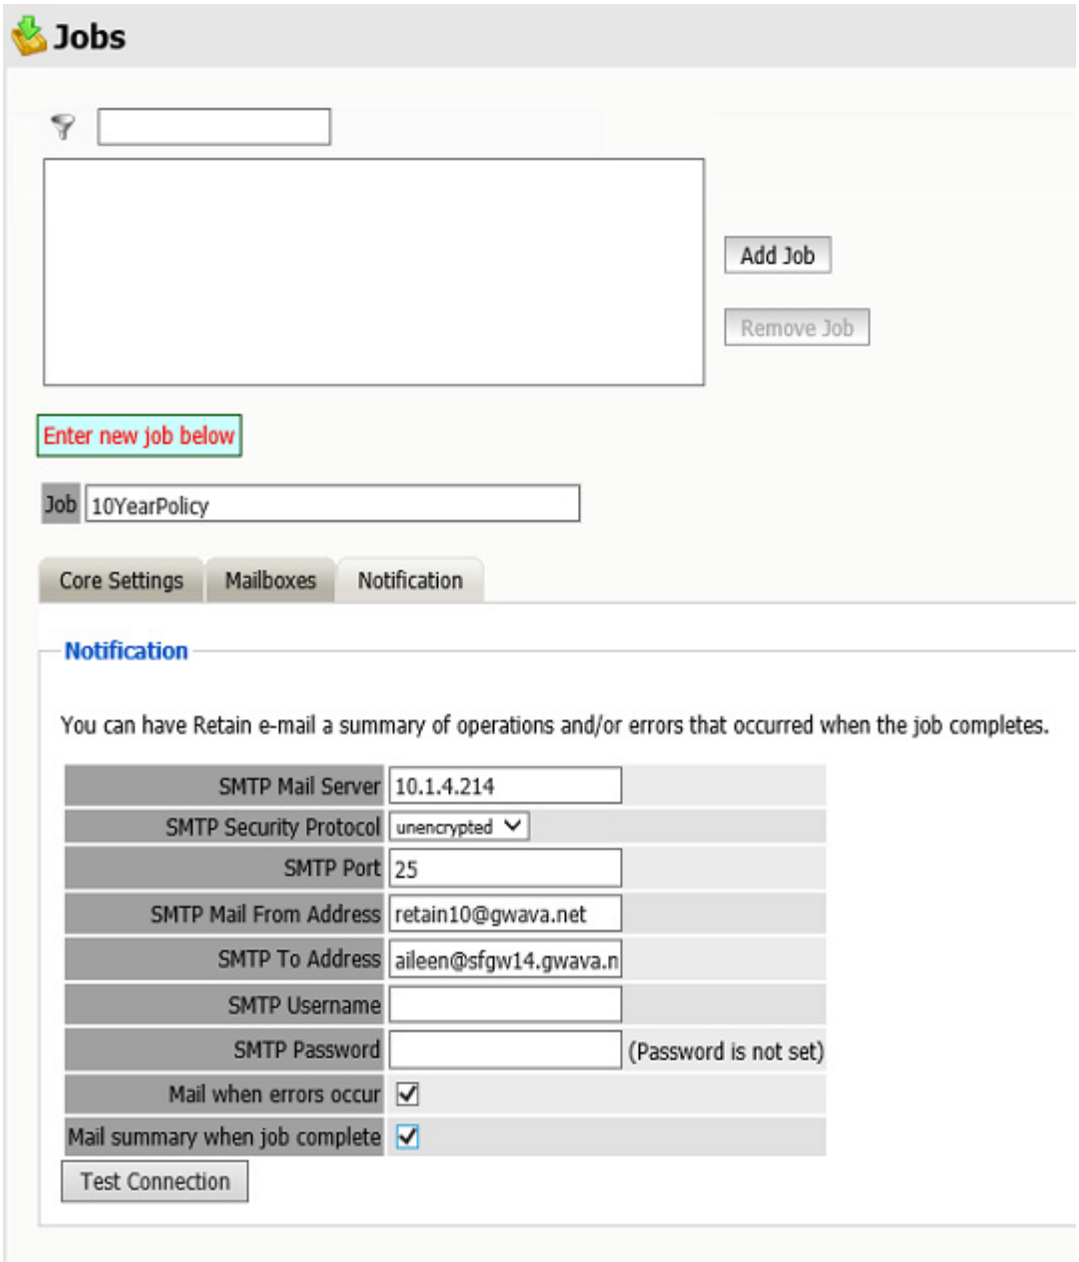
Select the objects to be included for the deletion job.

- Journaling Mailboxes may be excluded.
- An entire mail server or post office can be included.
- Individual users may be included or excluded.


Jobs
Add Job
Remove Job
Enter new job below
Job
10YearPolicy
Core Settings
Mailboxes
Notification
Mailboxes
Exclude Journaling Mailboxes
Include these objects:
Add Mail Server
Add User
Hayden Hernandez
Exclude these objects:
Add User

Notification

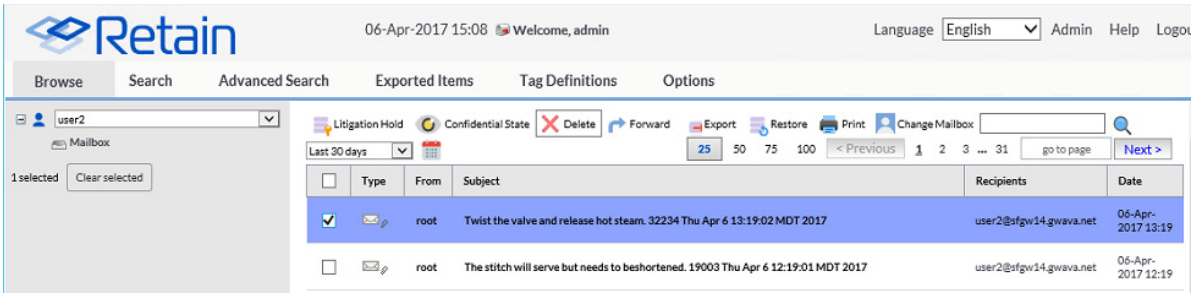
To be notify about errors and/or when the job is complete, fill out this tab. By default it is filled in with the notification data from Server configuration.



6.8.5 Delete Messages

Individual Messages can be deleted by the admin or a user given sufficient rights. Deletion actions will be logged in the RetainServer log.

Within the users mailbox one or more messages can be deleted by selecting with the checkbox and pressing the *Delete* button.



6.8.6 Deletion Logging

The job will be logged in the RetainServer log.

The resulting report CSV file will be placed in the archive directory.

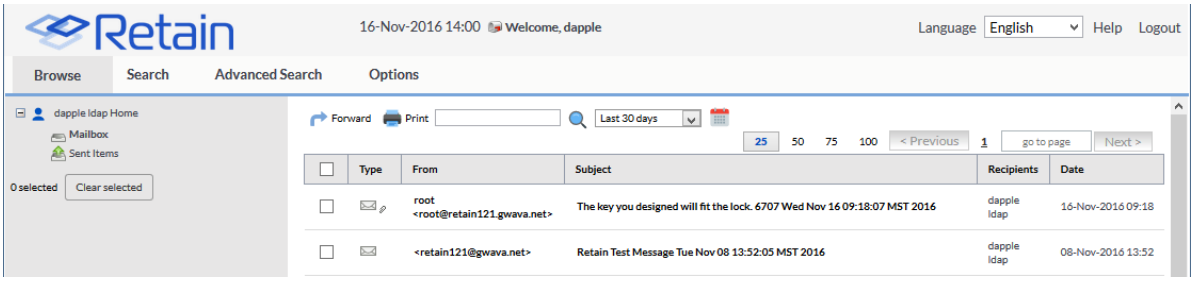
7 Retain User Guide

There are different ways to search the Retain Archives:

7.1 Retain Web Interface

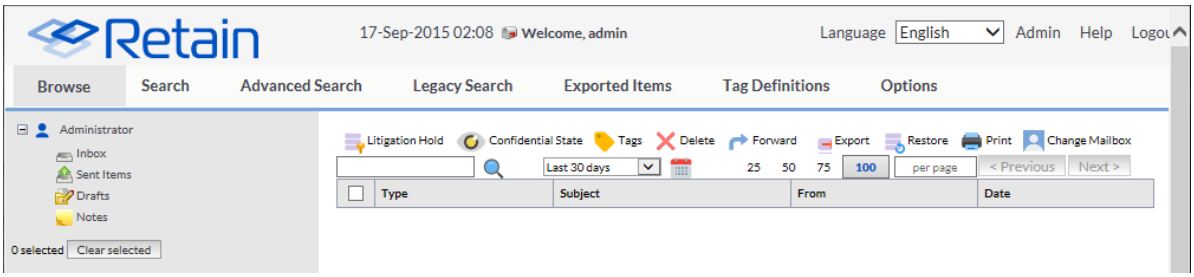
The Retain web search interface is the core of the information retrieval system for the Retain archive. Admin-istrators, auditors, and users search for what they want and can view, print, export, or forward the results. General users have the right to view their own archive, and will immediately be brought to the Retain mailbox interface, but limited to their own account.

A normal user will be directed to their Retain Mailbox. Users by default have the rights to View and Save attachments, and Forward and Print messages.



If a user has at least one administrative right assigned, the user will see the Retain Management console on login, and must select the ‘Search Messages’ option from the navigation window in order to access the archive.

The full administrator interface opens with the following tabs, users may or may not have access to all of these depending on the rights assigned by the admin:



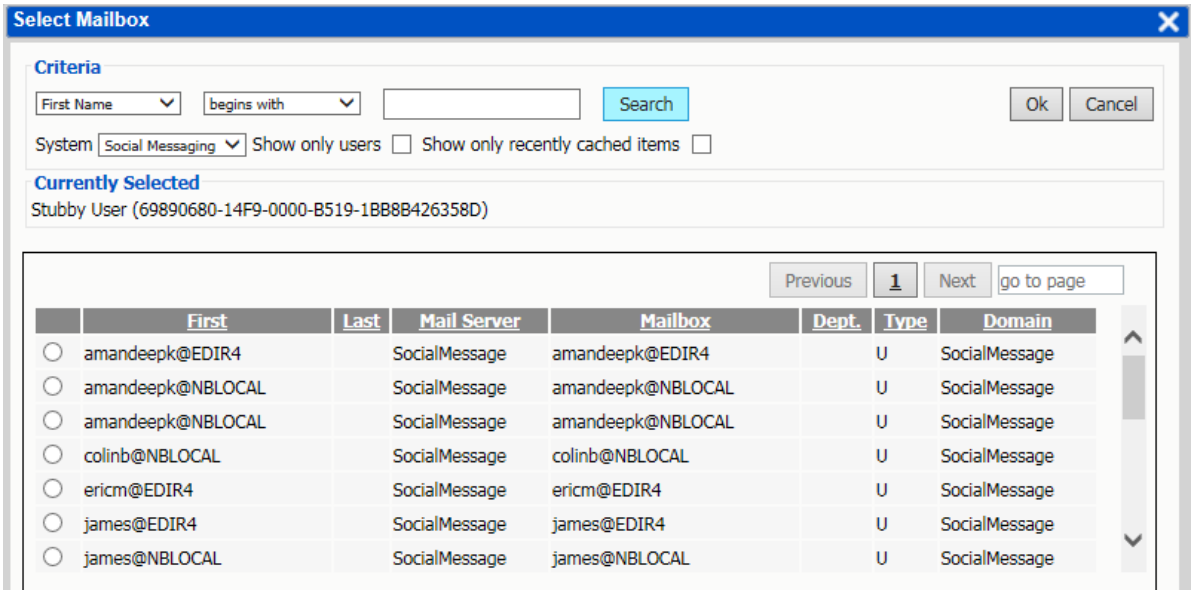
- Browse
- Search
- Advanced Search
- Legacy Search
- Exported Items
- Tag Definitions
- Options

7.2 Browse

The Browse tab immediately grants access to the selected mailbox. Individual users will only see their own mailbox, while users with the administrator right to search all mailboxes will have the option to change mailboxes to another user.

Retain is an archiving solution not a message management system. All items are stored in the location Retain found them in the production system and cannot be moved.

To change to a different mailbox, select the ‘Change Mailbox’ button in archive toolbar. Once clicked, the ‘Select Mailbox’ dialog will open.



The Select Mailbox dialog asks for specific information. The Select which mailbox you want to see by click-ing on its radio button and clicking OK.

When searching for a mailbox, the system of the mailbox must be selected as mailboxes from different systems may have the same user name or criteria. Select which mail system the desired user belongs to, specify any further criteria, or leave the criteria blank to display all possible mailboxes from that system.

(If Retain for Social Messaging is set to anonymous user, all Social Messaging data will be contained under the single user ‘?@?’, and separate user names and pertinent information is contained in the ‘from’ dia-log. Otherwise, individual user accounts will be displayed.)

If the search results are extensive, the system will have a ‘Next’ or ‘Previous’ button at the bottom of the search results window, which displays the next set of results.

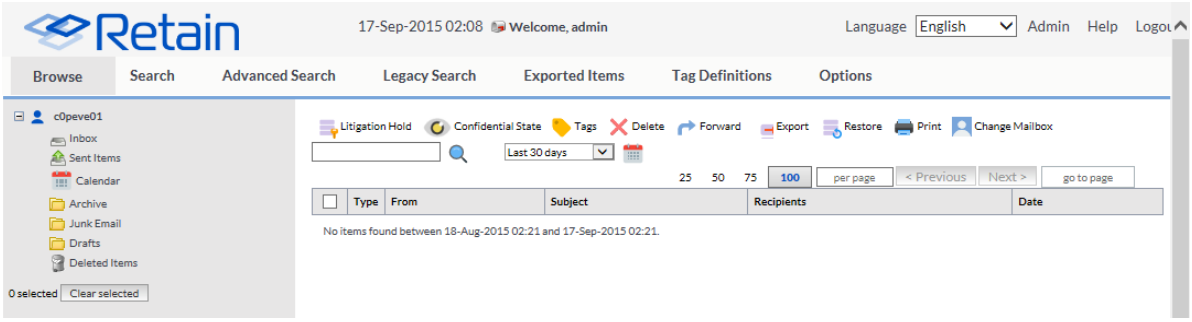
Refine search parameters to reach a manageable search result.

In Retain, you can only browse one mailbox at a time. To search multiple mailboxes simultaneously, use the search function.

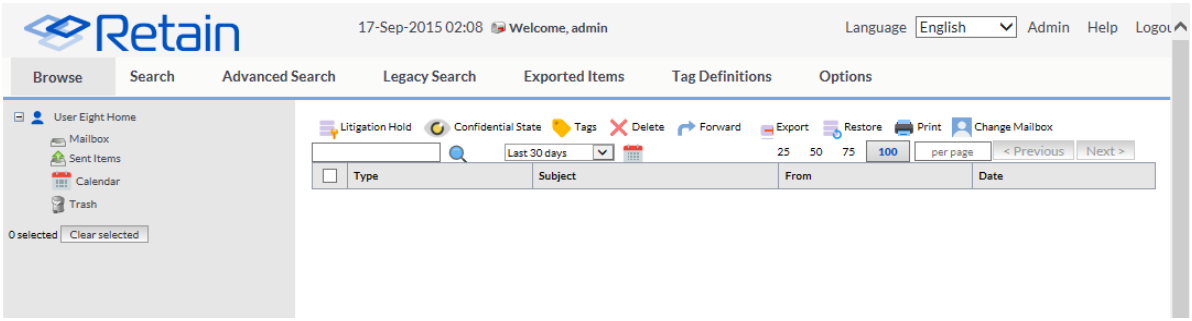
After selecting a mailbox, click ‘Ok’ to load that mailbox into the viewer.

How the browse interface appears will depend on the email system being archived.

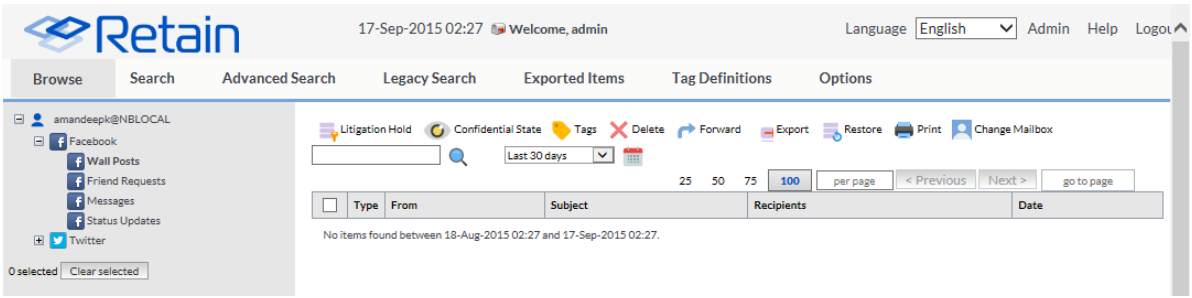
Exchange



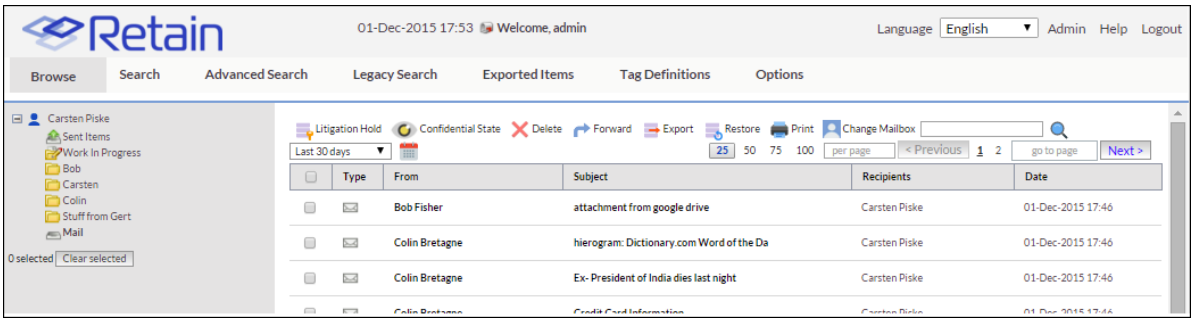
GroupWise



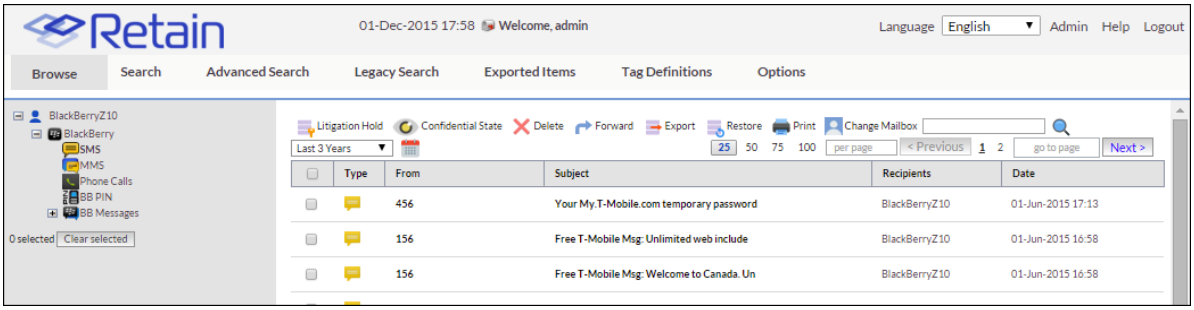
Social Messaging



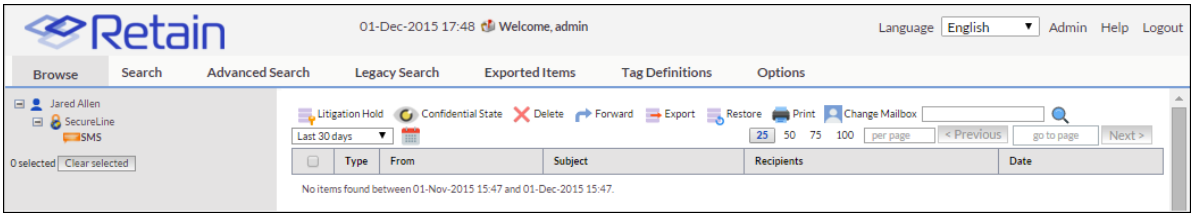
Google Apps



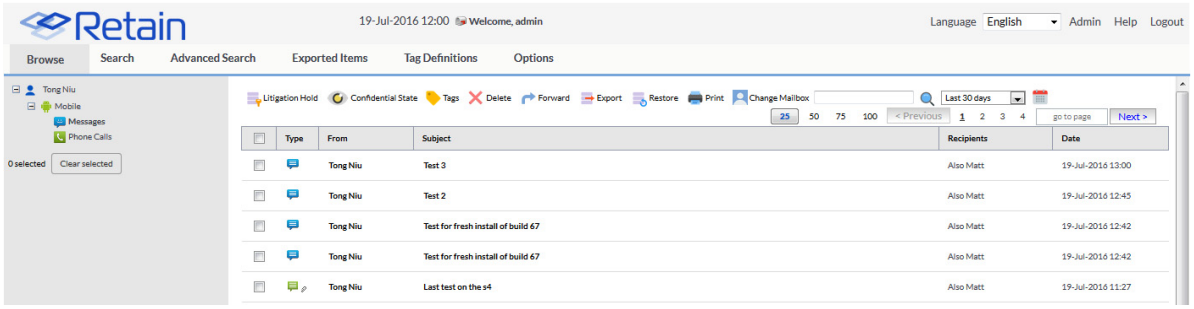
BlackBerry



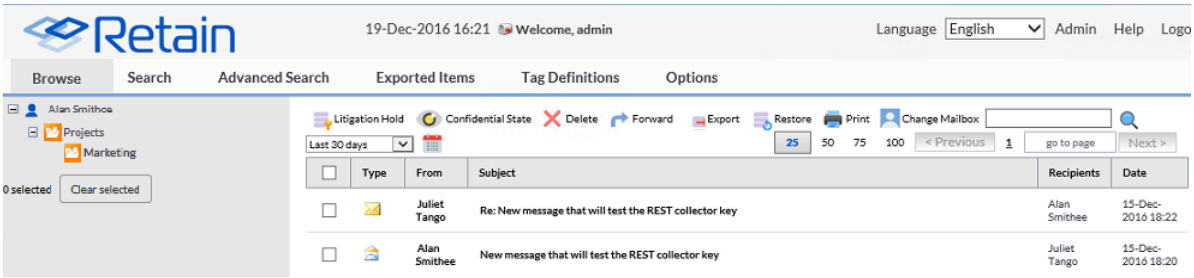
CellTrust



Mobile



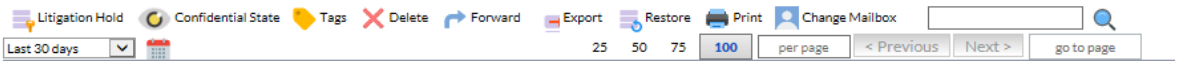
GBS Notes



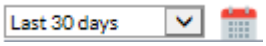
7.2.1 Toolbar

The toolbar has many different tools. Which tools are visible to each user will depend on what user rights they have been granted.

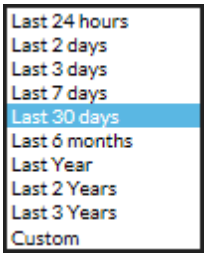
You will first have to select one or more items with the checkbox before these will function.



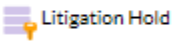
Date range



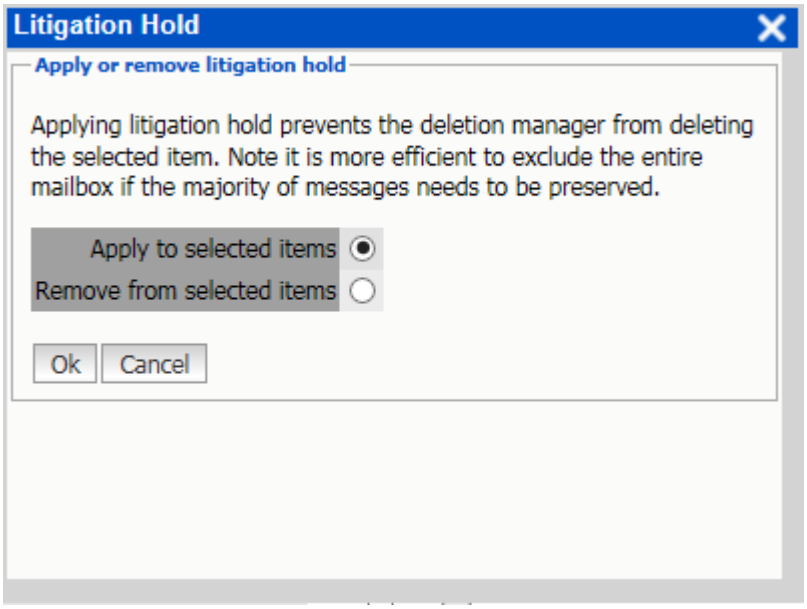
Pay special attention to the date range tool. Only messages which fall into this specified time frame will be displayed. The date range can be any of the drop-down options or it can be a custom range as selected through the ‘custom’ calendar button.



Litigation hold



The Litigation hold button applies the litigation hold property to messages and items. Items with a litiga-tion hold cannot be removed from the system. To apply or remove the litigation hold from items in the archive, a user must have the litigation hold right.



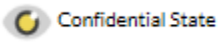
To apply or remove a litigation hold, first select the desired item or items, then click the ‘Litigation Hold’ button to open the litigation hold window.

Select whether to apply or remove the hold. Select ‘Ok’ to continue.

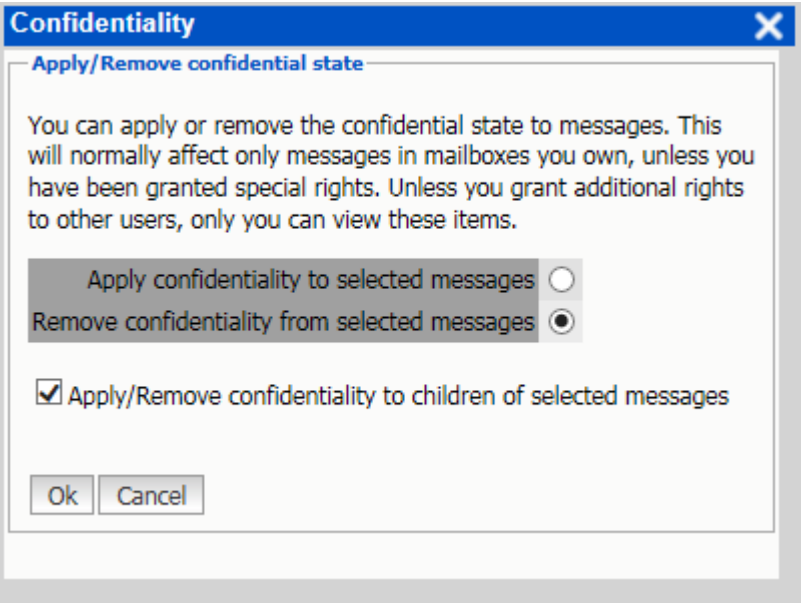
Items which have a hold placed on them display a miniature hold icon next to the message icon.



Confidential State



Toggles the confidential state. Confidential items are only visible with Admin rights, or to those specified by the user under the *Options* | *Confidential Exceptions* tab. Items which a user has marked in their box as Confidential may be viewed by other users who have the 'See Confidential Items [other mailboxes]' admin right.



To apply a confidential state to an item, select the item then select the 'Confidential State' button.

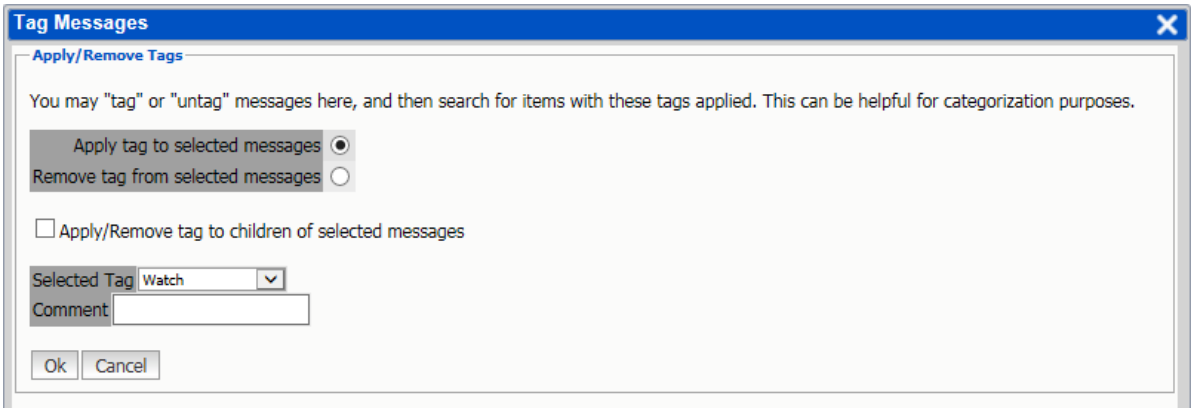
Items which have been marked confidential display the confidential icon below the message icon.



Tags



The tags icon allows the application or removal of tags from items. This is only visible with appropriate rights. Tags are an informative note which can be attached to any data item in the search messages interface. To apply a tag, select an item and click the 'Tags' icon.



Select the tag which is to be applied. All tags open to use by the user will be displayed. There is no limit to the amount of tags an item may have. An additional comment may be added to the tag for the specific item or items selected. Additionally, any associated items or messages, ‘children’, such as all messages in a thread, may also be selected for tag creation. Tags must be created previous to applying them to mes-sages. See the Tag Definitions tab for more info.

Messages with a tag applied to them, either a global or personal, will show a miniature tag icon next to the message or item icon.

Delete



The delete option allows users to mark items for removal from personal archives. This right is only visible with rights. To delete an item, highlight the item and then select the ‘Delete’ button.

The Delete right is not a default right, it must be specifically added to users or groups. Messages or data marked to be deleted will not necessarily be immediately removed, but will be queued for deletion. The actual item removal will depend on server availability and the amount of items queued for deletion.

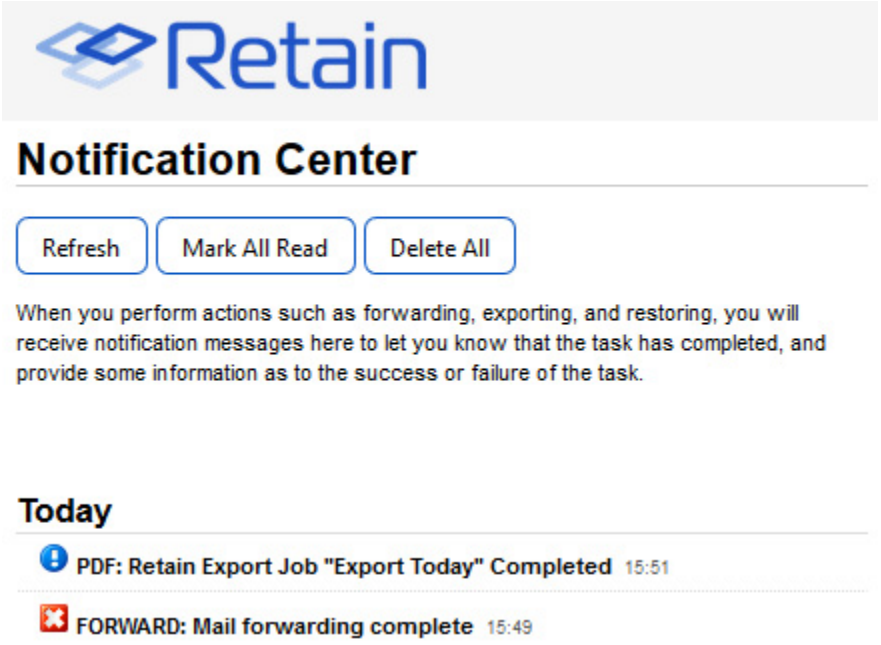
Forward



The Forward option allows users to forward messages out of their archive to a specified address. This right is only visible to users who have the Forward Messages right. The STMP server must be setup by the admin user under Server Configuration | Notification for this to function successfully.

Select a message or messages and then the ‘Forward’ button. Specify the address or comma-separated addresses the messages should be forwarded to, a desired subject, and any comments. Finally, select the ‘Forward Messages’ button and Retain will send the message via SMTP. This is a default user right.

There will also be a notification entered in your Notification Center found under "Welcome, [username]" at the top center of the Retain Mailbox web console.



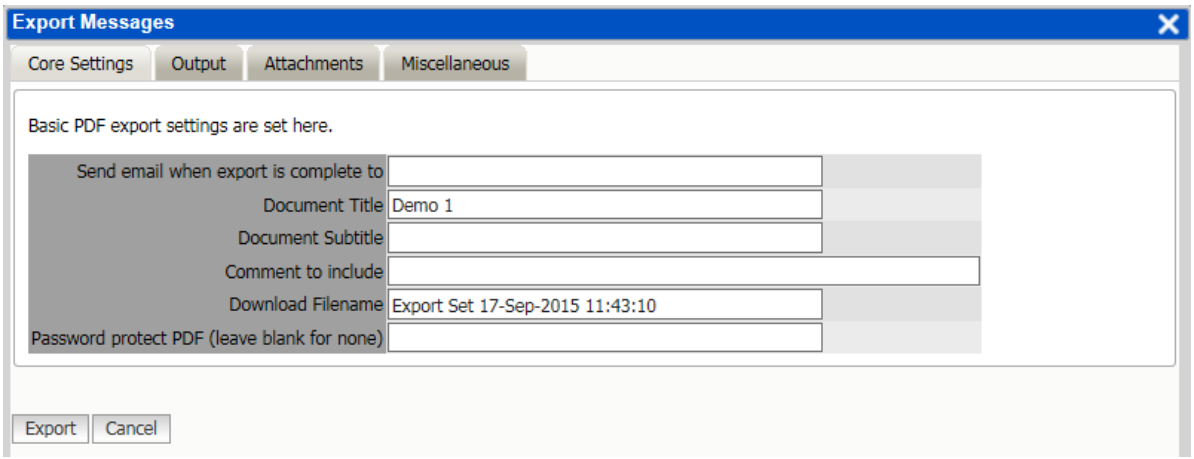
Export



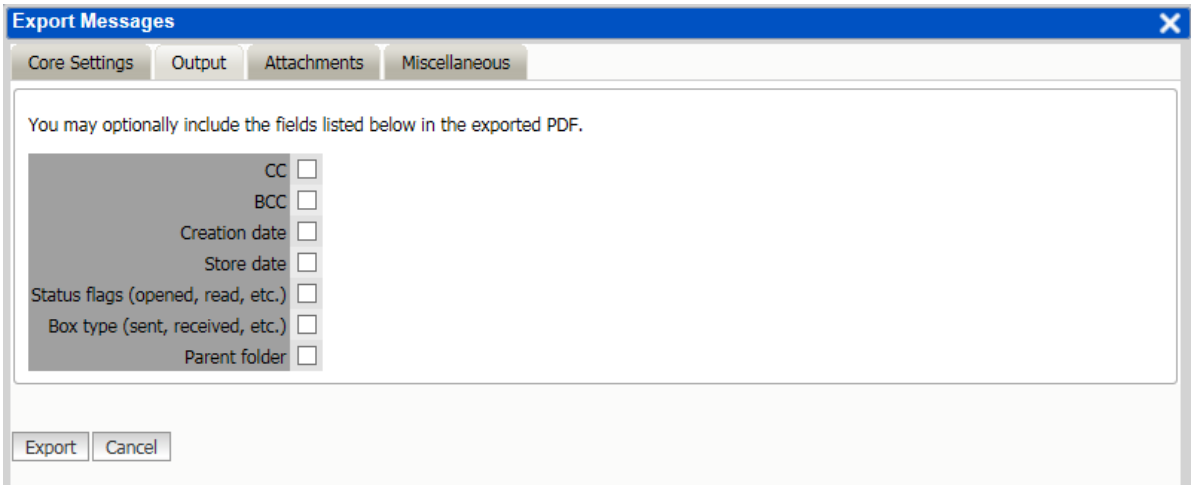
The Export option opens a PDF export window to create a PDF of the selected messages and emails a notification to the specified email address. Messages will be displayed as navigable files in an Adobe PDF Portfolio. The default view provides a list on top of the PDF viewer which works as a table of contents to provide access to the exported archive and requires Adobe Flash to be installed on the viewing workstation.

Core Settings tab

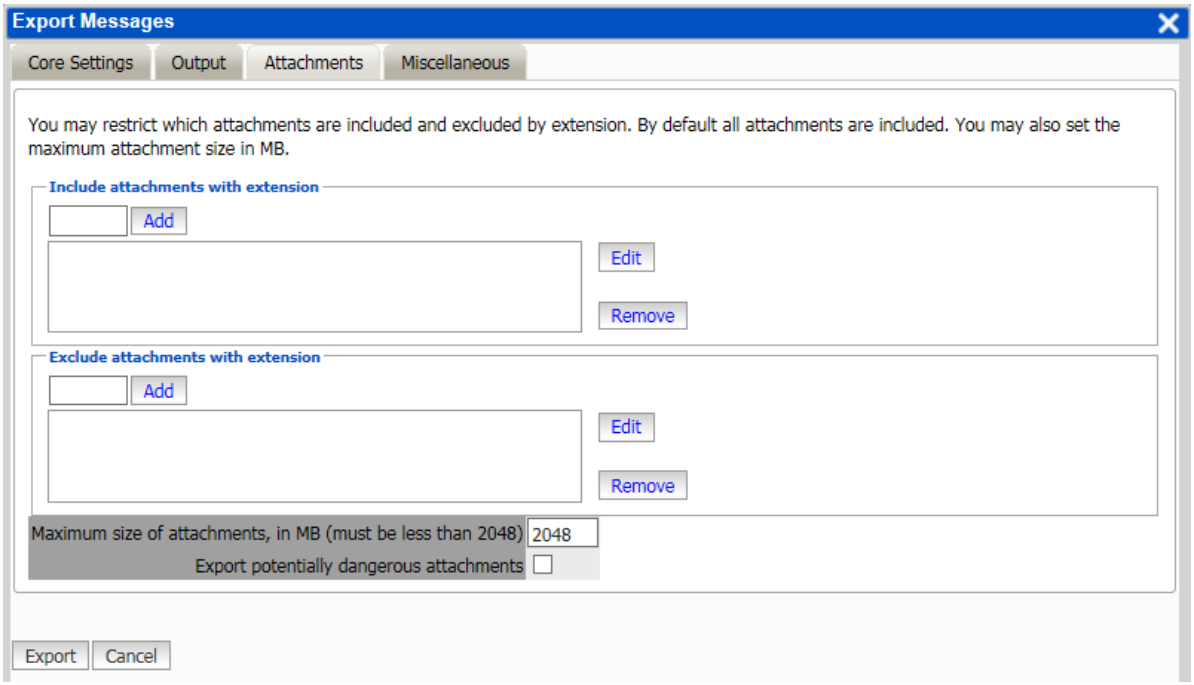
Select the desired message or messages and then the ‘Export’ button.



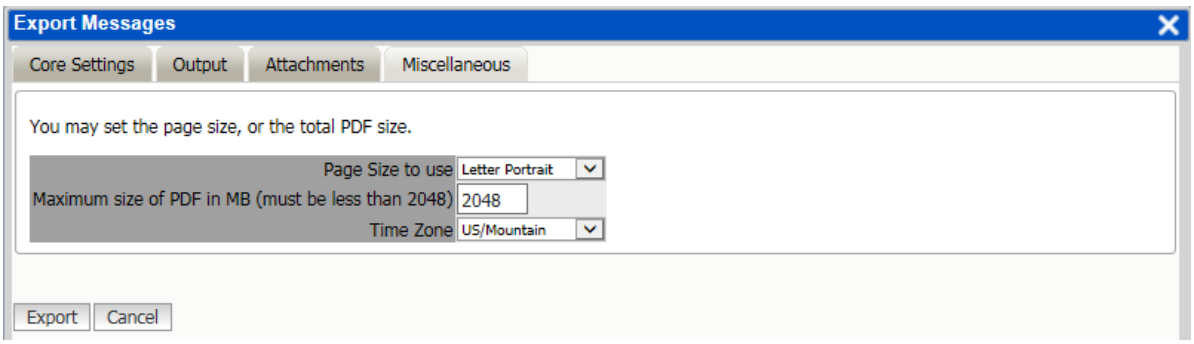
Fill out the various options as desired.



File output and included fields can be edited



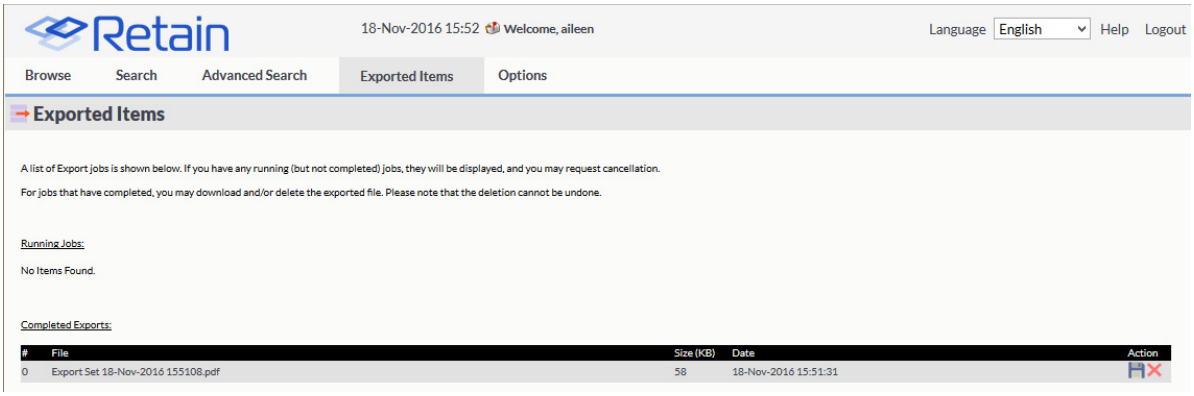
Select which attachments to include and exclude from the PDF.



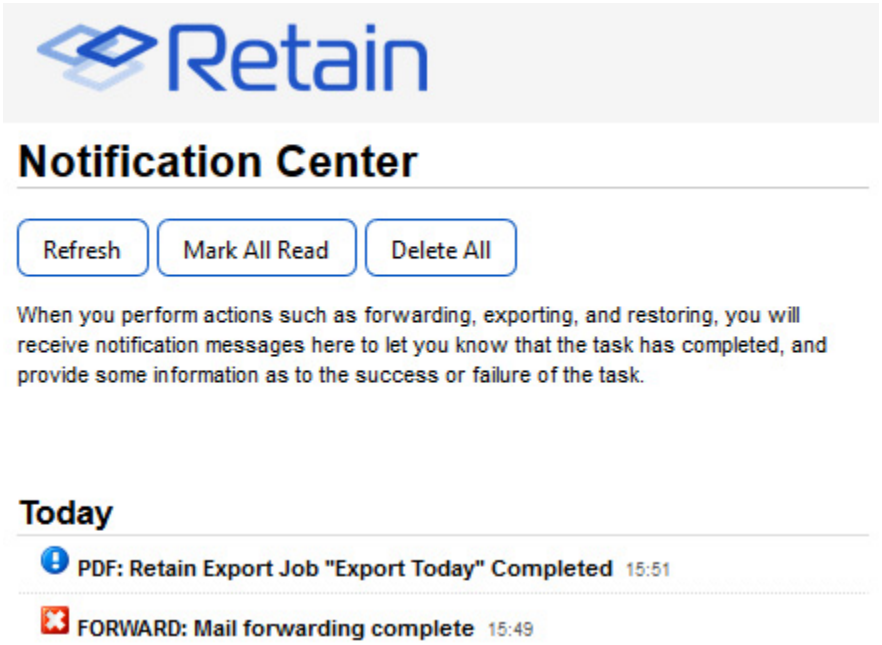
Additionally, any restrictions to the size or time zone for the PDF. If the archive size does not exceed the maximum size, then the archive will be contained in one file. If the archive is larger than the defined limit, it will be exported in multiple PDFs. Time zone information will be used for formatting any time and date information in the email headers. The time written to the cover sheet and date column in the PDF table of contents is written according to the time zone of the Retain Server.

When ready to export, select ‘Export’. Depending on the size of the selected archive, the export may take some time.

Currently running export jobs are listed under the ‘Exported Items’ tab.



Once complete, a notification email is sent, if a notification address was provided, and the PDF is available under the ‘Exported Items’ tab. Locate the desired export and select the disk icon to download the fin-ished PDF. There will also be a notification entered in your Notification Center found under "Welcome, [username]" at the top center of the Retain Mailbox web console.



Restore



This option restores messages to the message system. Forwarding is the preferred method to return mail to the email system. To use, select the desired messages and select the ‘restore’ button. Once there, select the target message system and mailbox. (If the user has rights to restore to any mailbox, the target mailbox must be selected. Select the ‘Restore Message’ button and messages will be restored to that account.

Restoring a message which currently exists in the target mailbox results in a duplicate message existing in that mailbox. It is recommended to only restore messages that have been deleted from the mail system, but exist in the Retain archive.

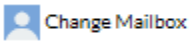
Restore works best with On-Premise Exchange, restoring the message completely to the Exchange Inbox, which is seen as a new message by the system. With GroupWise, it restores a link to the message, if stubbing server is configured in Retain but stubbing does not need to be enabled in GroupWise and we do not recommend enabling Stubbing in GroupWise. Restore does not work with GMail.

Print



This option allows users the ability to print selected messages. To use this option, select the desired mes-sages and then select the ‘print’ button. The ‘print’ function utilizes the user’s local print configuration; if there is no printer accessible to the local workstation, print will not function. This is a default user right.

Change Mailbox



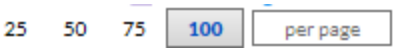
This option allows users who have the administrator right to see all mailboxes, or who have proxy rights to another mailbox, the ability to switch mailboxes.

Quick Search



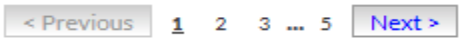
The Quick Search is a way to rapidly search for data in the subject field of messages in the selected mail-box.

Results per page selector



Results are automatically set to 100 per page. The selector allows the user to specify how many messages are to be displayed on each page. There are predefined values, or the user may put their own dictated limit in the field provided.

Previous/Next



The Previous/Next option allows users to navigate through the search or browse result pages. The total number of pages is displayed next to the ‘Next’ button.

Select a number to immediately move to the desired page

Go to Page



The Go To Page option allows users to specify what page they want to jump to.

Select the option box, specify the page number desired, and hit ‘enter’ to immediately go to that page.

Quick View



The Search interface incorporates a quick view option. This option allows users to quickly view the body of a message without having to actually open the message in the interface. To access the quick view, click on the message or data item icon next to the check box.



Once selected, the quick view window opens to quickly display the message contents.

The Quick View is a display only function and there are no actions available from this interface. When complete, close the window by clicking on the ‘X’ at the top right corner of the window.

Selection and sorting



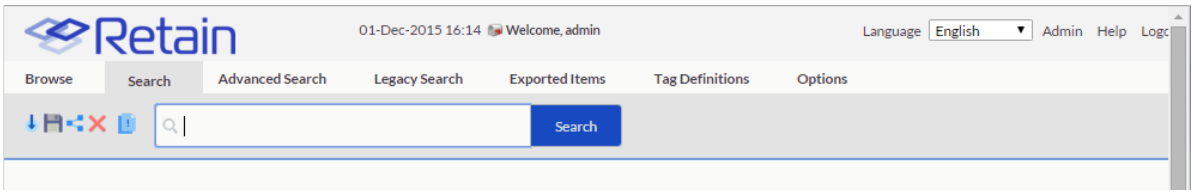
The top column bar contains all the different column headers and is located directly beneath the toolbar. This tool contains a select all/none checkbox as well as sortable column headers. Every column in this bar may be sorted by order except the ‘Mailbox’ and ‘Folder’ columns. (They are multi-value attributes and thus are not sortable.)

Browse Tab Defaults

By default, the browse window remembers which mailbox last browsed was, and will return to that mail-box when opened. When Administrators first login, they will be taken immediately to their mailbox. Afterwards, the browse will remember their last location.

7.3 Search

The search in Retain 4 is designed to work much like a Google search, with anticipated auto-correct search terms. The Search tab will allow searching across the entire Archive for administrators, while users will only be able to search their own archive.



The Search tab allows for management of searches. This is accomplished by Saving searches for later use. These searches can then be shared with others, loaded later, removed when no longer desired, and trimmed to remove duplicate data.



The icons for these functions are on the right, next to the search field. In order: *Load, Save, Share, Delete,* and *Collapse Duplicates*. Collapse duplicates removes all duplicate results but only functions with Group-Wise mailboxes.

To begin a search, type the desired search term into the search query field, or load a saved search. The search results will be automatically populated on specified search terms after a sufficient pause in typing. The auto complete suggestion feature requires at least 3 characters to be specified and a pause of several seconds before it will begin to work.

Tokenized Search Phrase

The indexing engine follows Unicode Standard Annex #29. This standard uses many common characters such as ' " . : @ + - * / and , as phrase ending or beginning characters. These characters will cause the system to read any character separated terms as individually entered items when the search is performed. Individually entered terms are treated as OR searches. For example, this means that using the search term 10/20 will be processed as 10 OR 20. Substituting a space instead of the

character will provide a logical 'AND' search. For example, searching for 10 20 will be processed as 10 AND 20 as the search term. For example, searching for an email address test@gwava.com will drop the @ even in an exact search with double quotes and with return test AND gwava.com.

Spaces between normal search terms are a delimiter between search terms:

The search algorithm follows these criteria:

1. Highest weighting: All delimited words together consecutively. (e.g. Searching for "quick brown fox" while the message contains: "The quick brown fox jumped over the lazy dog.")
2. All delimited words in the string although they may not be right next to each other. (e.g. quick AND brown AND fox)
3. Least weighting: Any of the delimited words in the string. (e.g. quick OR brown OR fox)

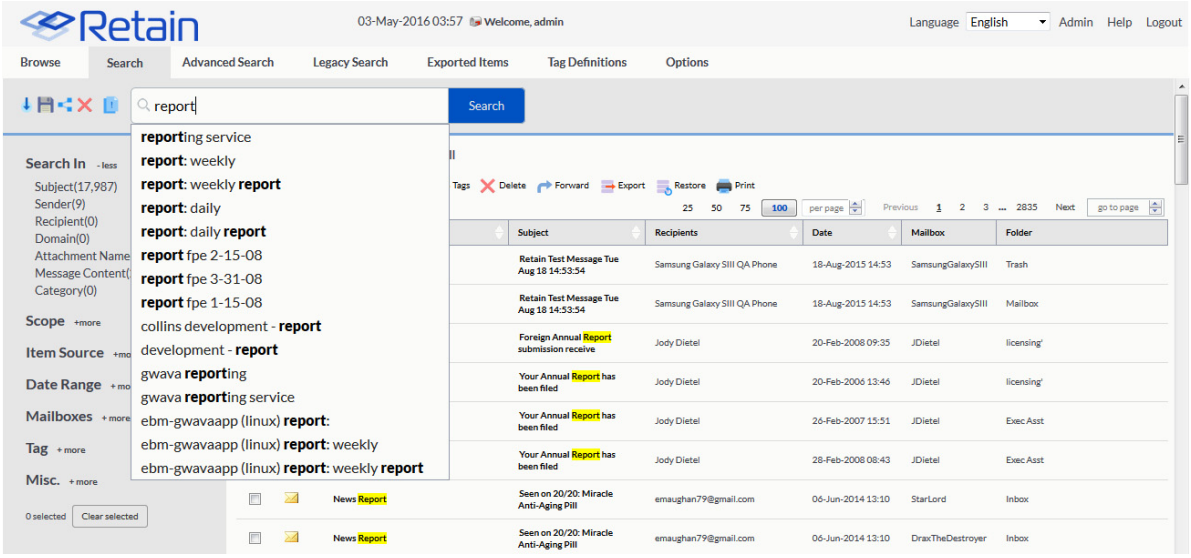
Results from these different terms are weighted in the order listed above. This behavior can be overridden by locating the solarcloud.indexing.properties file and changing the property phraseSearch.singleWordMatch. Setting this property to '0' turns this off, while leaving it set at '1' activates this search behavior.

Result terms are also highlighted in the user interface to ease visual confirmation that the term exists in the message. However, terms in attachments are not highlighted.

7.3.1 Wildcard Search

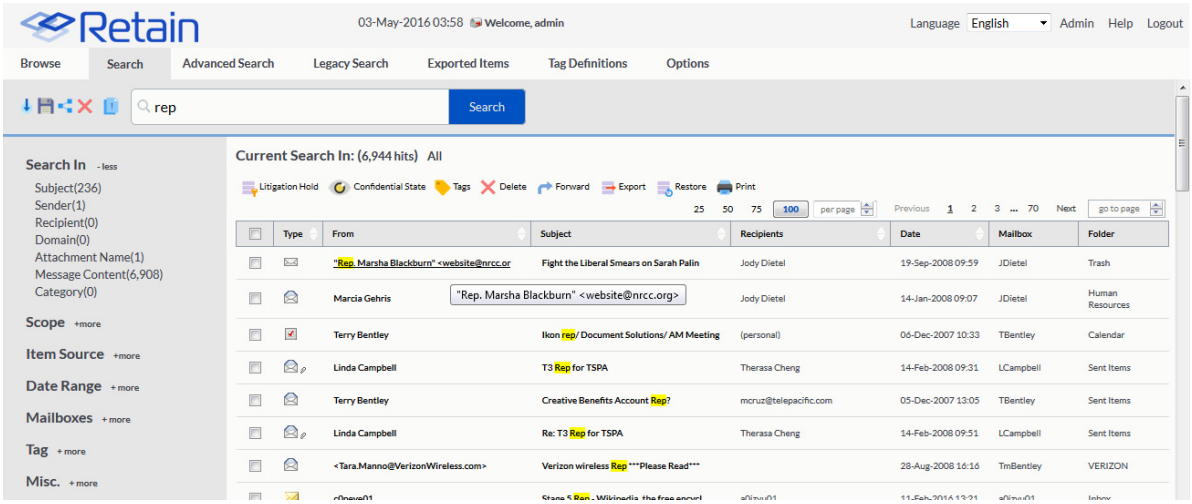
Accepted wildcard characters are: *, ?, and "". The '*' denotes 'any character or characters', while the '?' denotes 'any one unknown character'. The double quotes("") denote an exact phrase, and only that exact phrase. For instance, a search for the term "We all love spoons" will fetch a result of that exact phrase. Asterisks may be added inside the double quotes to allow for incomplete or unknown words: "*e all love sp*" will return both 'We all love spoons for our ice cream', and 'We don't like sporks but we all love space.'. While the '?' works like the asterisk, it is only for a single unknown character. This is particularly helpful when searching for exact phrases where the terms may be misspelled. For example, "Their going" would miss a misspelled 'there going'. However, if the search term were "The?? going", it would catch both.

Supported Regex characters are listed in the Advanced Search section.



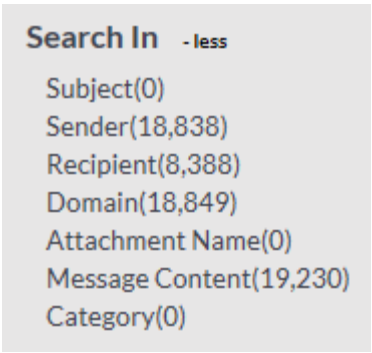
Select a final query, either by hitting enter, clicking with the mouse, or using arrows and hitting enter.

Search results will display the results with the search term highlighted for each message. The message type, Sender, Subject, Recipients, date, mailbox, and folder are all displayed.



7.3.2 Search In

Once search results have begun to populate, the left hand scope pane is populated with limiting and fil-tering options. If faceting is enabled the side bar will show numbers next to each section, indicating how many hits there are for each particular topic. The hits are total numbers of matching instances, not items. So if a message states a search term several times in the message body, it will be counted as that many hits even though it is only one message.



The Search In criteria limits the area in the message or data where the search is performed:

- The Subject indicates hits in the Subject field.
- The Sender field contains the sender of the message or data item.
- The Recipient is the recipient of the item or message.
- The Message Content will search in the following locations within a message:
  - - body
  - - attachment
  - - subject
  - - headers

Searching exclusively for the domain will be effective with search terms if a complete domain is provided, otherwise the term is recognized as text. If the top level domain is not known, (.com, .org, .edu, .etc) then the search term should use an asterisk afterwards. For example, searching for gwava.com will yield good results, as will searching for gwava.* or gwava*, though results will vary.

Searching for an item that has BCCed recipients has some special behavior. Retain acts like a normal email system, only the sender and the BCCed recipient can see themselves. If you are logged in as the sender, you will be able to see To, CC and BCC recipients as normal. If you are logged in as the To, or CC recipients, you will only see the Sender, To and CCed recipients. If you are logged in as a BCC recipient, you will see the Sender, To, CC, and no other BCC recipients. If you are logged in as admin and in a mailbox other then the Sender, you will not see the BCC recipients listed in the item even though they are returned from the correct mailboxes, except from the Sender folder, because admin is not one of the BCC recipients.

7.3.3 Item Type

The Item Type criteria option limits the type of message which is to be searched. All item types are available in the scope term. Again, the number of available hits is displayed to the side.

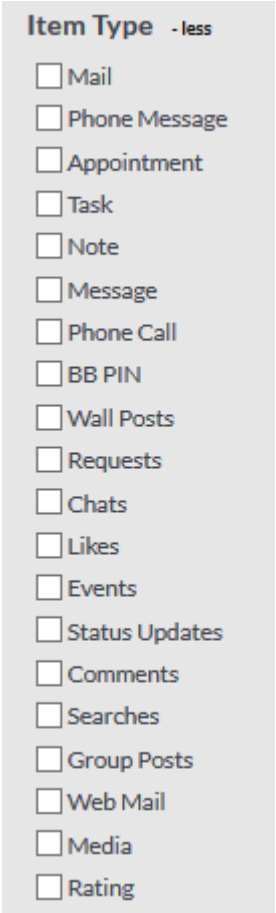
The Message Item Type includes:

Mobile: SMS & MMS

BlackBerry: BBSMS, BBMMS, BBMsg

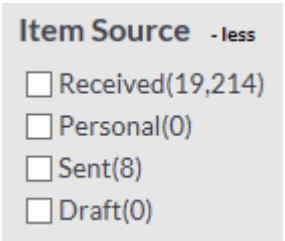
CellTrust: CTSMS

Bloomberg: BLMBMessage



7.3.4 Item Source

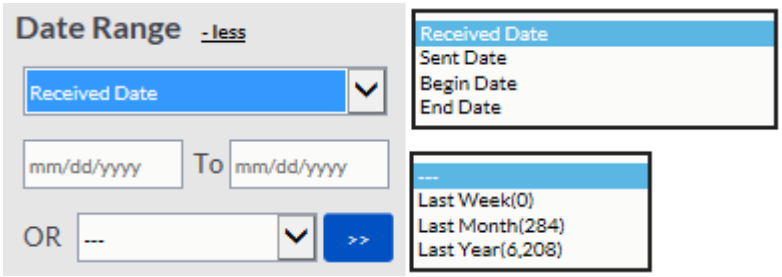
The Item Source criteria option limits the results to a particular source. This source can be limited to show ‘received’, ‘personal’, ‘sent’, or ‘draft’ items.



7.3.5 Date Range

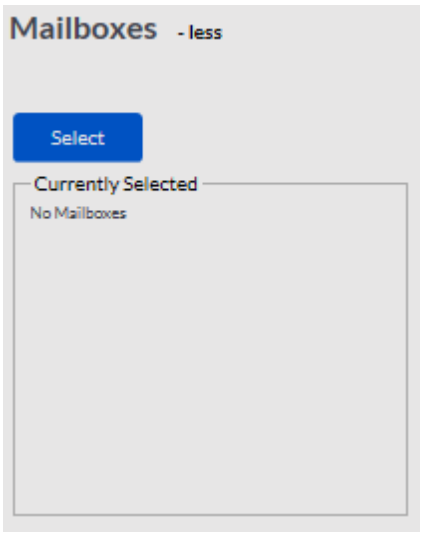
The Date Range criteria limits the time frame of item’s creation. Only messages which conform to the date range selected will be displayed in the results field.

The date range may be specified for any of the dates corresponding with an item. The range may also be selected form the drop-down menu, with pre-configured times for last week, month, or year. These time frames are for the past 7 days, the past 30 days, or the past year, not the previous calendar time frame.

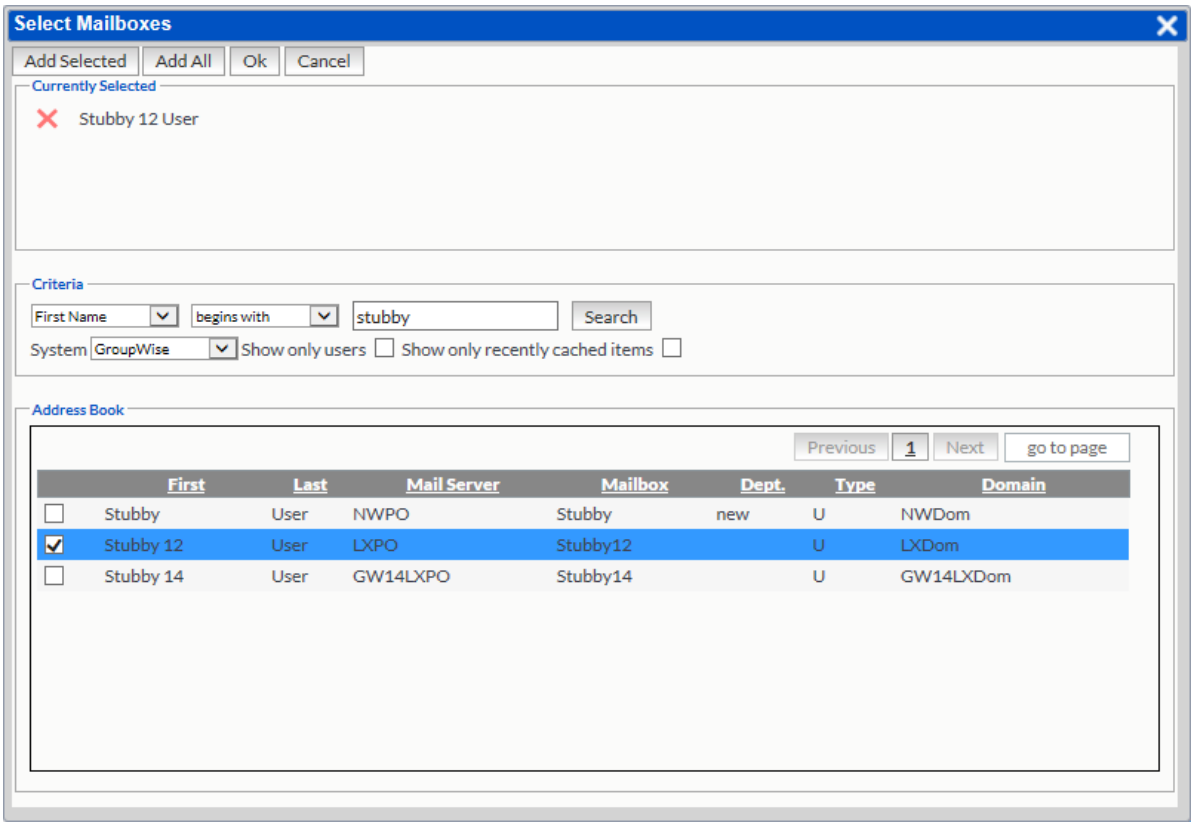


7.3.6 Mailboxes

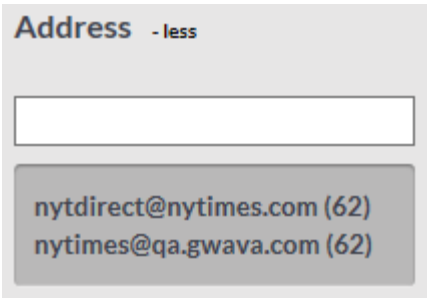
The Mailboxes criteria limit which mailbox or mailboxes the search will pull results from. To add users to the selected mailboxes list, click on the ‘Select’ button to launch the mailbox selection window.



The Select Mailboxes window allows for searching of every mailbox available to the user, the admin user can search all mailboxes. Mailboxes must be searched for by system and specified criteria. The results of the search are displayed below, while the active selected mailboxes can be added to the dialog through the use of the ‘Add Selected’ button along the top. Alternately, if the ‘Add All’ button is clicked, it will add all mailboxes dis-played in the search results, the Address Book field. Addresses which have been added to the top field may be removed by selecting the red ‘X’.

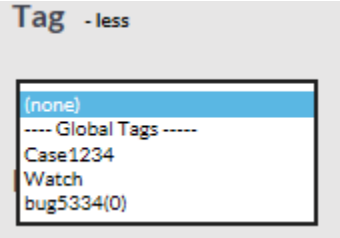


Once the desired addresses have been added to the top window, select the ‘OK’ button to load them into the search pane.



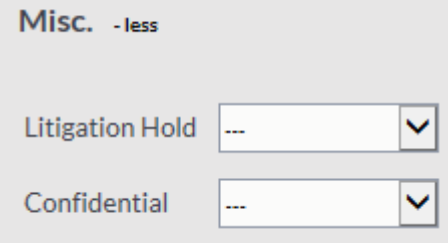
The Address option limits the results to a selected address. The addresses available are displayed below, and may be selected Addresses in the window are dictated by what is in the result set. Selecting an address adds it as a filter to the top of the search window, or the user may specify an address manually. Multiple addresses may be added to the search window at a time. To remove an address filter, select the ‘X’ next to the active address, and the result set will be reset.

7.3.7 Tag



The Tag option limits the search to items which have been tagged with a specific tag definition. The tag definitions may be personal or global. Tags must be specified in advance and applied before this option will work. Select the desired tag to limit the result set.

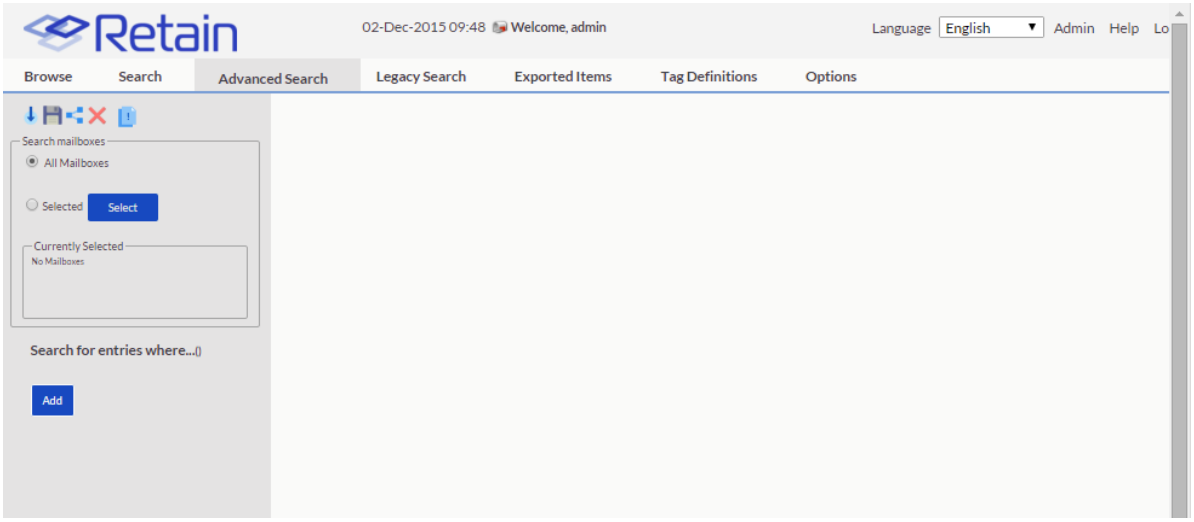
7.3.8 Misc.



The Misc option contains limits for the Litigation hold and the Confidential tags. They have two settings: True or False. A Setting of ‘True’ restricts all results to only items which have the selected tag.

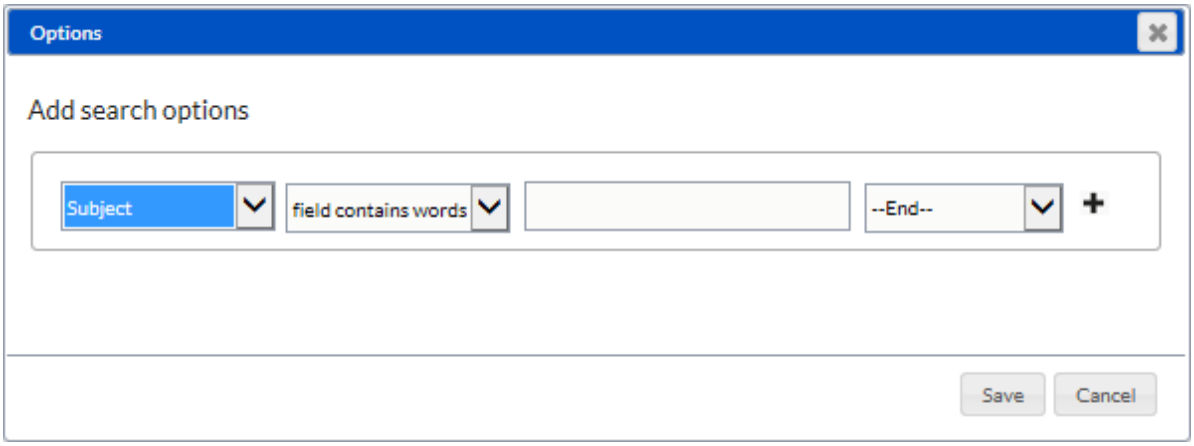
7.4 Advanced Search

The Advanced Search tab contains the ability to specify vast amounts of criteria and combine search terms to exclude and include various searchable items to retrieve specific information. This search works better, the more you know about what you are looking for, as it allows fine tuning of criteria.



Searches will be restricted to the mailboxes specified in the ‘Search Mailboxes’ window. By default, all mailboxes are set for searching. To limit the search to the mailbox or mailboxes specified, click on the ‘Select’ button to open the ‘Select Mailboxes’ window. The ‘Select Mailboxes’ window functions exactly the same in advanced search as it does in the standard search.

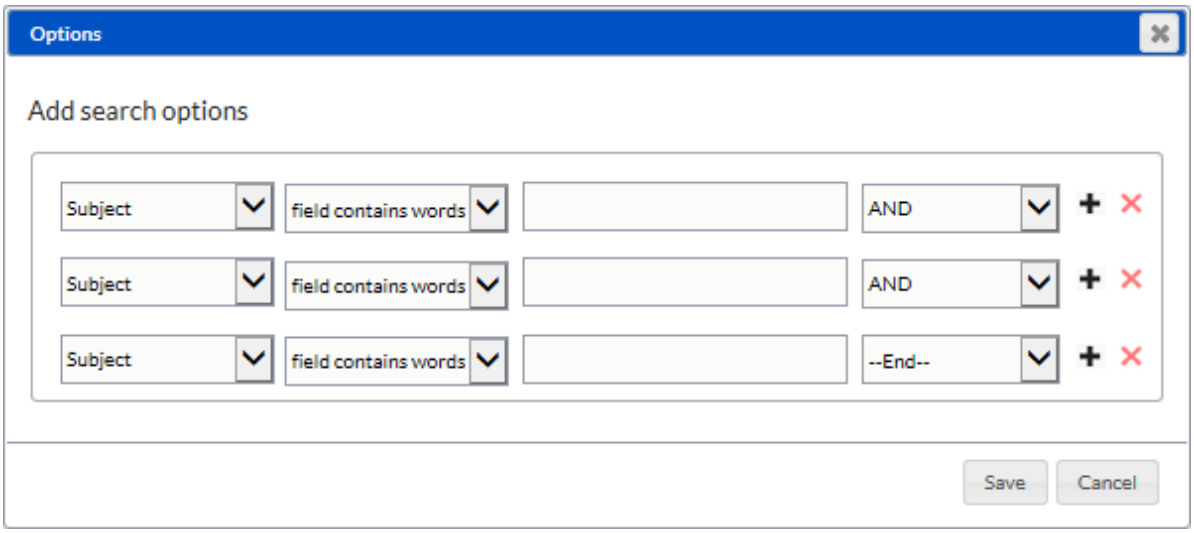
The system will begin to display results as soon as either the search mailboxes have been specified, or new search criteria has been added. To add new criteria, select the ‘Add’ button.



The Search Criteria contains the ability to specify where to search, operating criteria - (word ends with, word starts with, field contains words, field contains phrase), and the desired search terms. The list of search items and fields available to be specified in the drop-down list is shown. Each variable on the list is tied to appropriate search operators, (date range allows the specification of a date, Confidential tags have a true/false oper-ator, etc.)

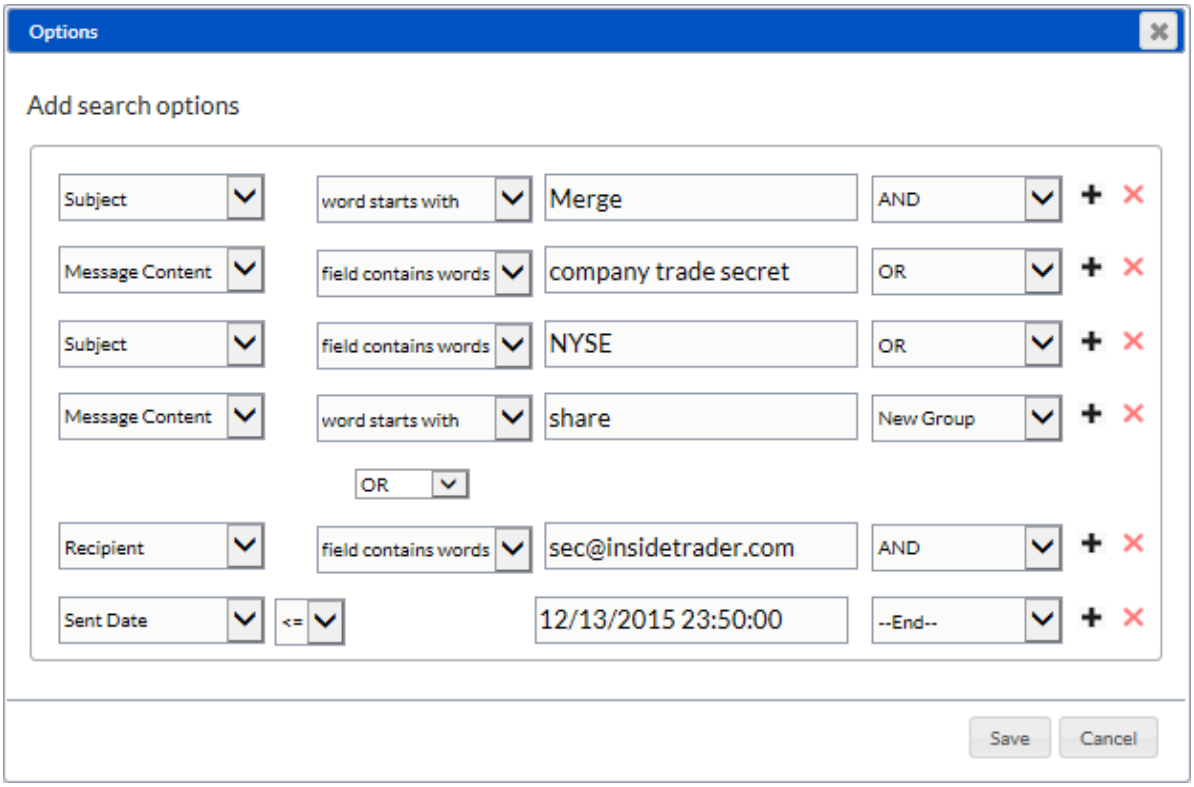


In addition, the interface allows for no limit of search terms. additional terms may be added to he search criteria and connected to the previous search terms. Additional criteria may be logically connected with ‘and’, ‘or’, or ‘new group’. To add a new search term and criteria, select the ‘+’ directly to the right of the existing search criteria.

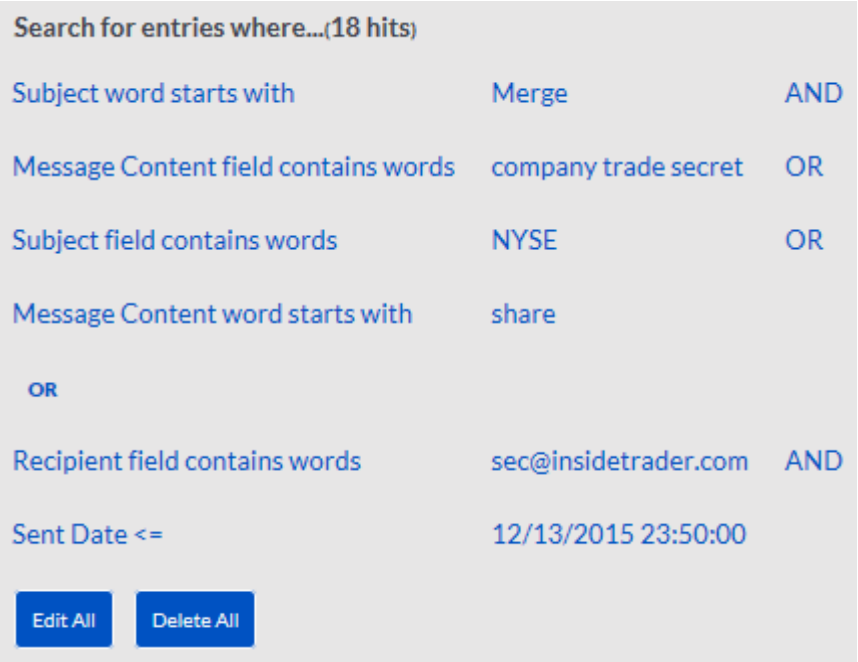


By default, when a new search term is added, it is automatically ‘AND-ed’ together with the previous search term. This allows you to be able to build complex search terms to fit known data.

When building complex search criteria, it is critical to know what you are looking for. For instance, if an insider trading tip was suspected, and the recipient was known as well as some details about the message and when it must have been sent by, the following search could be compiled:



In this search, any message sent which stated ‘merge’, or ‘merger’ in the subject, and contained a known company secret in the message body, or, discussed the name of an executive involved, would be dis-played. In addition, the search would also grab any messages sent to the suspected contact before the merger date. Additional criteria which could be added includes the company’s stock listing or any further details pertaining to the proposed leak.



To begin the search, select the ‘Save’ button at the bottom of the query window to perform the search. The active criteria is now listed in the left pane, and may be edited or removed. To add criteria, select the ‘Edit All’ button to add to or refine the search criteria.

7.4.1 RegEx and Wildcards

Both the Search and Advanced Search contain limited support for Regular Expression searches. To use Regular Expressions, simply put the desired regex string into the criteria window, denoted by a ‘/’ on either side of the regex. if the ‘/’ is not used, Search will not recognize it as regex.

The Supported Regex characters are:

Syntax	Description	Example
?	Zero or one occurrence	/abc?d/ matches abd, abcd
*	zero or more occurrences	/abc*d/ matches abd, abcd, abccd
+	one or more occurances	/abc+d/ matches abcd, abccd
{n}	n occurrances	/abc{2}/ matches abcc
{n,}	n or more occurrances	/abc{2,}/ matches abccc,
{n,m}	n to m occurrences, including both	/abc{2,3}/ matches abcc, abccc
[characters]	character class	/[abcde]{2}/ matches ab, ac, be
[^characters}	negated character class	/[^abcde]{2}/ matches fg, hi, jk
-	character range, including end-points	/[a-z]{2}/ matches ab, ac, bz
.	any single character	/ab.k/ matches abck, abdk, abek
|	union	/(ab | cd)/ matches ab, cd
&	intersection	/([a-z}*|cd+)/ matches cd, cdd, cddd

Syntax	Description	Example
@	any string	/@/ matches abc is abc
<n-m>	numerical interval	/<99-102> some dr/ matches 99 some dr, 101 some dr

Accepted wildcard search symbols are: *, “”, and ?. The ‘*’ denotes ‘any character or characters’, while the ‘?’ denotes ‘any one unknown character’.

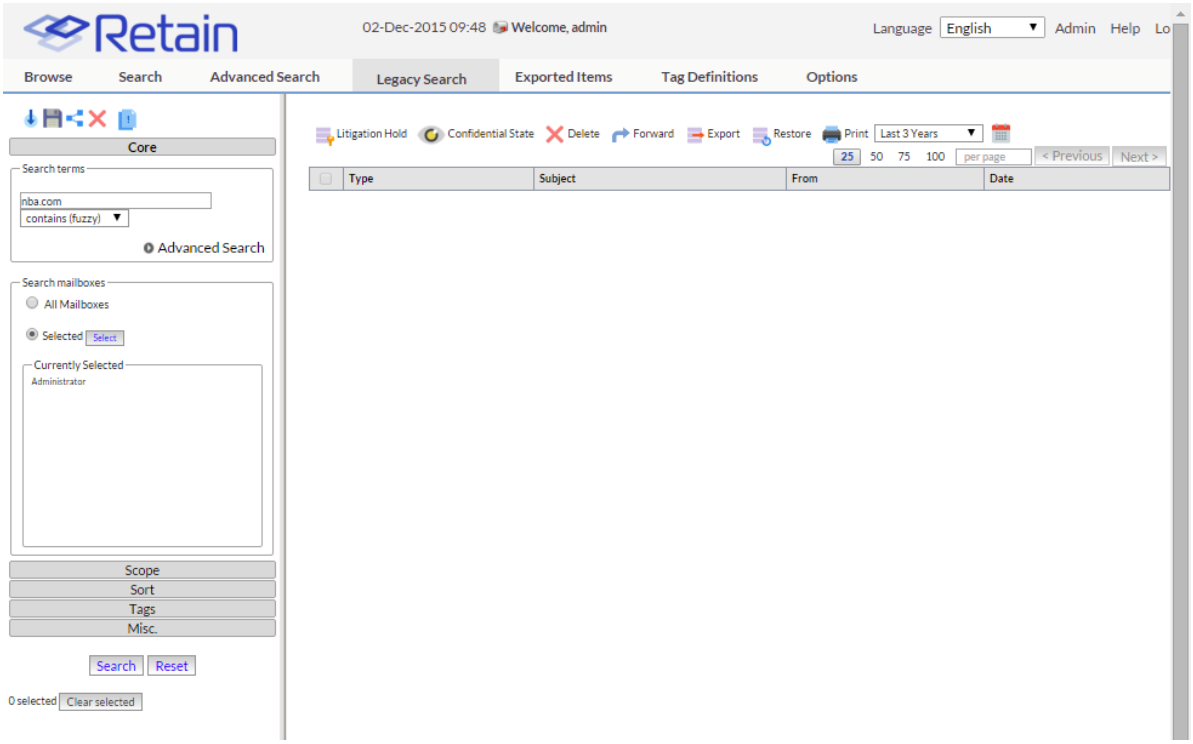
Special Characters

The Search has a list of special characters which cannot be searched for, and will cause erratic results with search criteria. The list of non-supported characters is: @,+,-,|,[],{},(),”,\,#,&,~. All of these characters are viewed as delimiters, and will break up the query. They are not supported.

As a result, if a special character is to be searched for, it must be placed in double quotes. Ie. searching for bob@gwava.com will result in any user at the ‘gwava.com’ domain. To search for bob@gwava.com, the exact phrase must be specified; ie. “bob@gwava.com”

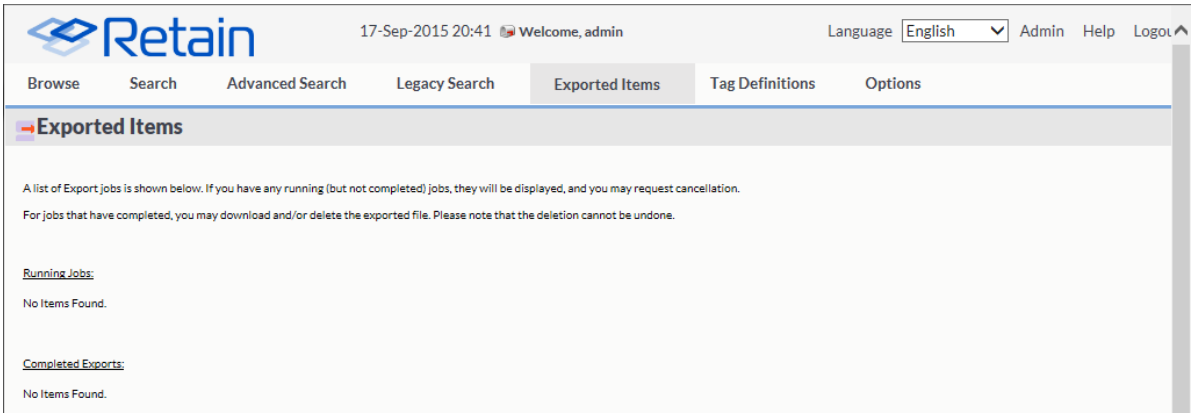
7.5 Legacy Search

The Legacy Search tab provides access to the older search engine, and is included as a courtesy for those who wish to continue using it. While the UI has had an update, the functions and search process remains the same.



7.6 Exported Items

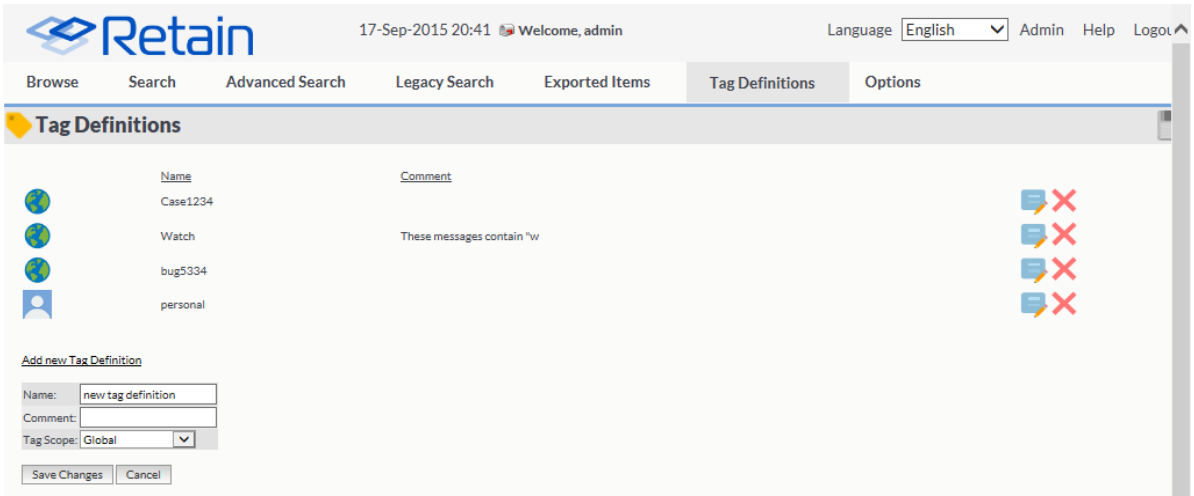
The Exported Items tab shows the export jobs which are currently running, or have been run in the past, on the system. If an export job has completed, this tab will contain the file and provides a link for down-load of the completed export job.



NOTE: The download is an Adobe PDF Portfolio which requires Adobe Flash to be installed to view items.

7.7 Tag Definitions

The Tag Definitions tab allows the creation and removal of Tags, their automatic comment, and name. Tags are an informative note which can be attached to any data item in the search messages interface. There is no limit to how many tags any one item may have applied to it, and there is no limit to how many tags a user may create. In addition, tags are also a serachable item, making this one of the most versatile ways to add long-term identification for items in the data store.



Before the tag icon will appear on the in the search interface tool bar, there must be at least one tag defined. To define a tag, enter the tag name and initial comment if desired, then, if the user has permissions to do so, define whether the tag is personal or global. Once saved the tag is available for use.

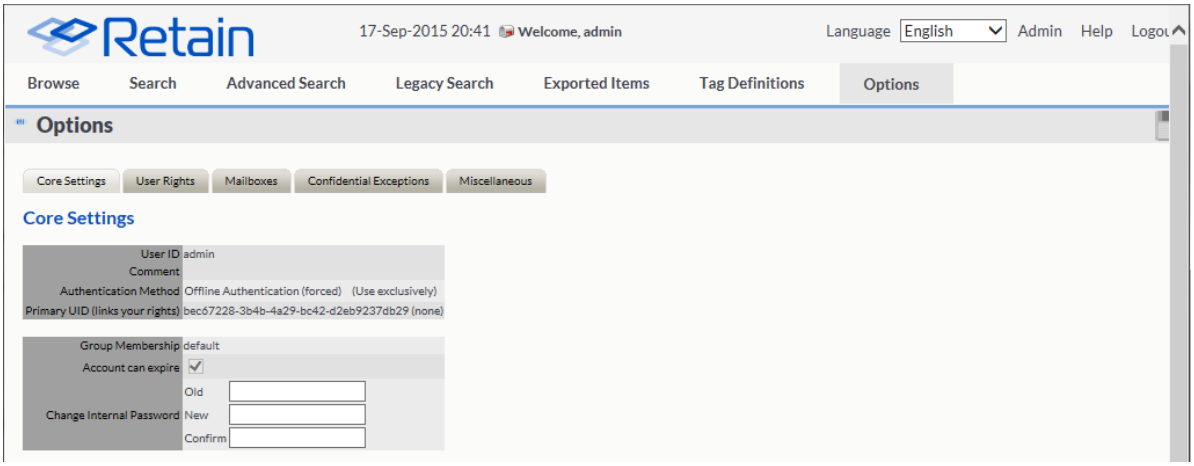
Global tags are tags that any user with the rights to see global tags will be able to view and apply. Personal tags are limited to the user who created them. Only tags visible to users will be available to be searched for by that user.

Any tags created or subject to manipulation by the user logged-in will be displayed under this tab.

To apply a tag to a message or data item in the search messages interface, simply select the data item or items, and then click the ‘Add / Remove’ tag button in the tool bar.

7.8 Options

The options section here is exactly like the section in the Administration | Users section. These settings here are specific to the currently logged in user. The current loaded user is shown next to “User ID” located below the Core Settings tab.



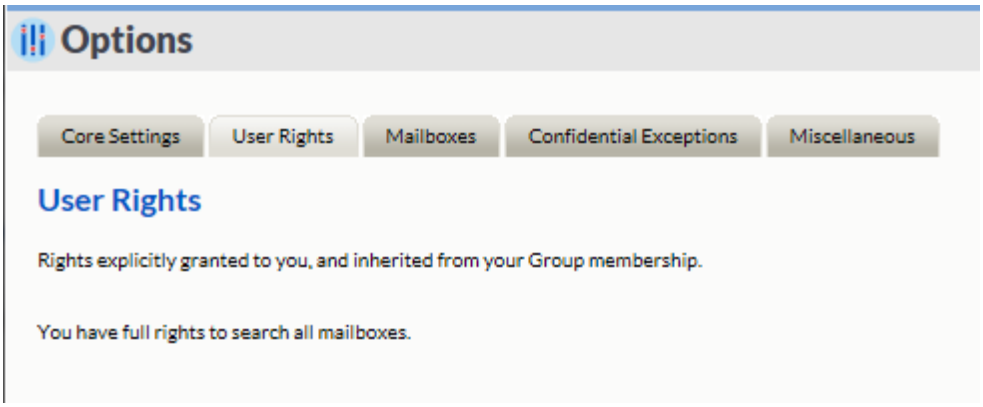
Core Settings

Among the information displayed is:

- Your User ID (internally generated by Retain)
- Your Group Membership (you cannot change that here)
- Your authentication method (you may change this if you have rights to)
- Your offline password if you don’t use SOAP to authenticate
- You may change the password only if you have rights to change it

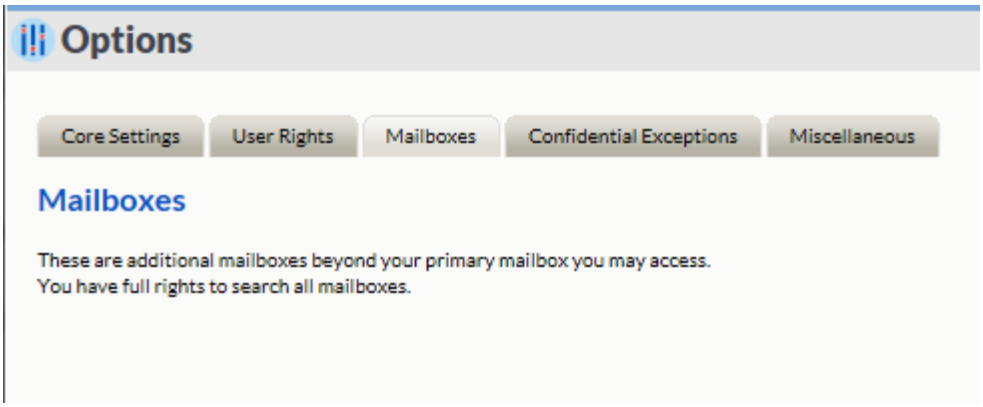
7.8.1 User Rights Tab

You’re shown the rights you have within Retain on this screen. Available rights may only be changed by a user administrator within the Administration screen.



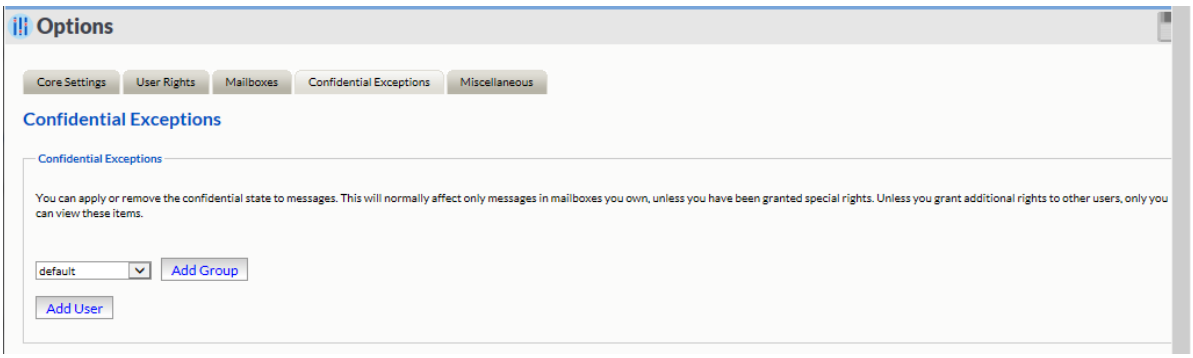
7.8.2 Mailboxes Tab

In this section, you're shown the mailboxes you have been given explicit rights to work in. By default, you have rights to only your own mailbox. If an administrator explicitly gives you rights to other mailboxes or you are a member of a group that has rights to other mailboxes explicitly, those mailboxes will be shown here.



7.8.3 Confidential Exceptions

When a user marks a message as 'Confidential' in the archive, the message becomes invisible to all except administrators who have been given the confidential right, or any user which is specified in the exception list.

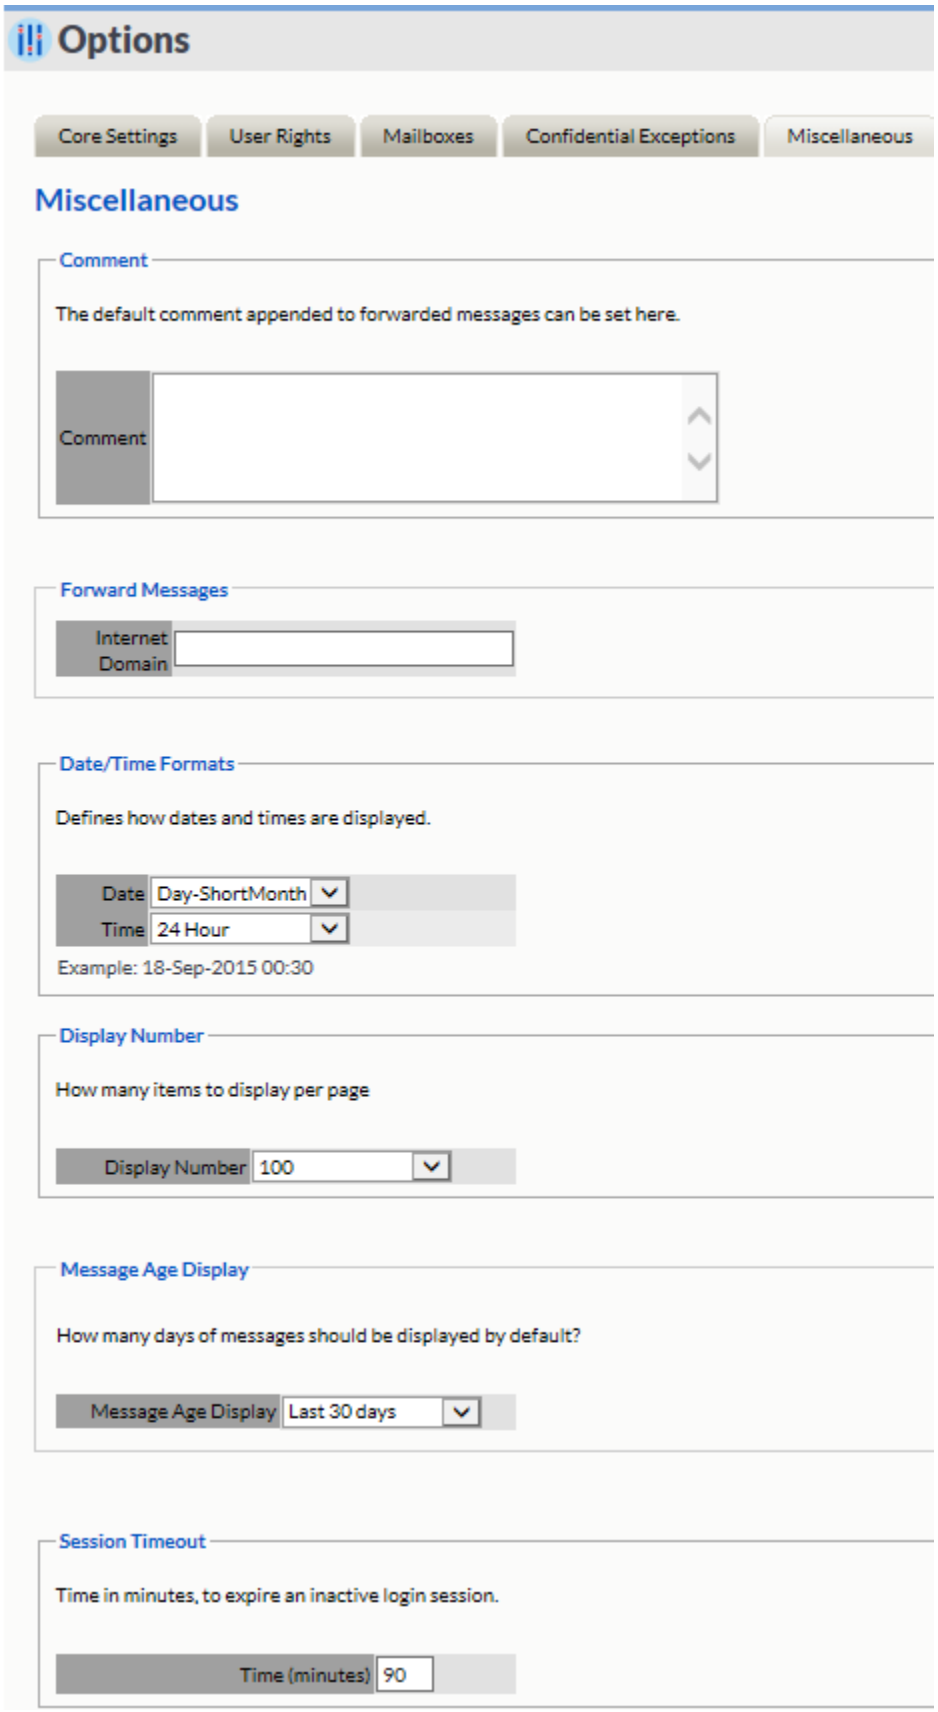


The ‘Confidential Exceptions’ allows users to add any necessary exceptions to the confidential tagging. Confidential tagging may be applied to protect sensitive. However, sometimes this information may need to be viewed by others and instead of granting that user rights to see confidential items for all users, a user may apply that right to only their items.

A group or individual user may be added to or removed from the list.

7.8.4 Miscellaneous Tab

This contains your display and session options


Options
Core Settings
User Rights
Mailboxes
Confidential Exceptions
Miscellaneous
Miscellaneous
Comment
The default comment appended to forwarded messages can be set here.
Comment
Forward Messages
Internet Domain
Date/Time Formats
Defines how dates and times are displayed.
Date
Day-ShortMonth
Time
24 Hour
Example: 18-Sep-2015 00:30
Display Number
How many items to display per page
Display Number
100
Message Age Display
How many days of messages should be displayed by default?
Message Age Display
Last 30 days
Session Timeout
Time in minutes, to expire an inactive login session.
Time (minutes)
90

These work just as described in the user display options in the Administration screen.

Options are:

Comment	Default comment for forwarding messages
Date/Time Format	How you want the dates and times to be displayed
Display Number	How many items to display per page
Message Age Display	Default date filter for searching - may be changed ‘on-the-fly’
Session Timeout	Inactive session timeout - set between 5 and 60 minutes

7.8.5 Tool Compatibility with Retain 3.x (Lucene) Indexes

Not all tools are compatible with the Retain 3.x (Lucene) indexes.

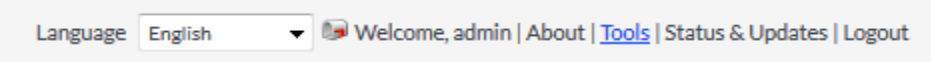
Tool	Compatibility
Retain web console	Compatible
Publisher	Not compatible
Outlook Plugin	Not compatible
Retain App for Android Phones	Not compatible
Retain App for iOS Phones	Not compatible
Retain Single Sign-on plugin for GroupWise Client	Does not use indexes
Retain Single Sign-on plugin for GroupWise WebAccess	Does not use indexes
GroupWise Archive Migrator/PAM Tool	Does not use indexes
PST Importer	Does not use indexes
Netmail and M+ to Retain Migration Utility	Does not use indexes
PowerShell Sync Script	Does not use indexes

8 Tools

8.1 Tools Overview

Retain has a number of tools available for importing and exporting data, as well as connecting to other programs.

To access the tools page, select the 'tools' link from the top right of the Retain Server administration page.



8.1.1 Tool Compatibility with Retain 3.x (Lucene) Indexes

Not all tools are compatible with the Retain 3.x (Lucene) indexes.

Tool	Compatibility
Retain web console	Compatible
Publisher	Not compatible
Outlook Plugin	Not compatible
Retain App for Android Phones	Not compatible
Retain App for iOS Phones	Not compatible
Retain Single Sign-on plugin for GroupWise Client	Does not use indexes
Retain Single Sign-on plugin for GroupWise WebAccess	Does not use indexes
GroupWise Archive Migrator/PAM Tool	Does not use indexes
PST Importer	Does not use indexes
Netmail and M+ to Retain Migration Utility	Does not use indexes
PowerShell Sync Script	Does not use indexes

8.2 Retain Publisher and Viewer

Retain Publisher and Viewer

Retain comes with the ability to export selected messages to a local archive for searching and viewing or to fulfill the need of a mobile archive for legal compliance. This can be extremely useful for larger systems or systems which have a high load, and where the need for review or legal compliance is being exercised, but access to the entire Retain archive is not necessary. Instead of allowing or facilitating constant access to the entire Retain Server, the Publisher can export and index groups of messages to a local archive, and the Viewer can search, view, and forward messages from the local archive.

The process is performed in two parts: the Publisher and the Viewer. The publisher, using an existing account with mail export rights, connects to the Retain Server and exports the messages complying with the search request, and creates a local database archive on the host machine. The viewer accesses the local archive and allows browsing, searching, and message exportation from the local archive.

Requirements:

- Network connection to Retain Server
- .Net 4.5 SP1
- Windows 7x, 8x 10x (32-bit or 64-bit)
- PST migration requires Outlook installed. (32 and 64-bit systems are supported.)
- Retain user with rights to export mail (See User management in Retain Administration)

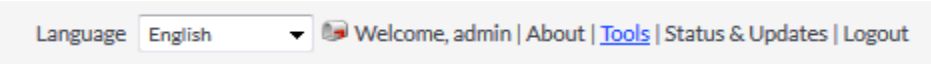
NOTE: Due to file access, the publisher cannot publish archives to a network drive

The Viewer can be run or be installed as stand-alone client wherever a published archive exists, but it is recommended to install the Viewer when you install the Publisher. While the Viewer and Publisher can be installed and run separately, the Viewer must have direct access to the published archive. Do not place the published archive on a network share. Accessing the published database from a remote machine may cause instability and is insecure. For this reason, it is best to use the Viewer local to the published archive.

Download the tools from the Retain site

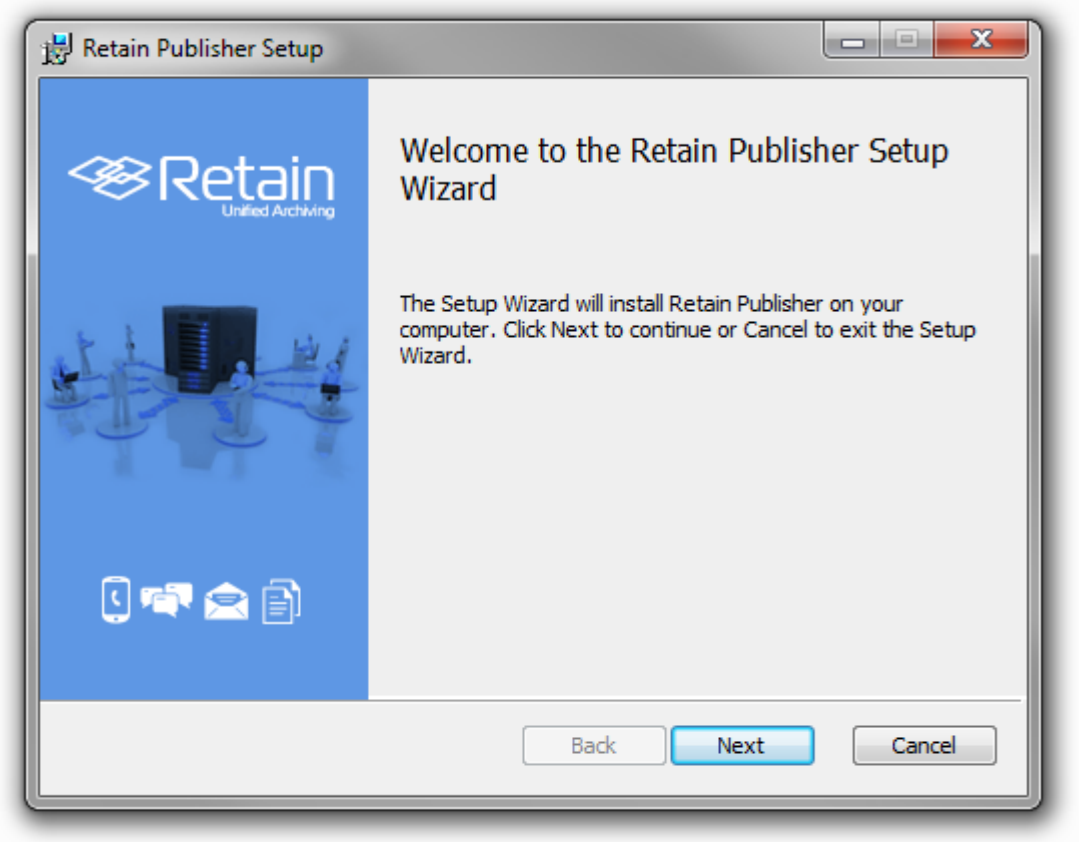
Run the installers on the desired machine.

The Retain Publisher and Viewer are found on the tools page. To access the tools page, select the 'tools' link from the top right of the Retain Server administration page.

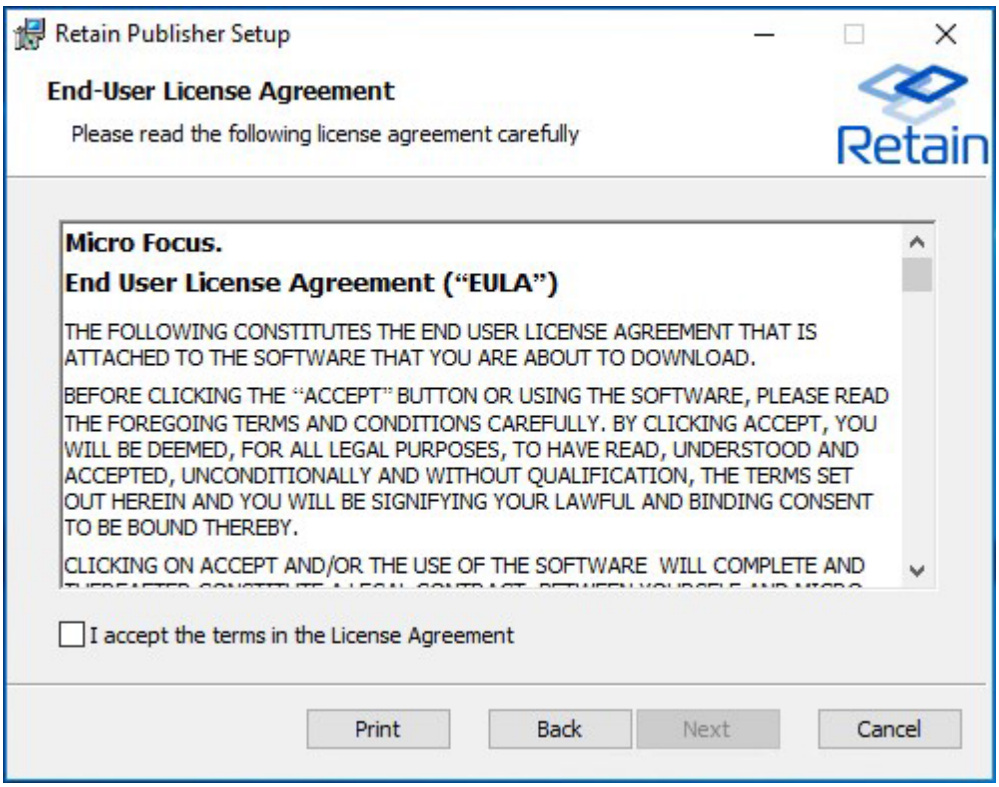


8.2.1 Retain Publisher Install

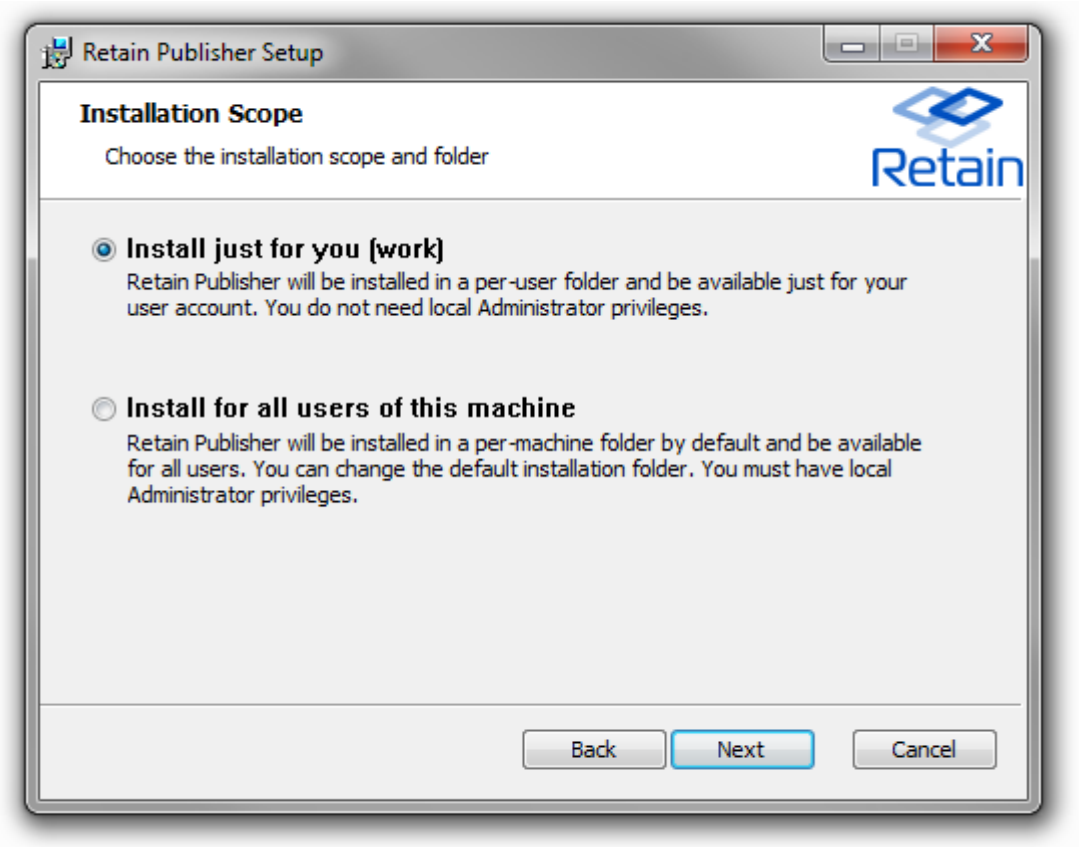
The Retain Publisher installation is very similar to the Viewer installation. Basic questions are asked, and the installer checks for .NET 3.5 SP1 before copying or installing any files. Click ‘Next’ to continue.



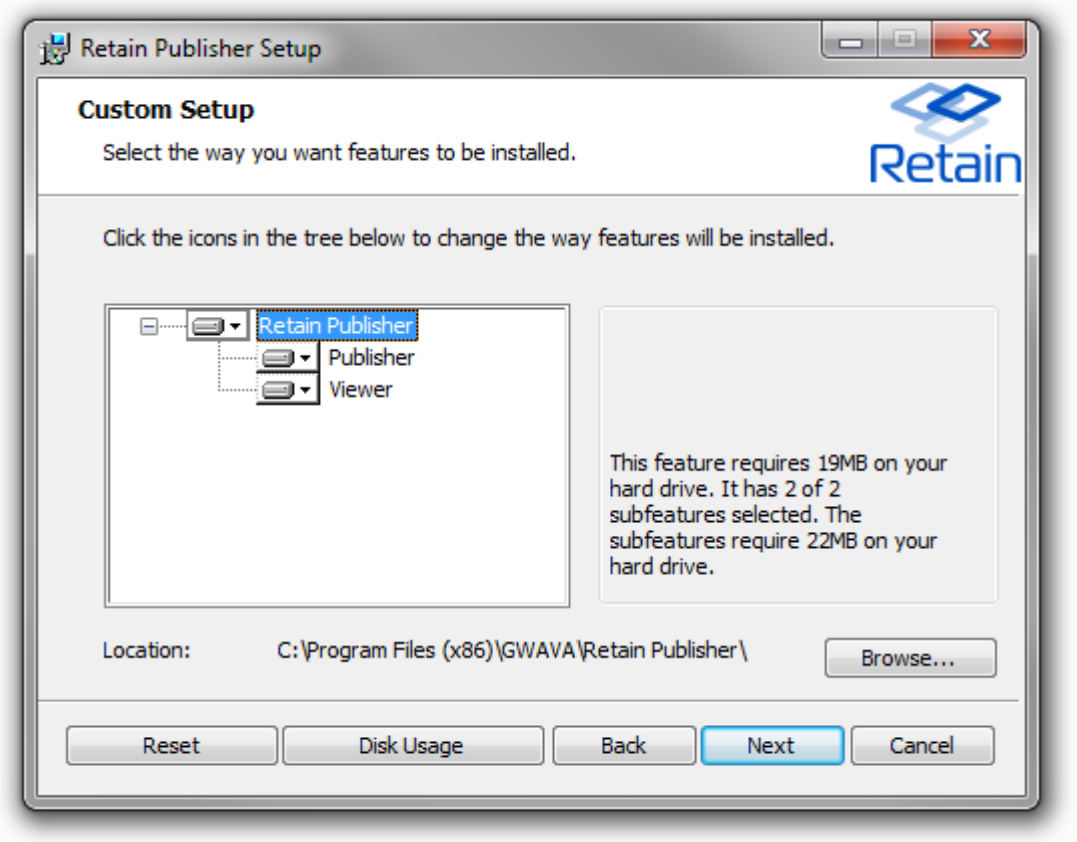
Read and accept the license agreement. Select ‘Next’ to continue.



The installer checks for any programs that need to be closed for installation. Close any applications speci-fied and click ‘Next’. Otherwise you will need to choose which users to install the software to.



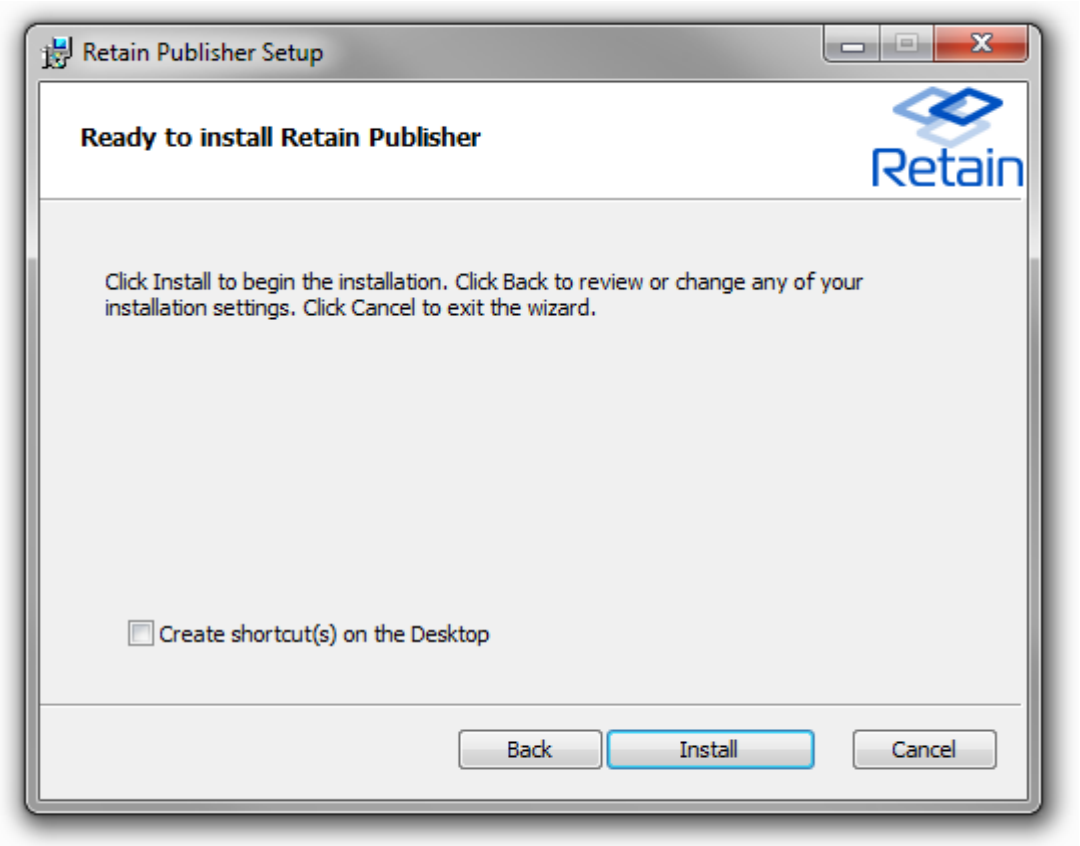
Select the scope of the install



Select which components to install.

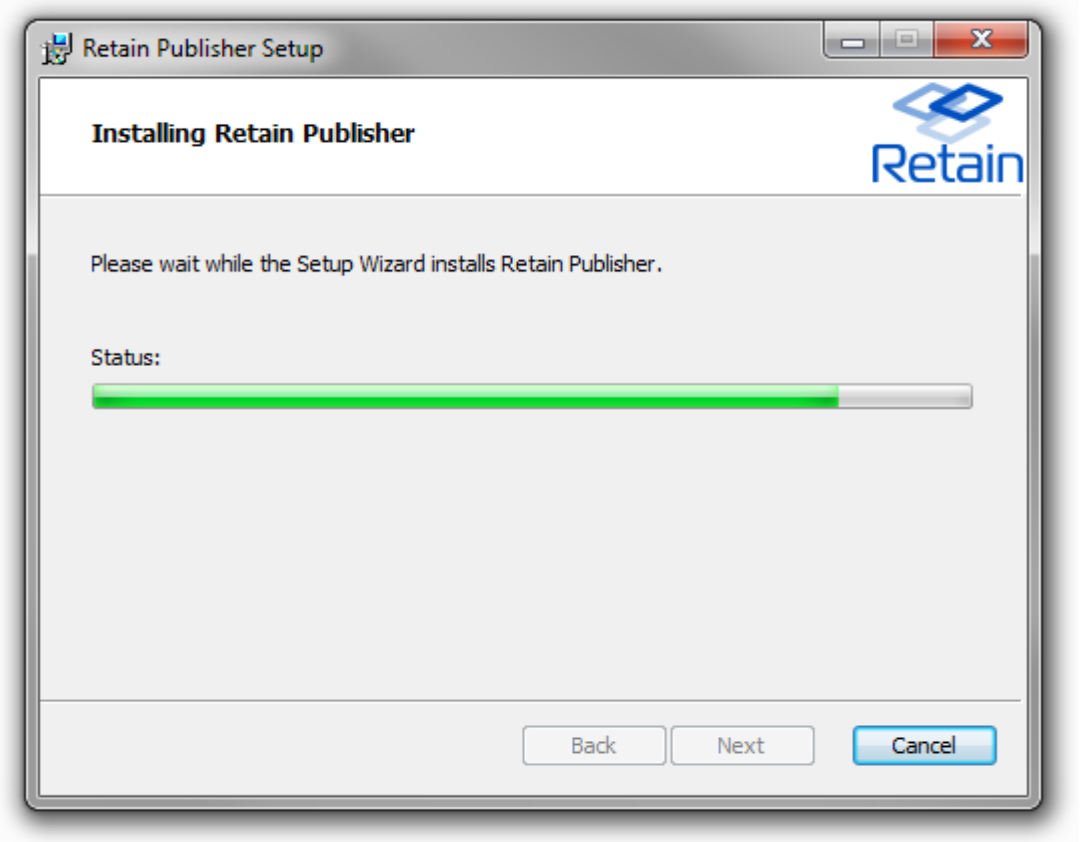
The Publisher installation provides the option to install the viewer alongside the publisher. If the viewer has already been installed, this is not necessary. It is recommended to have both the viewer on the same machine as the publisher. To only install the viewer, unselect the publisher. Select the desired setting and click ‘Next’ to continue.

Select the install location. The default is shown. If the default location does not work for the system, browse to, or specify the desired install location.

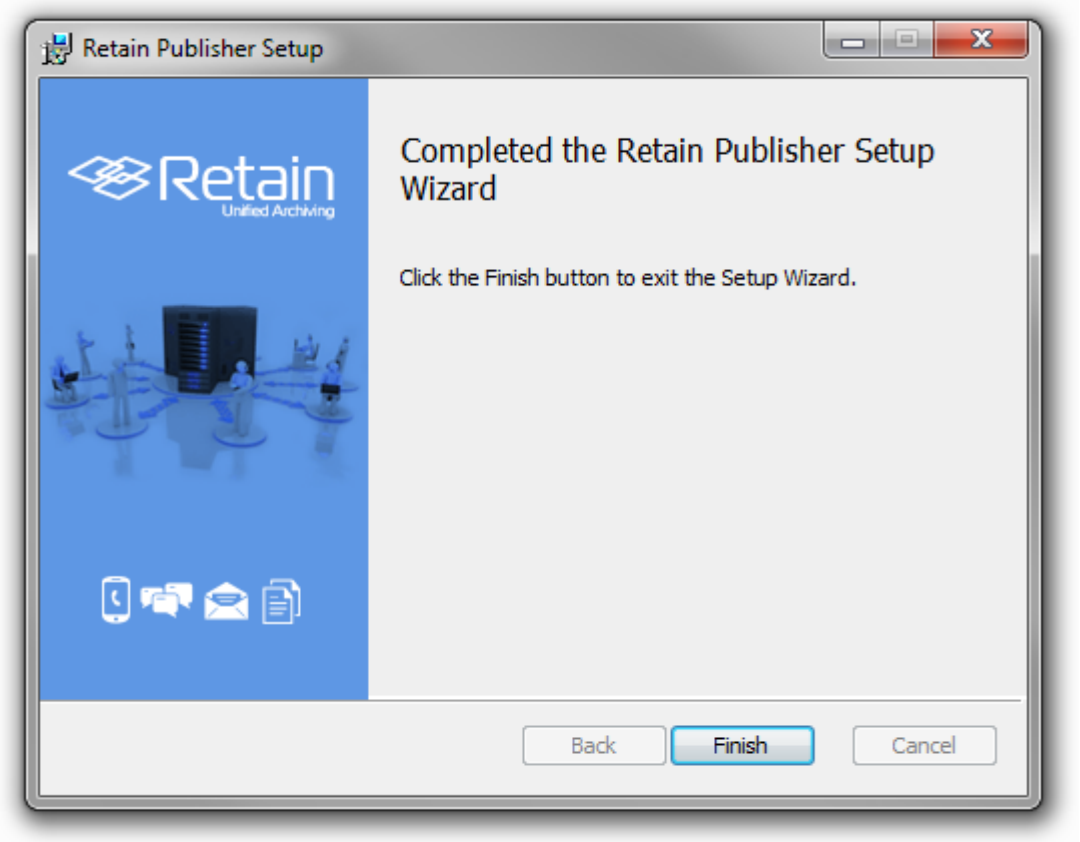


Select whether to create shortcuts or not. Default is to not create shortcuts. Select the check box to cre-ate shortcut(s) on the Desktop.

Click ‘Install’ to begin installation.



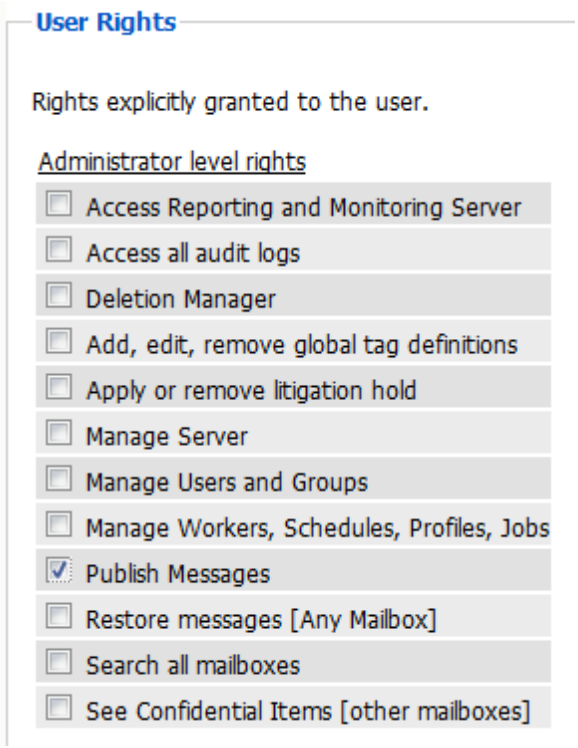
Wait for installation to complete.



After install has completed, select ‘Finish’.

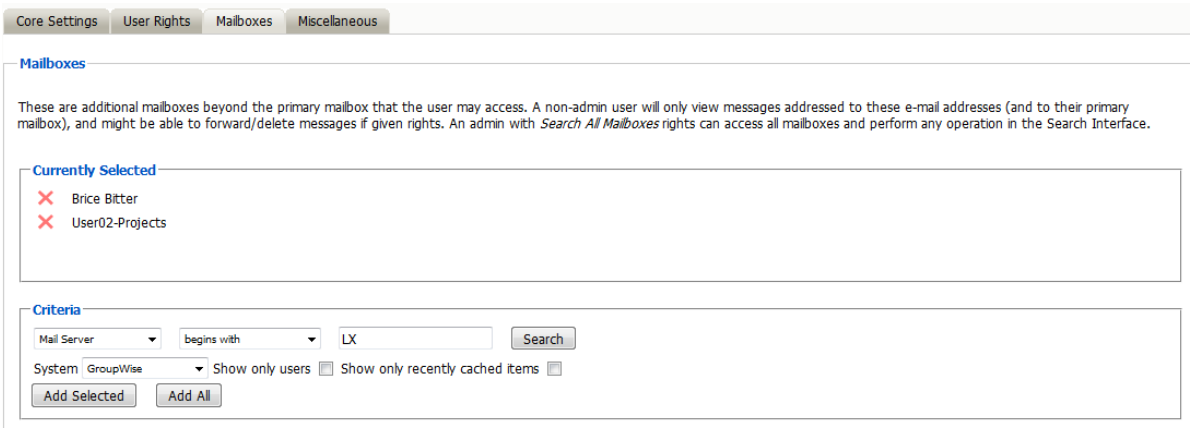
8.3 Retain Publisher

The Retain Publisher must be run to create the local archive for the Viewer to connect to. The Publisher does not display messages, it just accesses and exports messages into a portable message archive.



The Publisher must be used in conjunction with a user that has administrator rights to publish messages. Because the Publisher connects over the network, it requires an open network connection to the Retain Server.

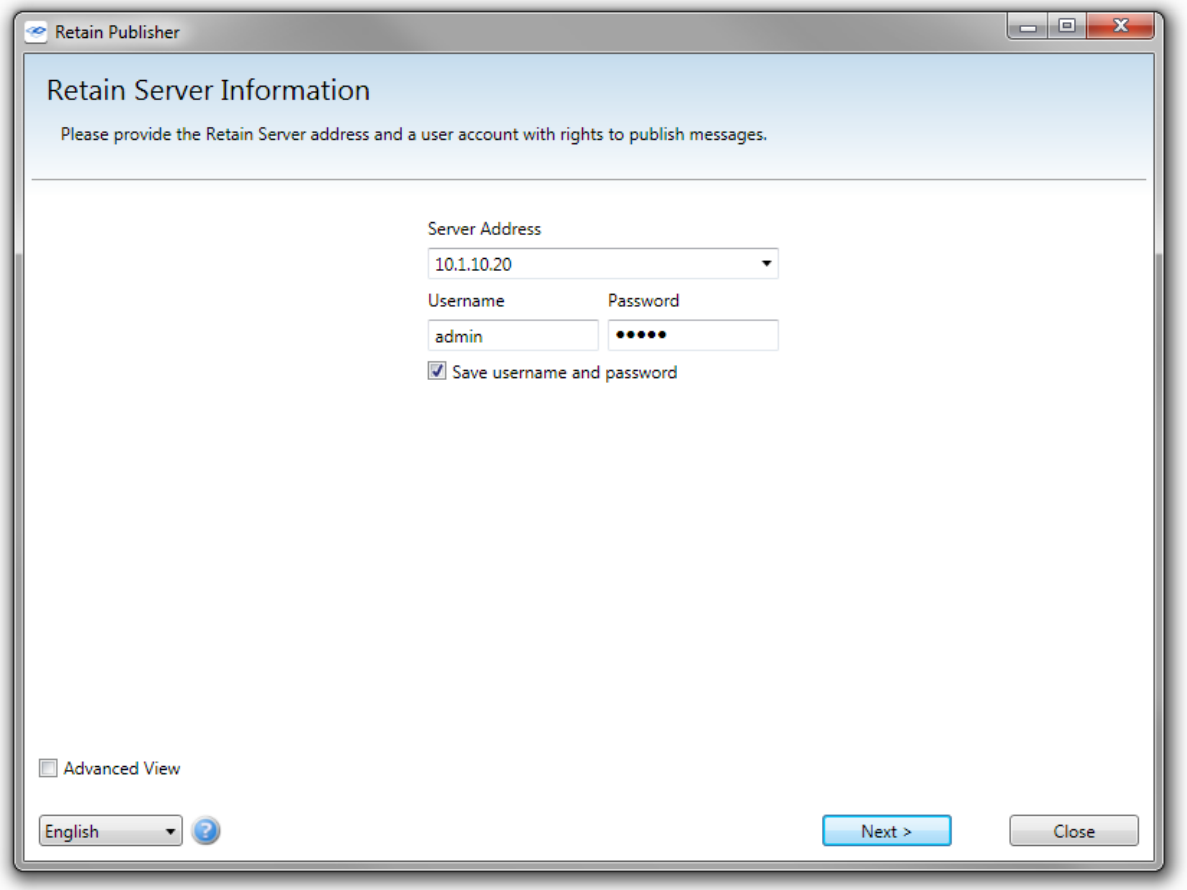
The Administrator account for Retain automatically has the publish messages right, and can be used here, though it is highly recommended to create and use an auditor account with the mail export right and access to the desired mailboxes or post offices required.



Start Retain Publisher by opening the program Publisher.exe.

8.3.1 Retain Server Information Page

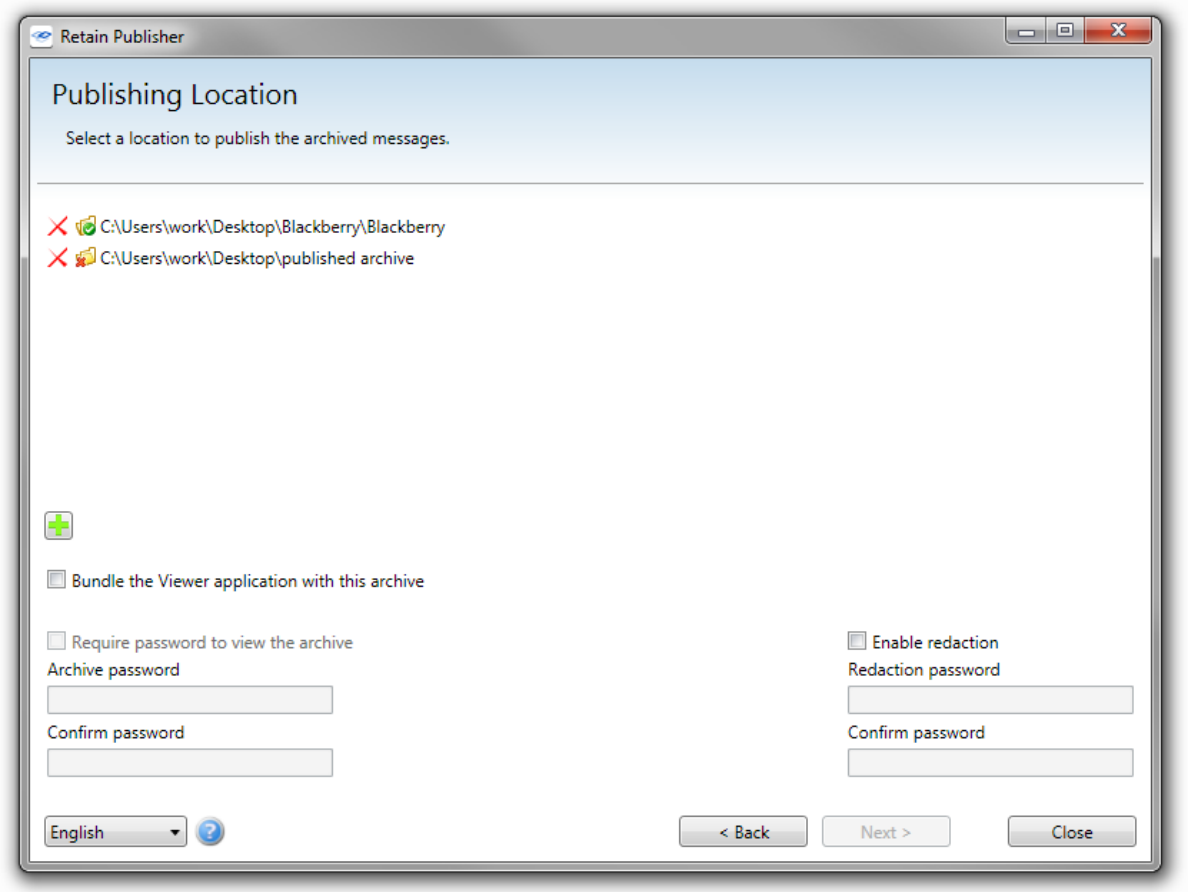
Input the Retain Server's DNS or IP Address and the login for the account with message export rights and rights to the desired mailboxes. Unless the login account has the administrator level right to 'search all mailboxes', only the mailboxes granted to the auditor account will be accessible. Granting rights to mail-boxes other than the active user's own mailbox are specified in the user rights section, under the 'mail-boxes' tab. (The advanced view provides a protocol and port options.)



Select 'Next' to login to the Retain Server.

Publishing Location Page

The Publisher then asks for the desired location for the exported archive messages. The messages extracted from the Retain Server will be saved in a database at this location. Select an existing location or select the green plus button and then browse to, or create a new folder for the published archive destina-tion. If a destination is selected which already contains an existing archive, a warning of overwriting an existing archive will be displayed.



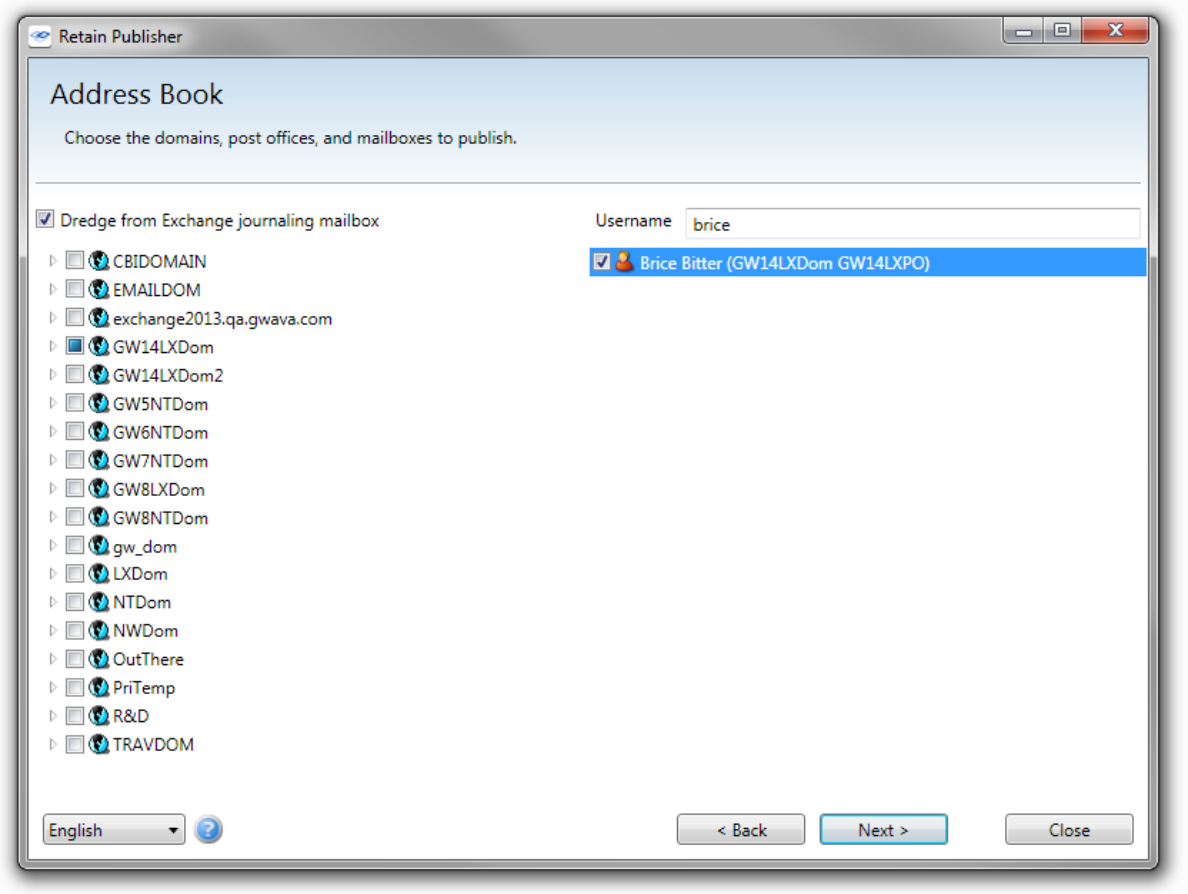
The options to password protect and use Redaction, both require passwords. To use them, select the checkbox and specify a password. The passwords can be different and both options may be present on the same archive.

When the 'Bundle the Viewer application with this archive' option is selected, the Viewer installation file is copied into the archive as well, preparing it to be completely mobile; the entire destination folder can then be copied or sent to any system with all that is required to read and search the archive. If the archive is to be distributed on a burned DVD or CD, further redaction will not function from the disk. All redaction comments should either be made before the archive is burned to a disk, or the archive should be copied from the disk to a local location where changes can be saved. When writing the archive to a read-only medium, it is always prudent to finalize the compilation.

Making the archive mobile can be very useful in situations where direct access to the Retain archive, such as during legal discovery, when email information is required to be surrendered for a user or group of users, an archive can be created and provided without disrupting current Retain operation. Click 'Next' to continue.

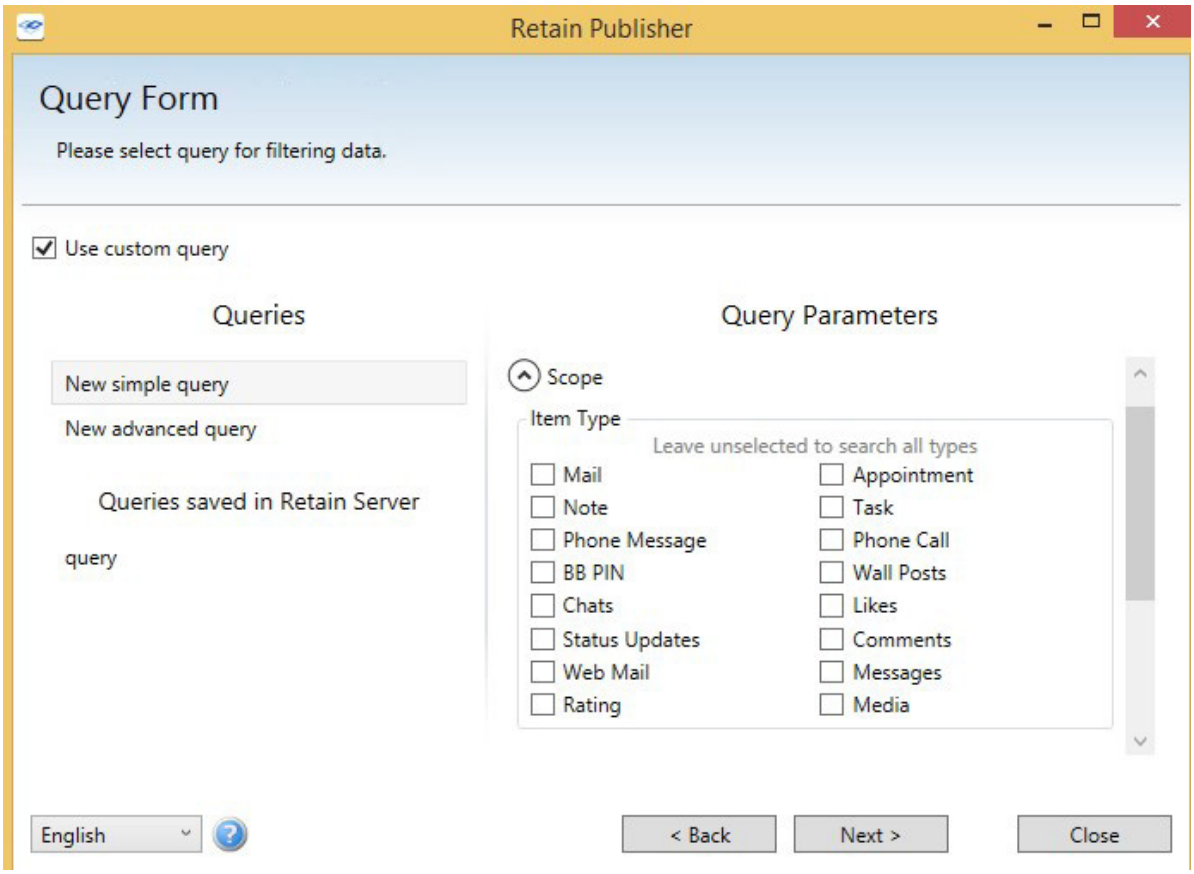
Address Book Page

Select the Post office and/or user(s) to dredge. Search for users by their display name.



Query Form Page

Select the scope of the export job. Single or multiple users, mail boxes, message servers, domains, and social media, even crossing message systems, may be added or excluded from any export job. Select ‘Next’ to continue.



The Publisher may restrict what mail or items are published into the message archive by specifying search or query terms. If a message or item does not apply to the search terms, it will not be added to the pub-lished archive. There are three opportunities to specify search terms or criteria; Core (query), Scope, and Misc.

Each of the three different criteria areas at this step may be shown or hidden by clicking on the hide/show arrow. Select the ‘Use custom query’ and either the simple or advanced query options to enter search criteria. Previously saved searches or searches that have been shared with the logged-in user will be displayed in the queries menu.

To add criteria to the Core query parameters, simply select the green ‘+’ icon, then select the different desired operator, search type, and phrase.

For instance, a very exclusive ‘contains’ search of the message contents will only publish messages that have the exact specified word or phrase in them. Fuzzy searches for approximates to the specified phrase, producing results including the words, but not in the specific order, or similar words.

The Scope details what types of messages will be searched for. If no items are checked, all are allowed, and that setting is default. To restrict to specific messages and types, (some are mail system specific), at least one item must be selected. If one item is selected, or something is desired to be excluded from the published archive, select all desired item types. The Item Type, Item Source, and Attachment size, are all independent factors and must be specified independently or left blank. A blank item source and attach-ment size will not restrict those criteria when combined with a specified type.

Attachment sizes can be restricted or allowed into the message archive, and are accessible through the drop-down menu at the bottom of the query screen, under the ‘Attachment Size’ menu.

```
vm-sles10-retain:/retain # l /retain
total 1
drwxr--r--   3 tomcat root   72 Mar   7 16:06 ./
drwxr-xr-x 25 root   root  608 Mar   7 16:01 ../
drwxr--r--   2 tomcat root   48 Mar   7 16:06 archives/
vm-sles10-retain:/retain #
```

Queries saved in Retain Server

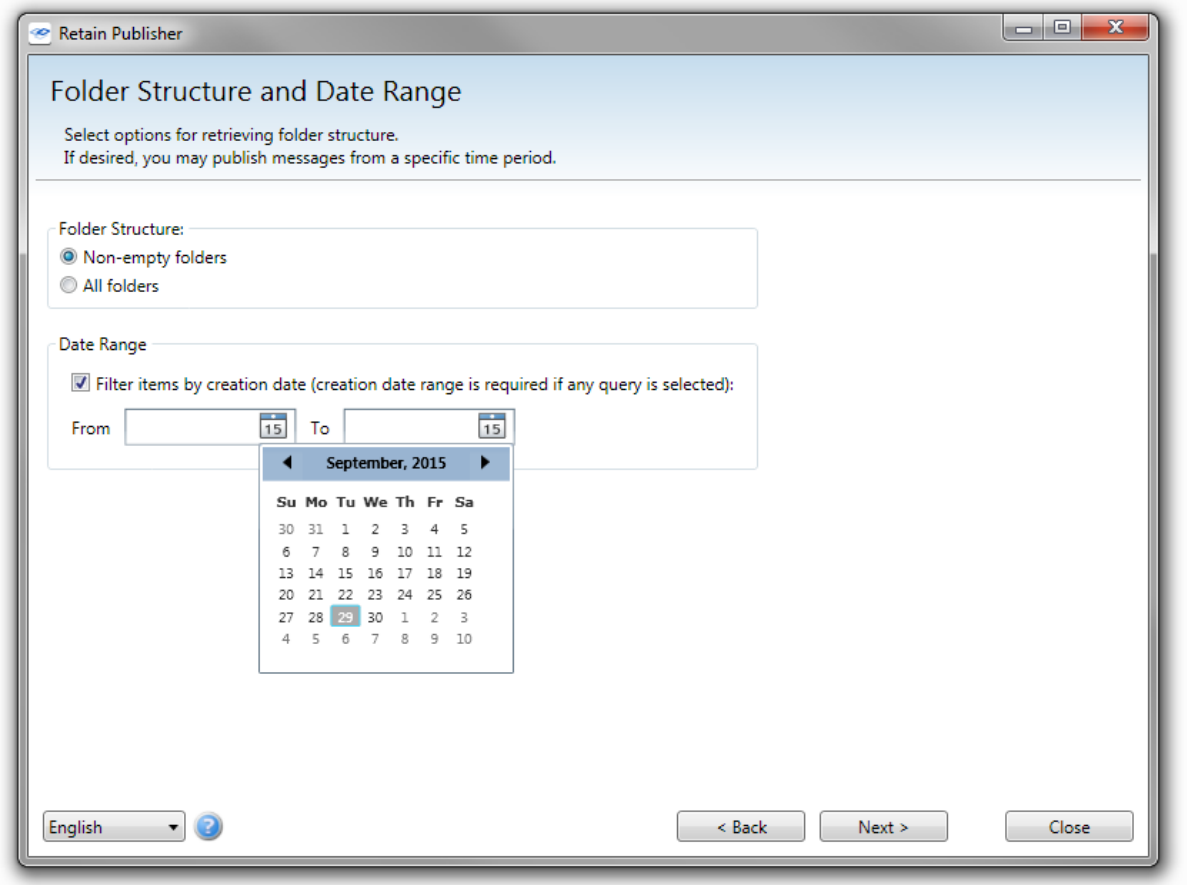
A query saved on the Retain Server can be loaded at this point.

Queries can be saved under the Search and Advanced Search tabs in the Search Messages Web interface. Publisher requires you to log in to the same account as the one used in the web interface to access the saved query.

A saved query will disable the date range selector on the next page. All criteria needs to be set in the web interface. The query can be changed in the web interface and will appear in Publisher next time it is started.

Folder Structure and Date Range Page

Choose how much of the folder structure and date range you wish to download. By default, the entire date range of the system will be searched.



Simply selecting the desired setting will make it active in the search criteria. Set as desired and continue.

The publisher can also be restricted to exclude all empty folders from the published archive.

The Date Range allows for specific item date ranges to be selected or excluded. The query will cause the publisher to only look for the specific mail required within the specified time frame. The time range may have a start date, and end date, or both. This feature is disabled if a saved query is being used.

This setting differs from the following time restriction in that the start and end times may be specified in a range, instead of a specific hard date. Only messages strictly adhering to the range will be published. If a date range is desired, enable the date range and select the desired range.

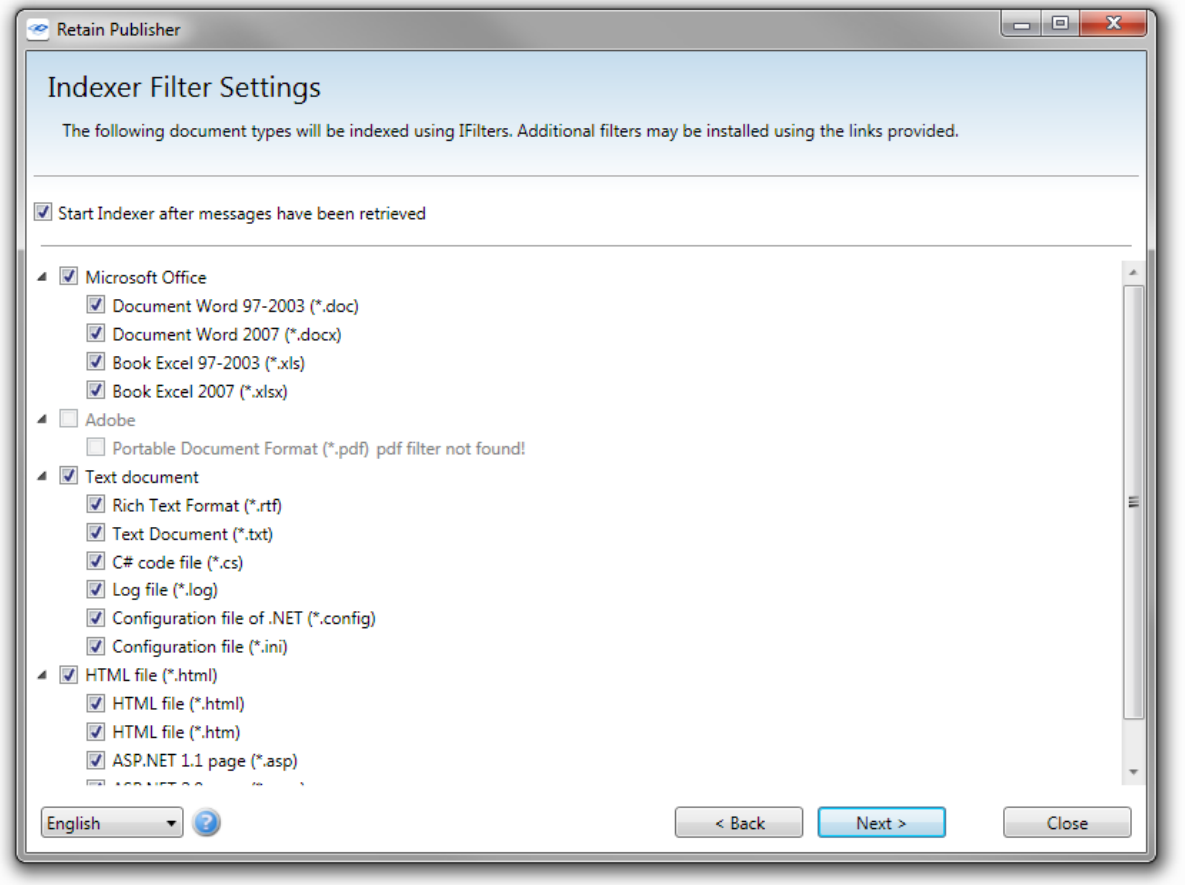
A date range is required to be specified if the active query is anything other than a previously saved query, so make sure it is set to an acceptable window for the desired data. The publisher automatically defaults the date range to one year.

Click on the ‘From’ or ‘To‘ dates to reveal the data selection calendar and specify the date as desired.

Click ‘Next’ to continue.

Indexer Filter Settings Page

Select what types of documents to be indexed. It is best to index all types of documents so all items may be searched.



The Publisher can create and index the archive for faster searching and browsing by the Viewer. It is highly recommended to index a published archive. All default filters are shown. To enable Adobe PDF indexing you will need to install Adobe PDF iFil-ter found on Adobe's web site.

Click 'Next' to continue.

If previously published messages are present in the selected publish location then you will be presented the option to overwrite the present information, or to append the new messages to the existing message store.

The Publisher will connect to the Retain Server and export the qualifying messages from the selected user(s) mailboxes and builds the local archive.

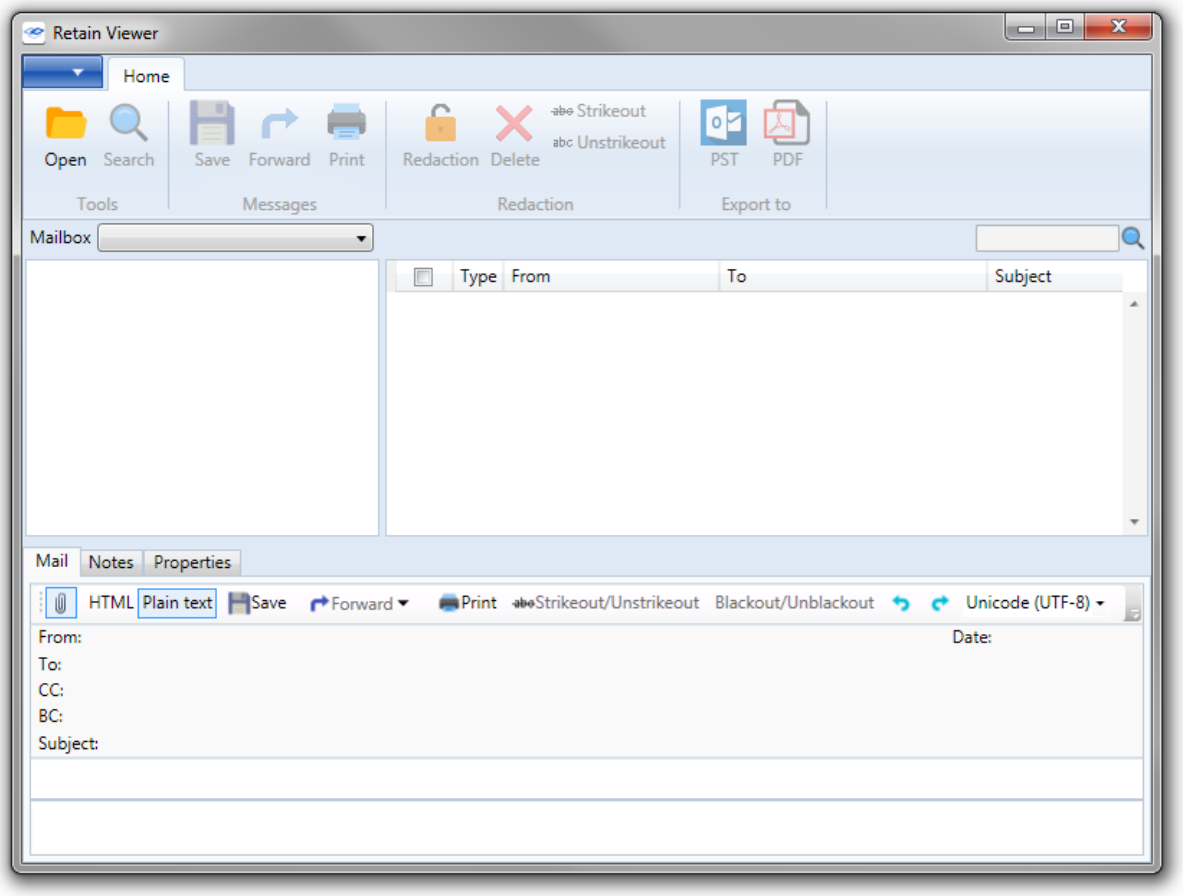
When the Publisher is done, it will display a report on the job. Depending on whether the exported archive was selected to be indexed after the job is done, the indexer will run. A .net error may occur at the end of Indexing, caused by an error in Adobe iFilter, if this error occurs, please install the latest Adobe PDF iFil-ter found on Adobe's web site.

Select 'Close' to exit the publishing wizard and the indexer will automatically run if that option was selected. With a successful publishing job, the archive is now ready to be connected to the Viewer.

If the Indexer was not automatically run, it must be run before the Viewer can search the archive.

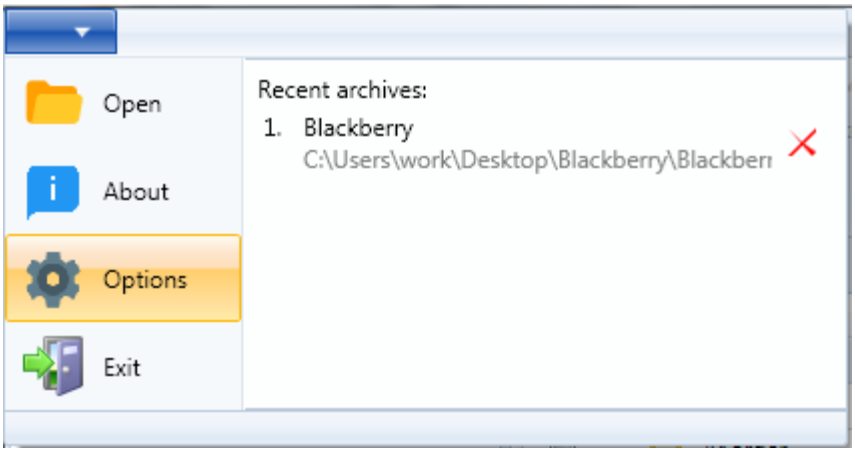
8.4 Retain Viewer

To connect to the published archive, locate the shortcut to the installed Retain Viewer and start the Viewer.

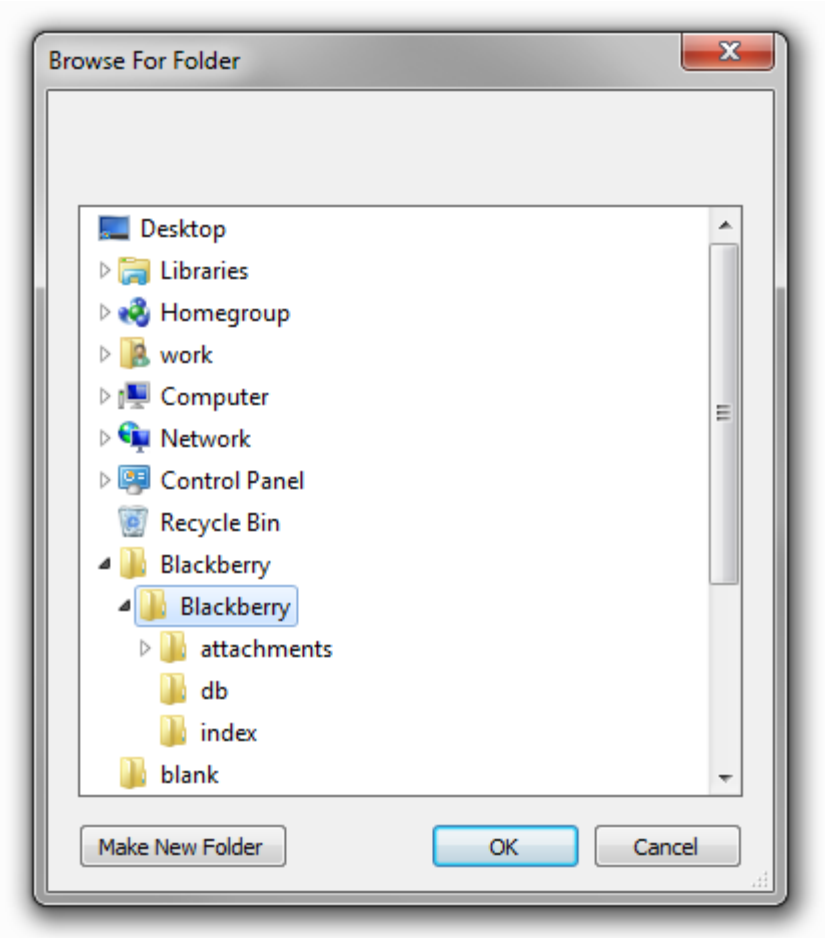


The viewer must be pointed to a valid database on initial startup. If the viewer does not prompt for the location to a published database, or if a new database is desired, select the ‘Open’ button and browse to and select the containing folder. If the Viewer has been previously opened, it will remember the last database opened.

The viewer is setup similar to an email client, with the mailbox and account on the left, the selected account’s contents fill the space on the top right, and any selected message’s contents are displayed along the bottom. Tabs allow access to the message properties or text, and highlighted buttons below the tabs determine what is shown and how, (Attachments, HTML, Plain Text). The Viewer also has the ability to forward the selected message out of the portable archive to any specified address.



To access an archive, select the ‘open’ button from the top toolbar, or select the ‘Open’ option from the ‘File’ dropdown menu.

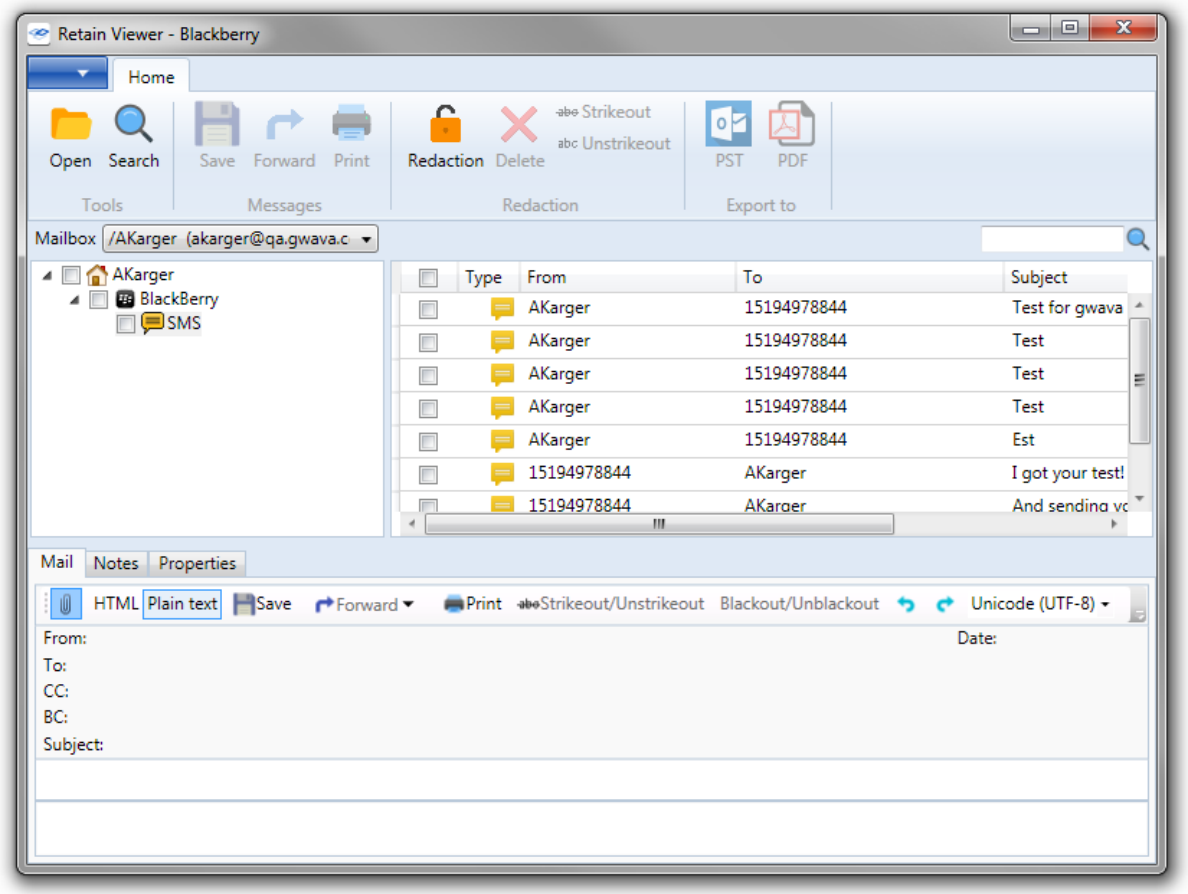


The Viewer only needs to know the base location of the archive, or the folder selected in the Publisher as the archive location. The ‘Open’ menu starts a browse window.

Browse to the location where the portable archive is located, select it, and click ‘OK’.

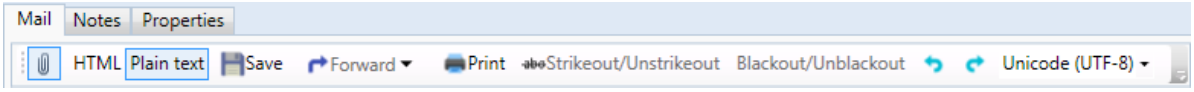
If the archive was password protected, the correct password must be entered before the Viewer can gain access.

Once the Archive has opened, all mailboxes contained in the archive are displayed. Select a mailbox to access the mail in the archive.

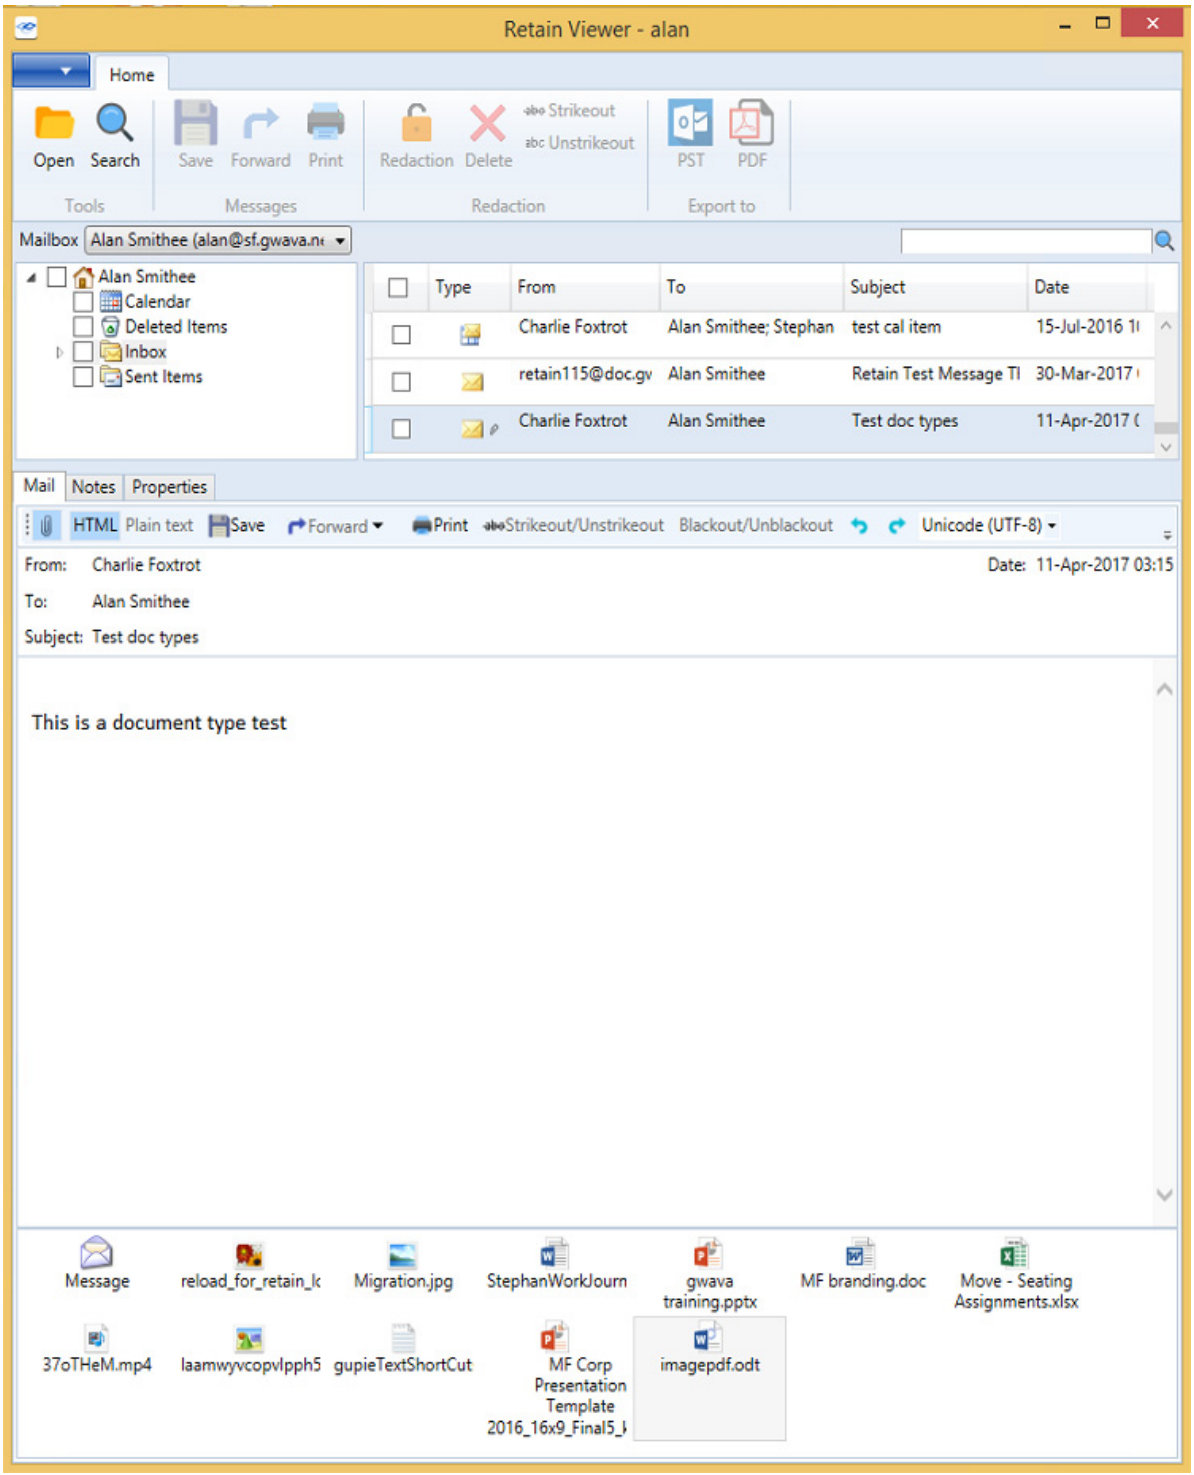


It is important to note that while the Publisher stores time in UTC, the time displayed in the viewer, located in the ‘Date’ column for each message, is relative to the viewer’s installed time zone, it is not in UTC. Messages published in PDF also display according local time zone. PST is exported in UTC, but Out-look rounds to the nearest minute, and displays the UTC time stamp in local time.

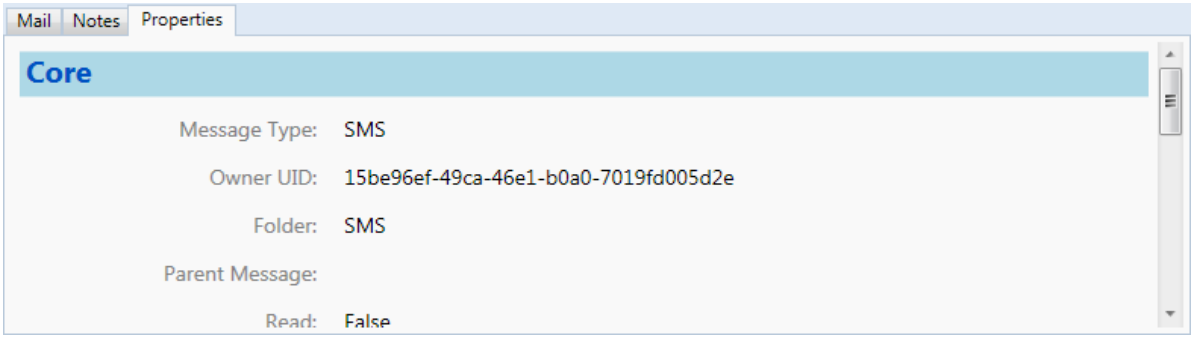
The text of the selected message will be displayed in the bottom viewer pane. The viewer contains options to show or hide the attachments, forward the message, as well as whether to view the message in plain text, or in HTML when available.



Attachments are shown along the bottom of the message window. Double-clicking on an attachment will send it to the operating system to open in the appropriate program.

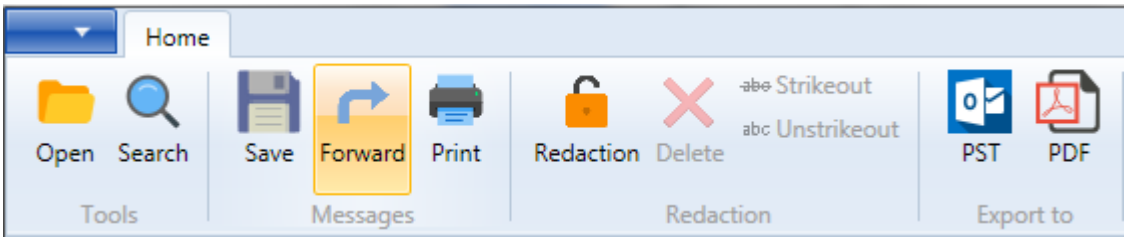


The properties tab displays the essential information on the mail item in question. The created, delivered, read status, and store date are displayed along with the identification number, message source, etc.

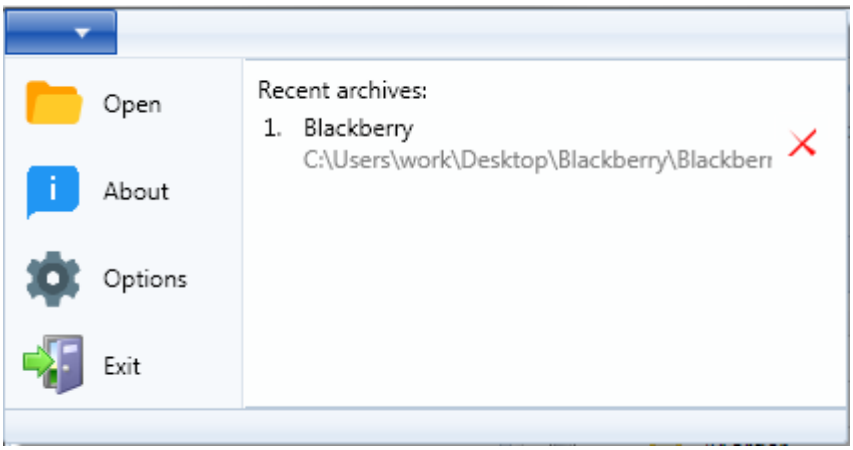


8.4.1 Forward

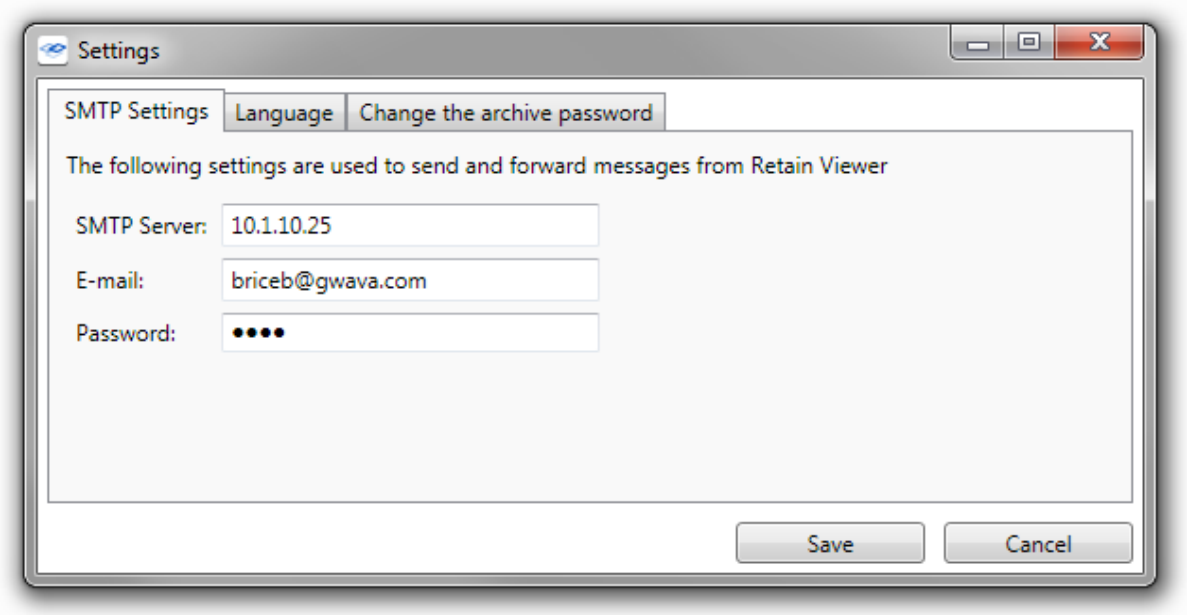
The Forward, and forward as attachment options, (found both on the main toolbar, as well as a drop down menu in the toolbar in the view window), allow the Retain Viewer to send the selected message from the archive to a destination account. 'Forward' functions exactly as the forward function in an email client. The 'Forward as attachment' creates an attachment from the selected email and attaches it to a message sent.



In order to utilize the 'Forward' option in the viewer pane, the Viewer must be configured with a mail server. To tell the Viewer which SMTP server and account to use to send messages, select 'Options' from the drop-down menu.

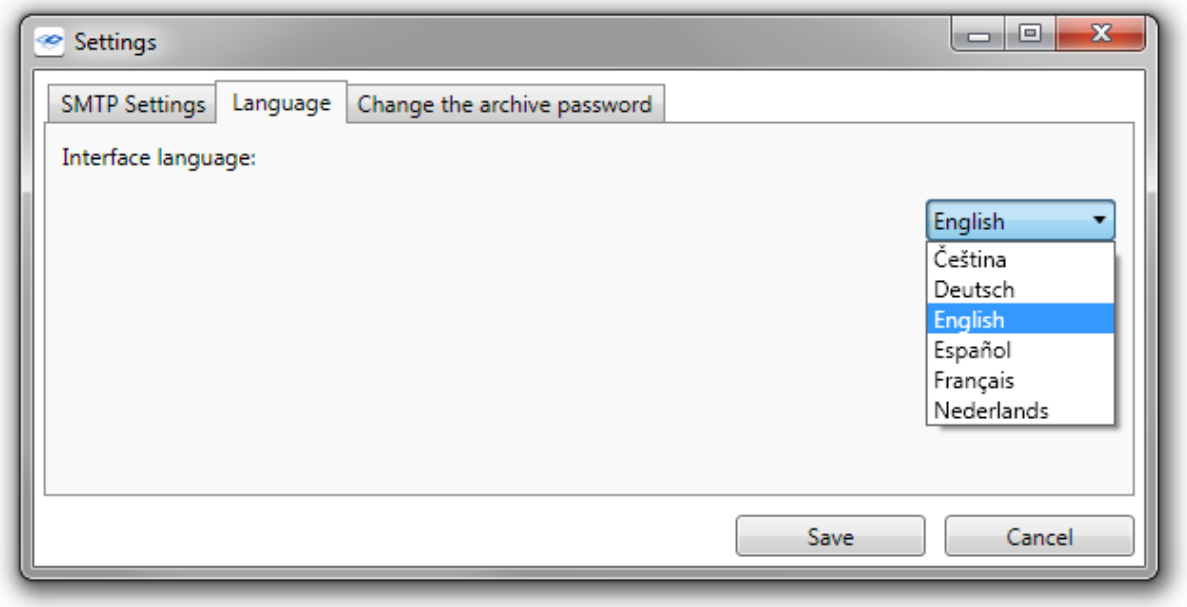


8.4.2 Settings

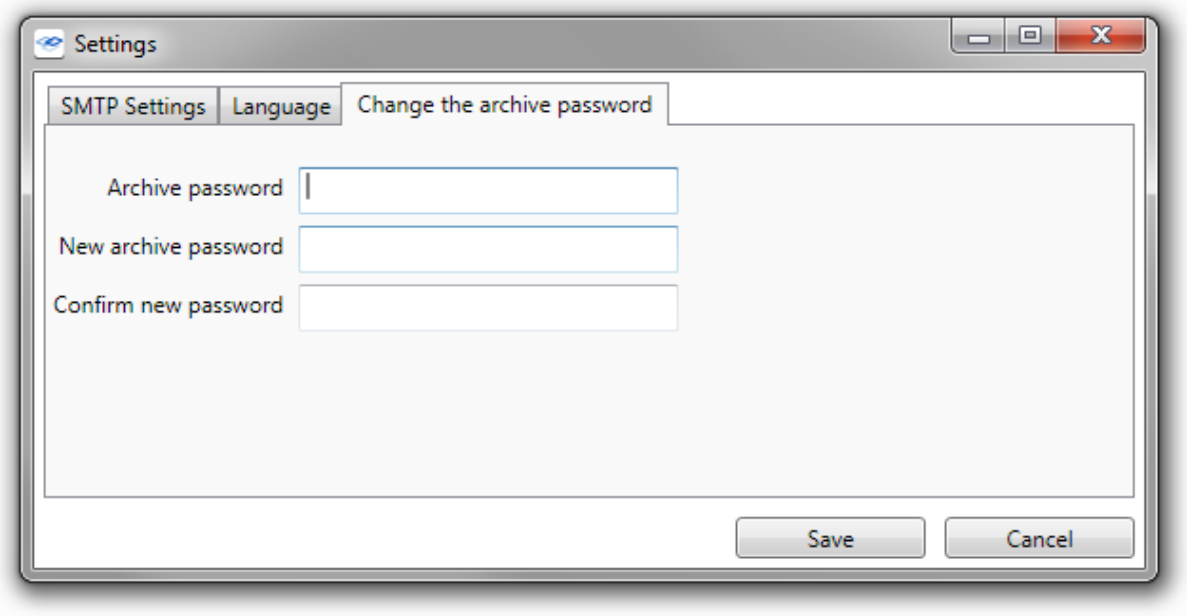


The ‘Settings’ configuration menu will appear. Enter the DNS or address of the SMTP Server desired to use for the Viewer system, and an appropriate account, (Username is required, password is optional), to connect and send messages.

Language and displayed document format settings may be changed as well for the specific viewer.



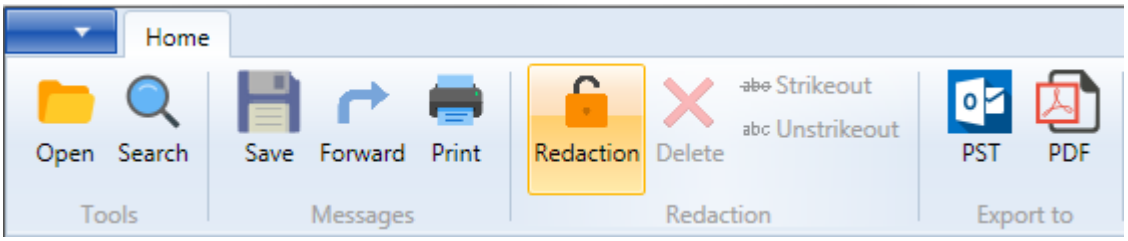
The ‘Change the archive password’ tab allows access to modify the password of the current archive. In order to change a password on an archive, the original password must be supplied first. If there is no password on an archive, a new password may be created; simply leave the current password field blank and input the desired password in the provided fields. While passwords may be changed, or created, the requirement for a password may not be removed.



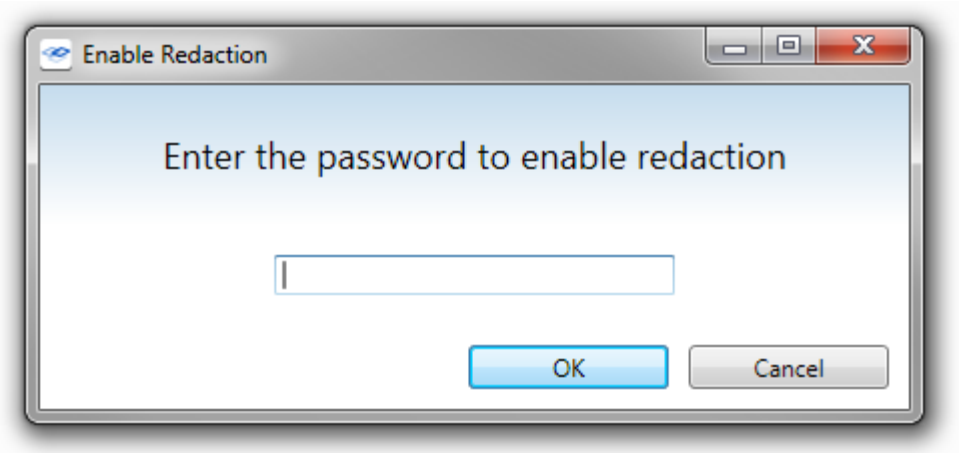
8.4.3 Redaction

Redaction allows the viewing auditor to compile notes on the archive.

NOTE: Attachments cannot be redacted directly because they require third-party software to open. However, the archive can be exported as PDF and a third-party plugin can be used to redact the PDF portfolio.



The notes are appended to the entire archive, and not any individual messages, which makes Redaction extremely useful to compile messages and identify messages of interest.

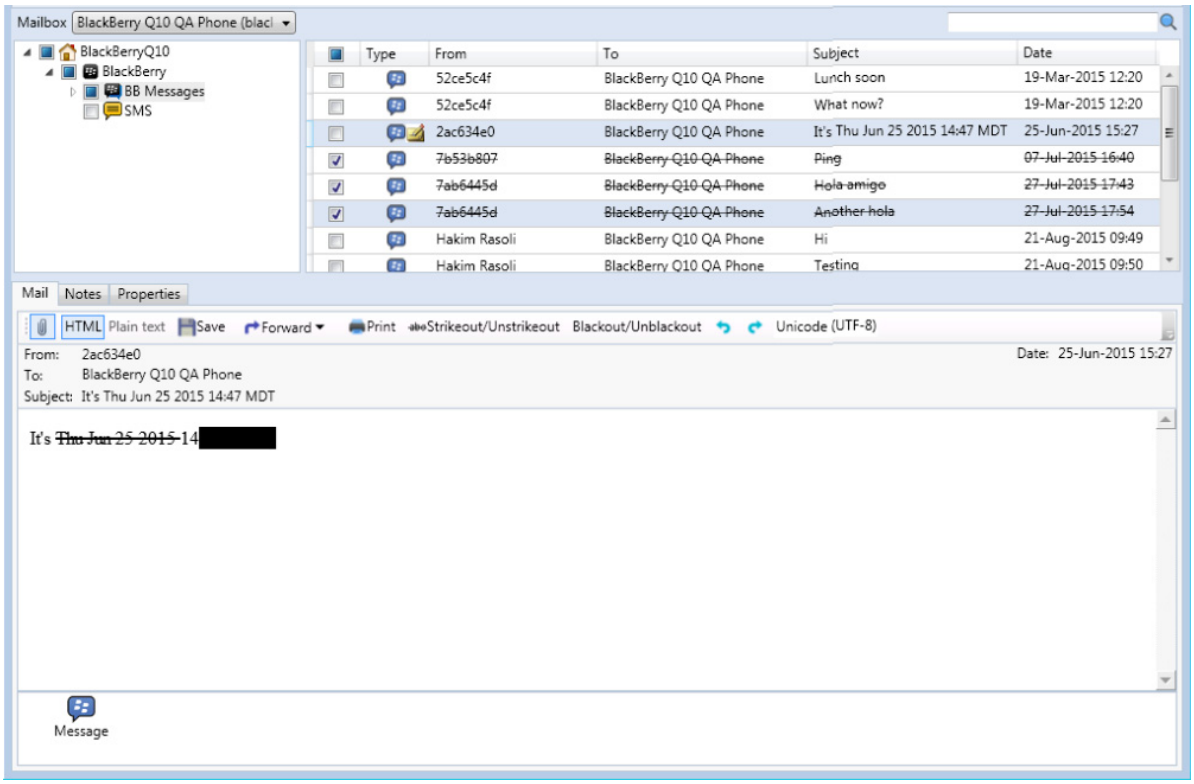


To enable redaction, the archive must first be exported with the redaction option enabled in the publisher, (if the archive was not published with Redaction enabled, the option will appear grayed out), and a redaction password must be entered.

On a redaction-enabled archive, to access the feature, it must be ‘enabled’ in the viewer. This option is found on the toolbar. Select ‘Redaction’ and enter the redaction password.

When Redaction is enabled, the ‘Strikeout/Unstrikeout’, ‘Black out’ options, and ‘Notes’ tab become active. The Notes tab works much like a notepad that is always connected to the archive. Items of inter-est or whole messages’ texts can be copied and compiled in the redaction notes. Messages with notes on them are marked in the archive viewer with a notepad icon in the message list.

Redaction notes are attached to the message. Notes are only accessible for the messages which have notes, indicated by the note icon.



8.4.4 Strikeout/Unstrikeout and Blackout

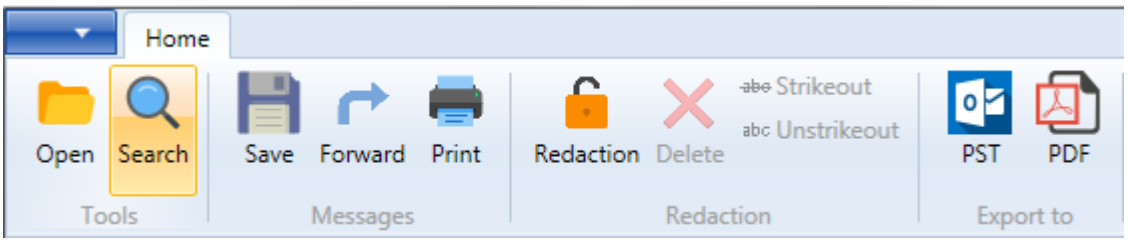
The Strikeout option enables a review of the archive before handing it over to legal scrutiny. This allows the user to denote items and messages which do not apply, or are unnecessary for the published goal. For example, if an investigation is being performed on specific activity, or internally on the account and all correspondence with a customer, the strikeout will allows messages or even sections of messages which do not apply, to be marked out.

To use the Strikeout, first Redaction must be enabled. After redaction is enabled, the options become active. To Strikeout or Unstrikeout messages in the message list, select the messages by placing a check in the checkbox and then select the ‘Strikeout’, or ‘Unstrikeout’ button from the top bar. To strikeout sec-tions in the body of the message, highlight the desired section and select the ‘Strikeout/ Unstrikeout’ but-ton from the view message bar.

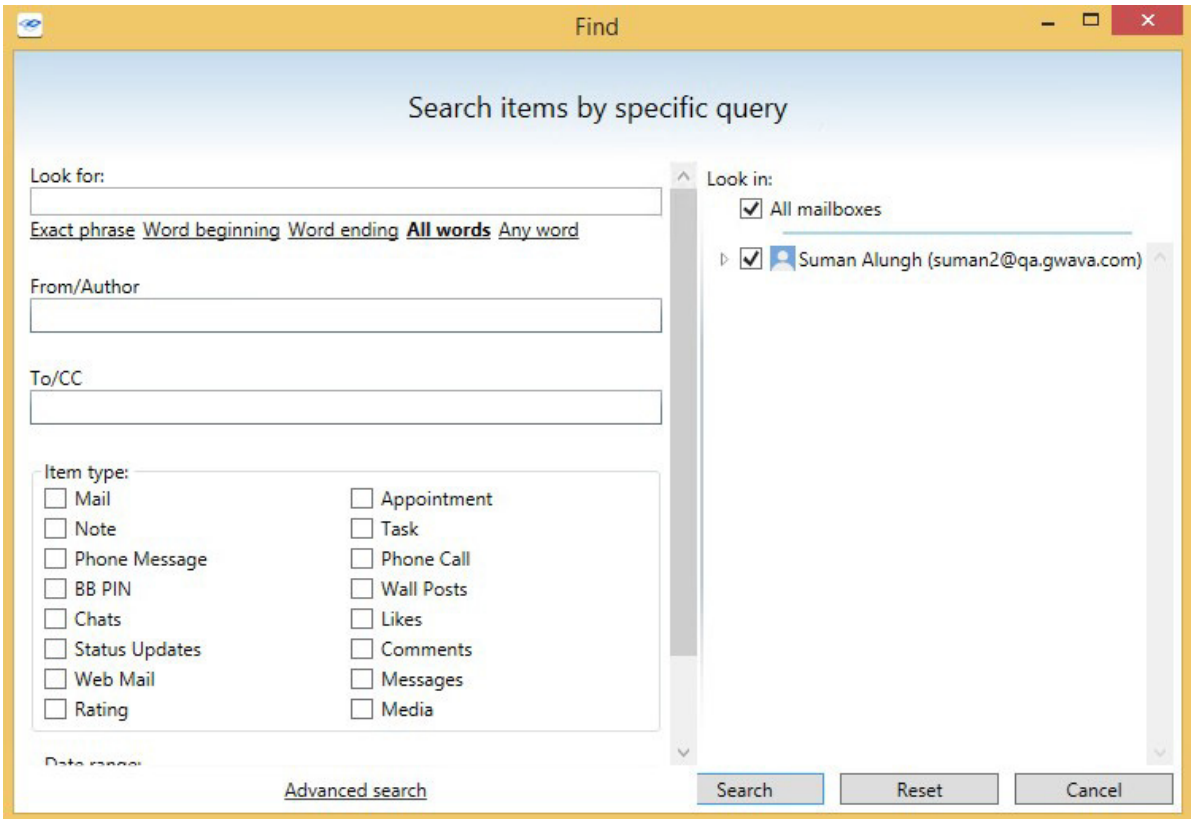
The Blackout option allows a reviewer to remove sensitive information, such as social security numbers, from published messages. To use the Blackout option, highlight the desired text and select the ‘Blackout’ button. Blackout cannot be removed from text.

8.4.5 Searching the Archive

The Retain Viewer provides full search functionality.



The search function is accessed through either the shortcut ‘Ctrl+F’ or by selecting ‘Search’ from the toolbar.



The different options for the search are shown with the criteria input below the options. Any text string or value can be searched for in the full text of the mail, or simply in the subject line. The following options are functional for both the search messages option included here, as well as the publisher search criteria. In the search ‘look for’ line:

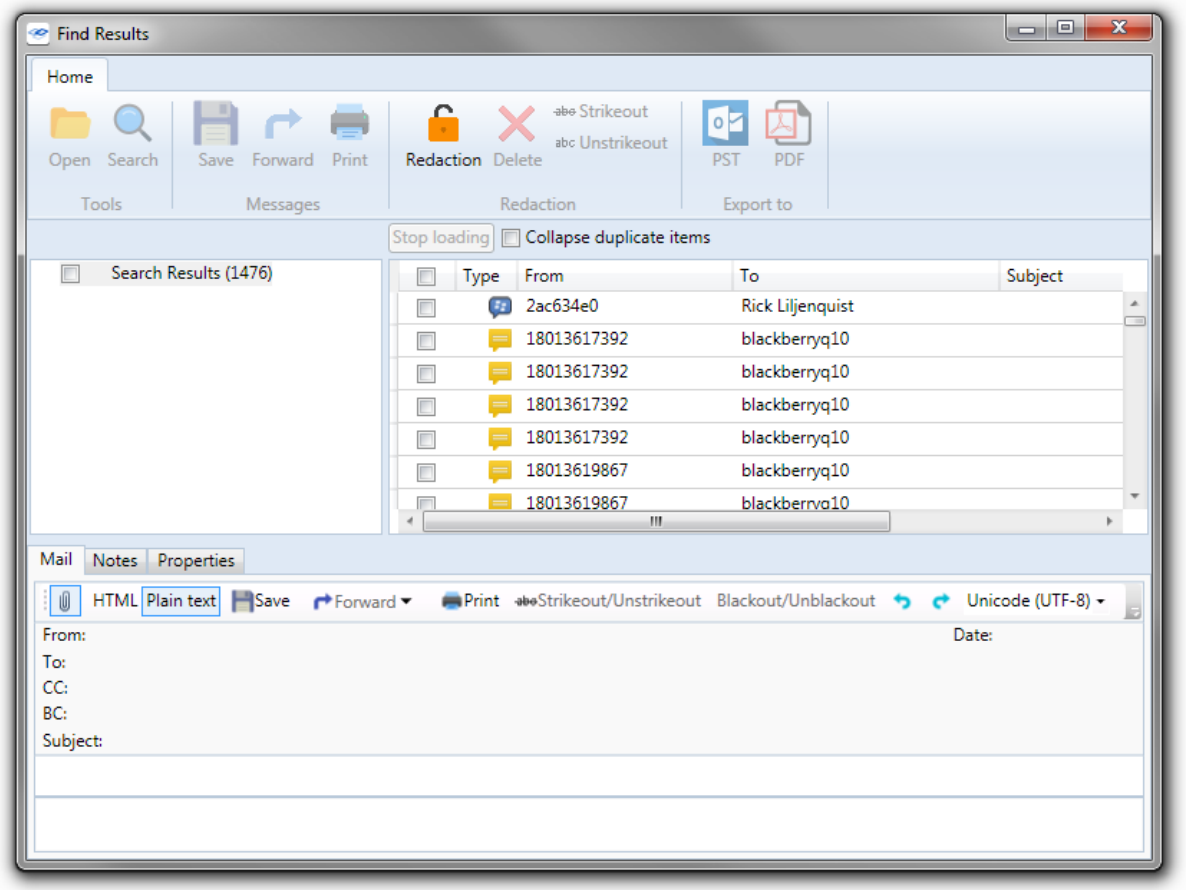
- && represents ‘and’
- || represents ‘or’
- No operators work as an exact match search.
- For example:
- Boy girl = exact search for “Boy girl”
- Boy && girl = Boy and girl search
- Boy || Girl = Boy OR girl search
- Boy && girl || dog && cat = Boy and girl OR dog and cat search.

NOTE: These search operators DO NOT work with the search option ‘starts with’, but can be utilized with all of the other base modes of the search.

NOTE: The search engine does not recognize the following English ‘stop words’ or articles of speech:

---

a, an, and, are, as, at, be, but, by, for, if, in, into, is, it, no, not, of, on, or, such, that, the, their, then, there, these, they, this, to, was, will, with



These words are not indexed, and thus are not taken into account if they appear in the list of words spec-ified for the search. This may result in some unanticipated results in “exact” searches.

In addition to the operators and text, the sending Author or the destination and carbon copy recipients of a message as well as item type and the date range can specify or restrict a search. Selection of the date range is done through an interactive calendar.

Be sure to select the desired mailbox or mailboxes to apply the search through. (The depicted example archive only contains one mailbox.) Select ‘Search’ to begin.

When the search is finished, the results will be displayed in a separate results window, identical to the main viewer interface.

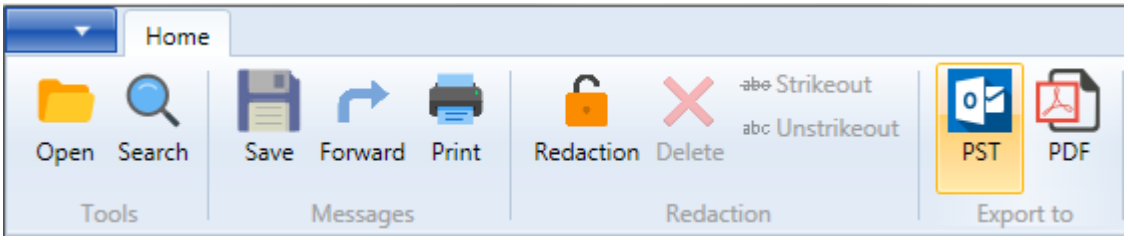
8.4.6 PDF and PST Export

Retain Viewer can export selected files and messages from a mailbox or archive to a PDF or PST archive file.

NOTE: PST export requires Outlook, or the associated plugin. The viewer will prompt for plugin instal-lation if necessary. It is highly recommended to have Outlook installed. The export PDF is a PDF Portfolio and requires Adobe Flash to be installed on the workstation to properly display the message selector header.

PST

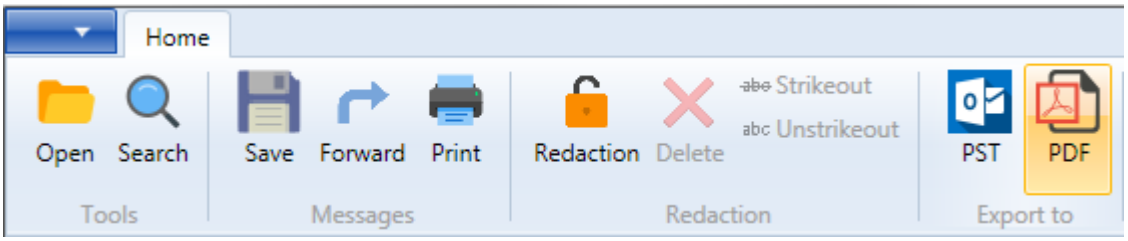
To migrate to PST, select desired files in the view window then select ‘PST’ from the toolbar.



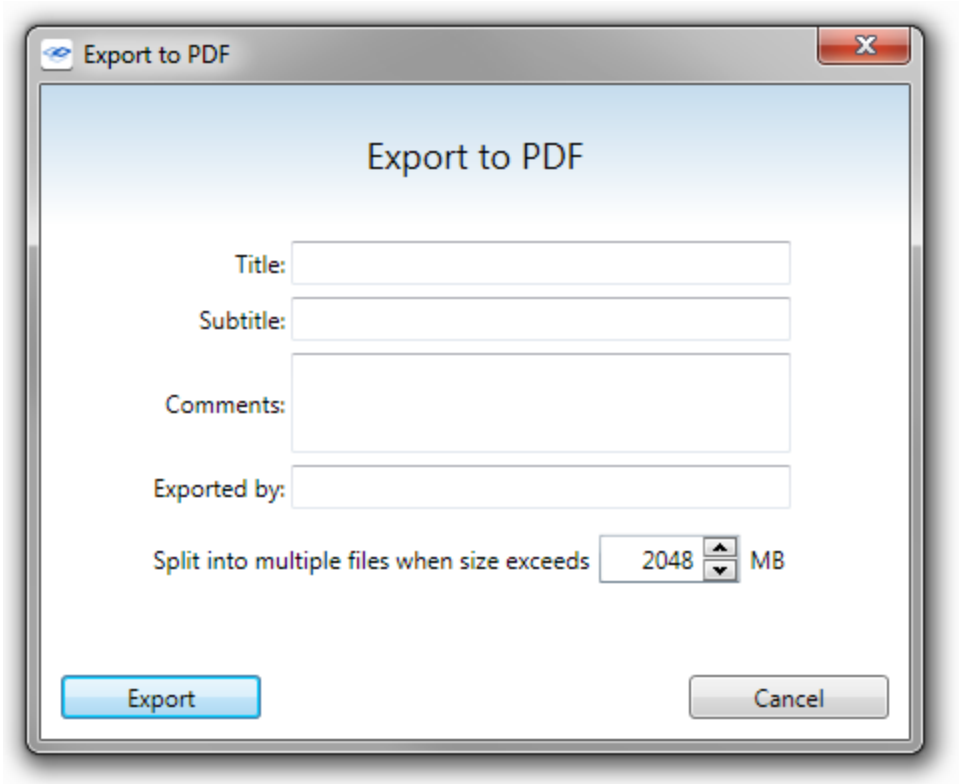
A window confirming the location and file will open. When the ‘Save’ button is selected the PST file is cre-ated and available for use.

PDF

The Retain Viewer can also save selected messages as PDF. Select the desired messages and click ‘PDF’ in the toolbar.



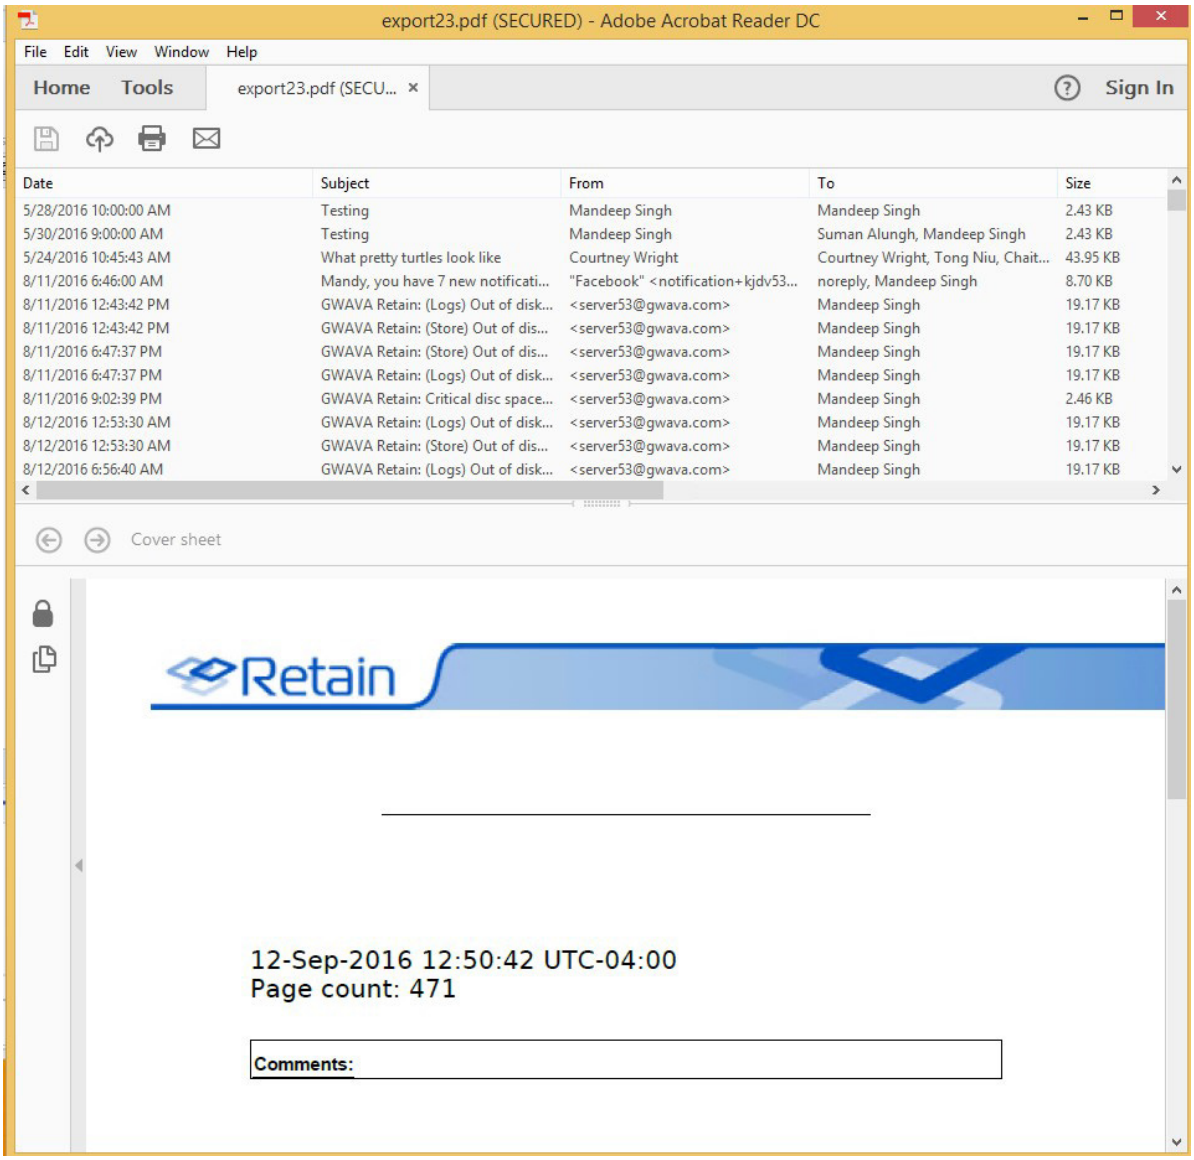
The selected messages are then published in a single PDF. This is an Adobe PDF Portfolio, this requires Adobe Flash to be installed on the workstation for full functionality, or you will only see the title page.



In addition to the standard file name and loca-tion, the PDF has title, exported by, and comment sections that the user is prompted to specify before the viewer publishes the PDF.

The finished PDF contains all the information displayed in the viewer, as well as an option to save individ-ual text or message sources. (This information includes redactions, litigation tags, confidential tags, strikeouts etc. from the Retain Server archive as well as the Viewer.)

The published PDF displays the separate emails along the top in a list and displays the message informa-tion in a lower pane, as shown below. Flash must be installed on the workstation viewing the PDF Portfolio file.



8.4.7 Enabling Debug Logging

On occasion, extra logging will be needed for troubleshooting purposes.

The Publisher and Viewer logs are found in the *%APPDATA%/Gwava/Retain Publisher* folder

To enable debug level logging:

1. Close Publisher, if open
2. Open File Explorer and enter *%APPDATA%* into the address bar
3. Enter the *Gwava* folder
4. Enter the *Retain Publisher* folder
5. Edit *Publisher.settings.xml*
6. Find the line <LogLevel>INFO,WARN,ERROR,FATAL</LogLevel>
7. Add the DEBUG level <LogLevel>INFO,WARN,ERROR,FATAL</LogLevel> *,DEBUG*

```
<?xml version="1.0" encoding="utf-8"?>
<PublisherSettings xmlns:xsd="http://www.w3.org/2001/XMLSchema" xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance">
  <LogLevel>INFO,WARN,ERROR,FATAL,DEBUG</LogLevel>
  <StartIndexer>true</StartIndexer>
```

8. Save and exit
9. Run Publisher

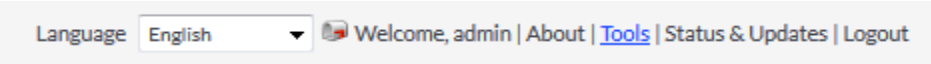
8.5 M+ and Netmail to Retain Migrator

8.5.1 System Requirements

- Java JRE 8 required
- 4 GB RAM recommended
- Dual-core processor or better recommended
- Windows 7.x (64-bit), 10
- Retain 4.0 or higher

All archives must be local to the Migrator, or located on an active mapped drive. The migrator requires direct file access.

The M+ and Netmail to Retain Migrator is found on the tools page. To access the tools page, select the 'tools' link from the top right of the Retain Server administration page.



8.5.2 Directories and Logs

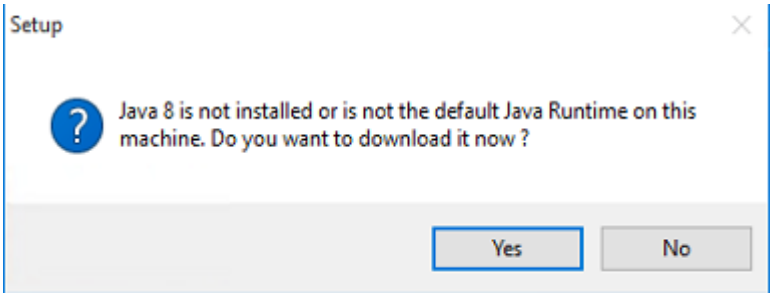
Logs are created, one for each day, in the following location by default: "C:\Users\Administrator\AppData\Roaming\GWAVA\M+ and Netmail to Retain Migrator\logs.

Logs contain all errors as well as migration info for all users. After setting up the Migration Context, a .conf file can be saved anywhere and will be loaded automatically the next time the migration tool is loaded.

8.5.3 Install

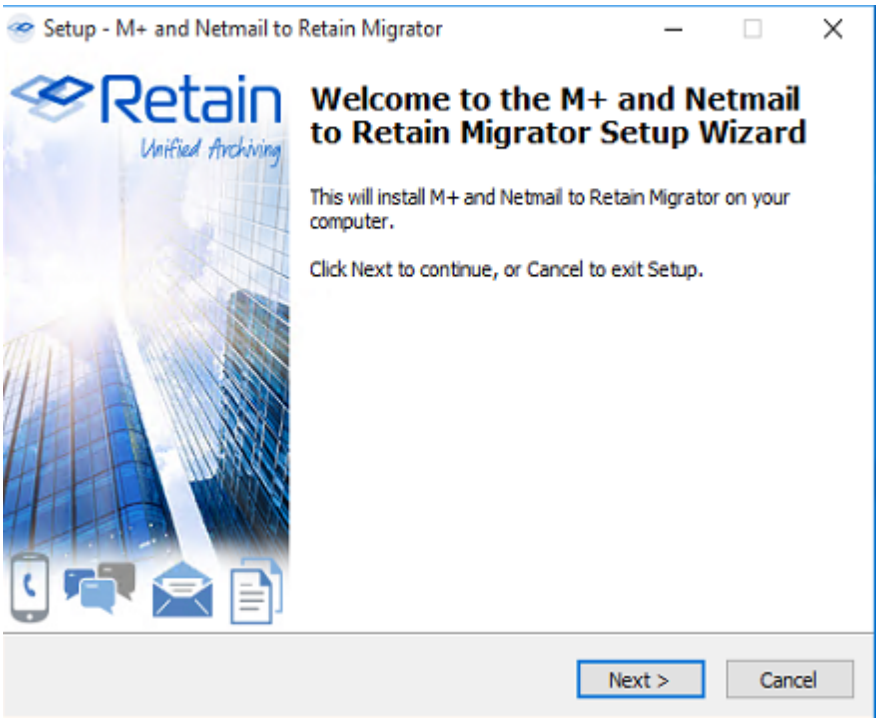
The install is straightforward. Simply execute the installation file and follow the prompts.

If Java 8 is not installed, the installer will prompt you to download it.

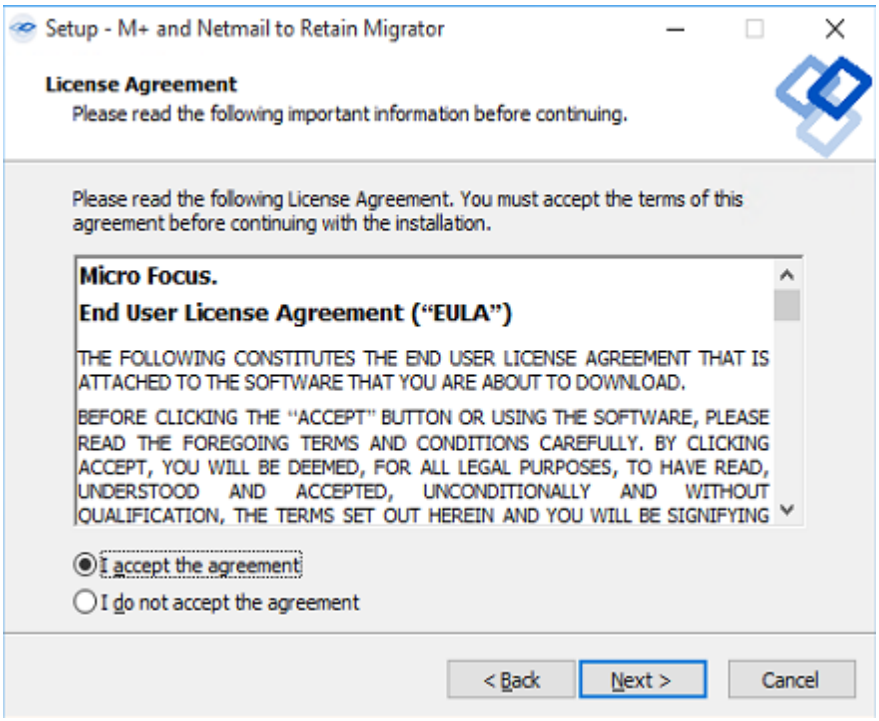


Press Yes Retry to be taken to the Java 8 download page, please choose the appropriate version for your system and install it, then press *Retry.*

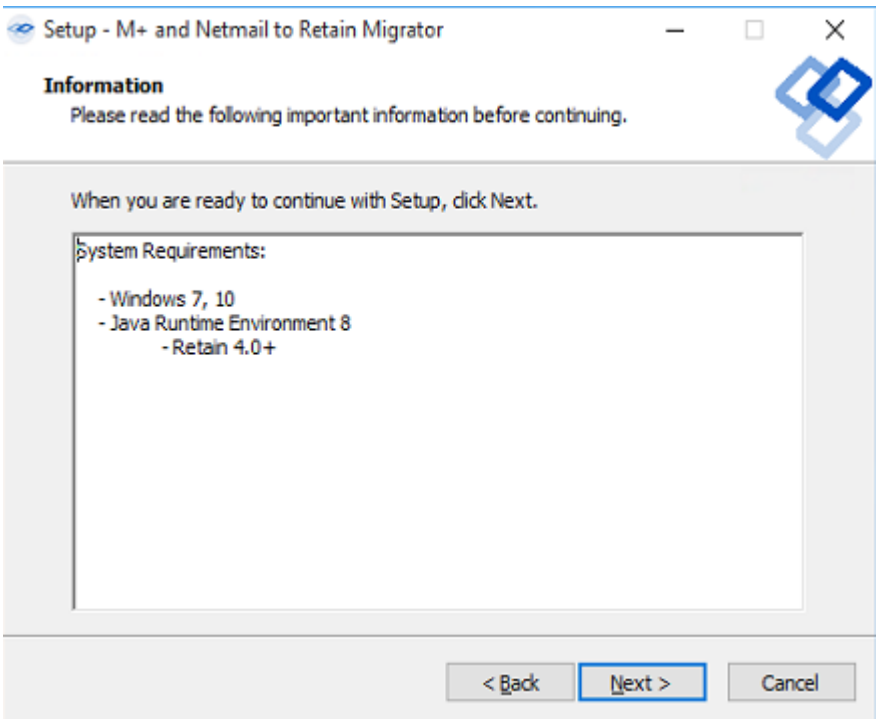
The install will welcome you.



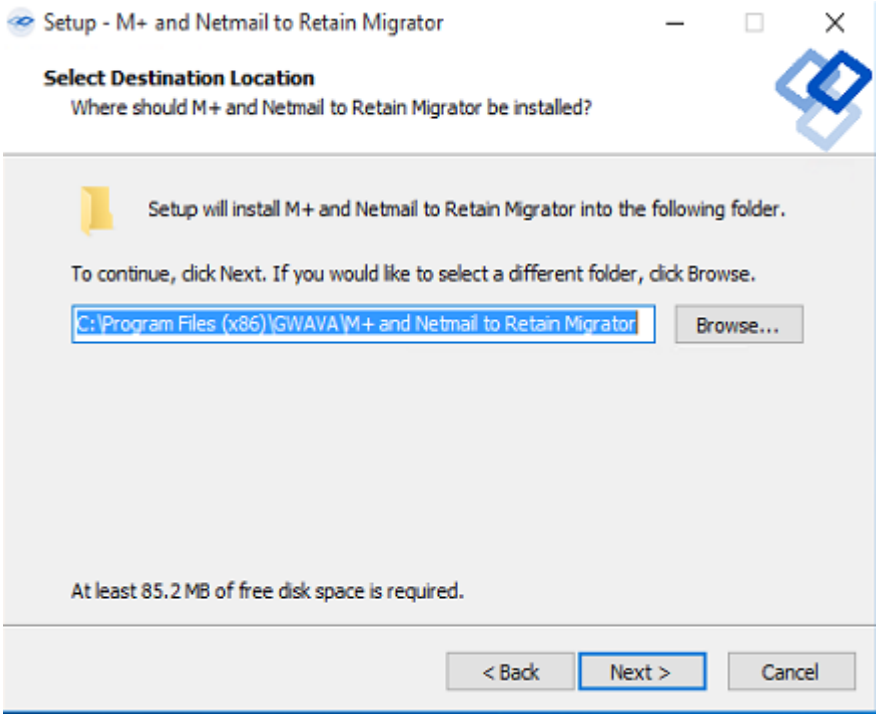
You will need to agree to the license.



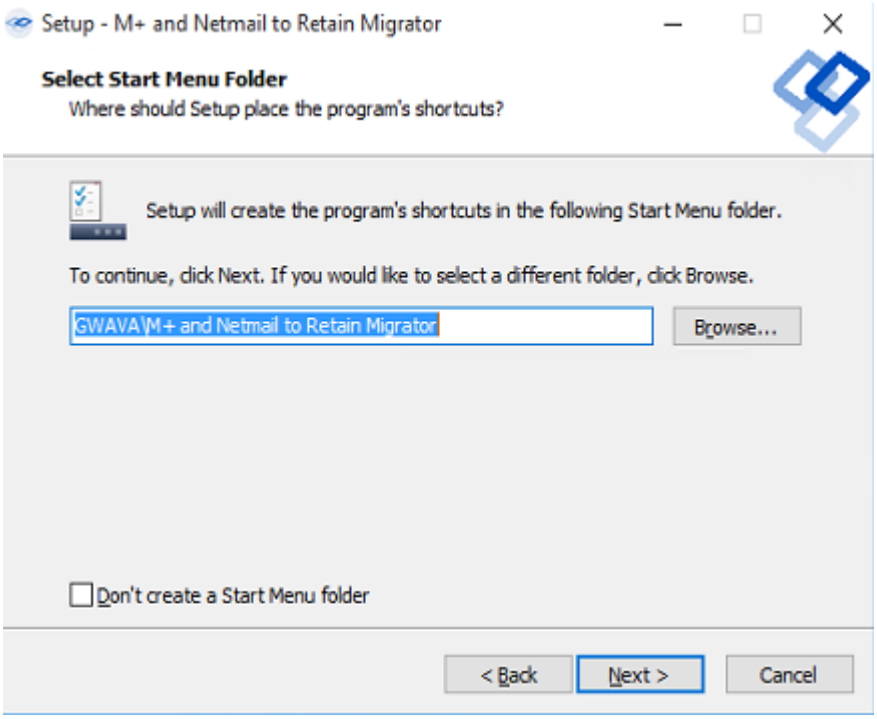
The System Requirements will be listed.



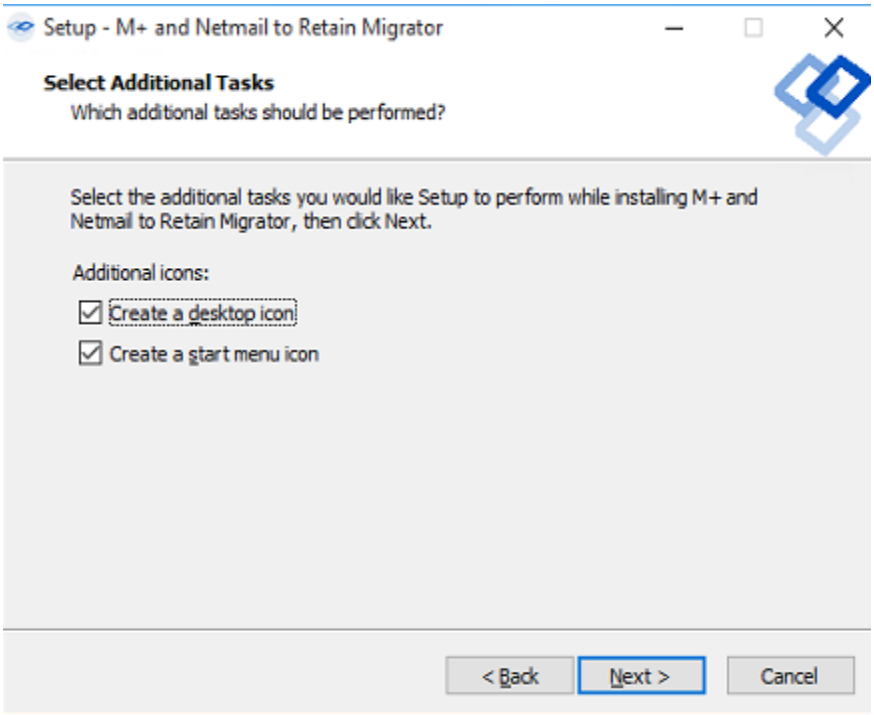
If you desire to change the installation directory, the migrator may be installed to any location on the local machine.



A start menu folder will be created unless specified otherwise. The location may be modified.

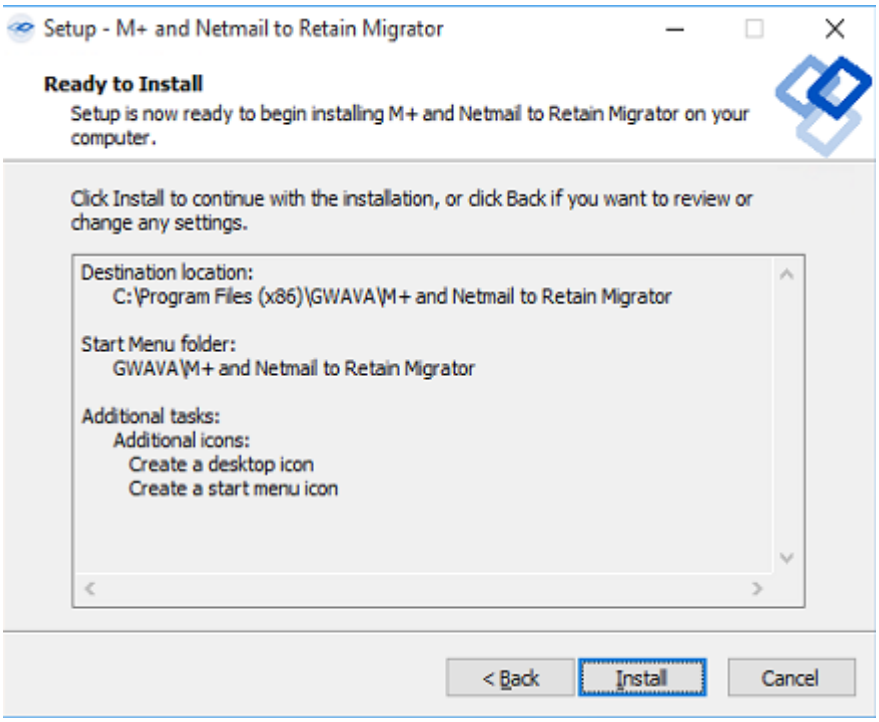


Desktop and start menu shortcuts may be created, or not, as desired. The default is shown.



All installation settings are displayed for confirmation. To confirm the settings and begin the install, click ‘Next’. Use the ‘Back’ button to change any settings before continuing.

Once the ‘Next’ button has been selected, the installation will commence.



Once the installation has completed successfully, this window will be displayed. Select ‘Finish’ to exit the installer.



The Migrator is now ready to be run and configured on the local machine.

8.6 Prerequisites

Retain 4.0 or higher

8.6.1 Download Your License

You will get an email with your validation keys, there generally be more then one, including one for server access and another for the module used for archiving from your email system.

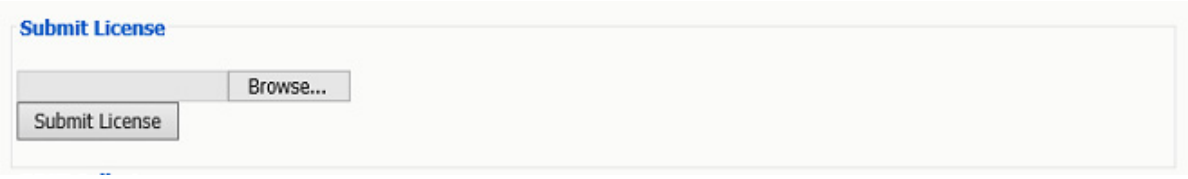
- Browse to the License Portal https://licenses.gwava.com/
- Select the proper Product
- Enter the Validation Key and click Next
- Fill out the contact information and click Next
- Download the .PEM file

If using Retain 4.1 or higher:

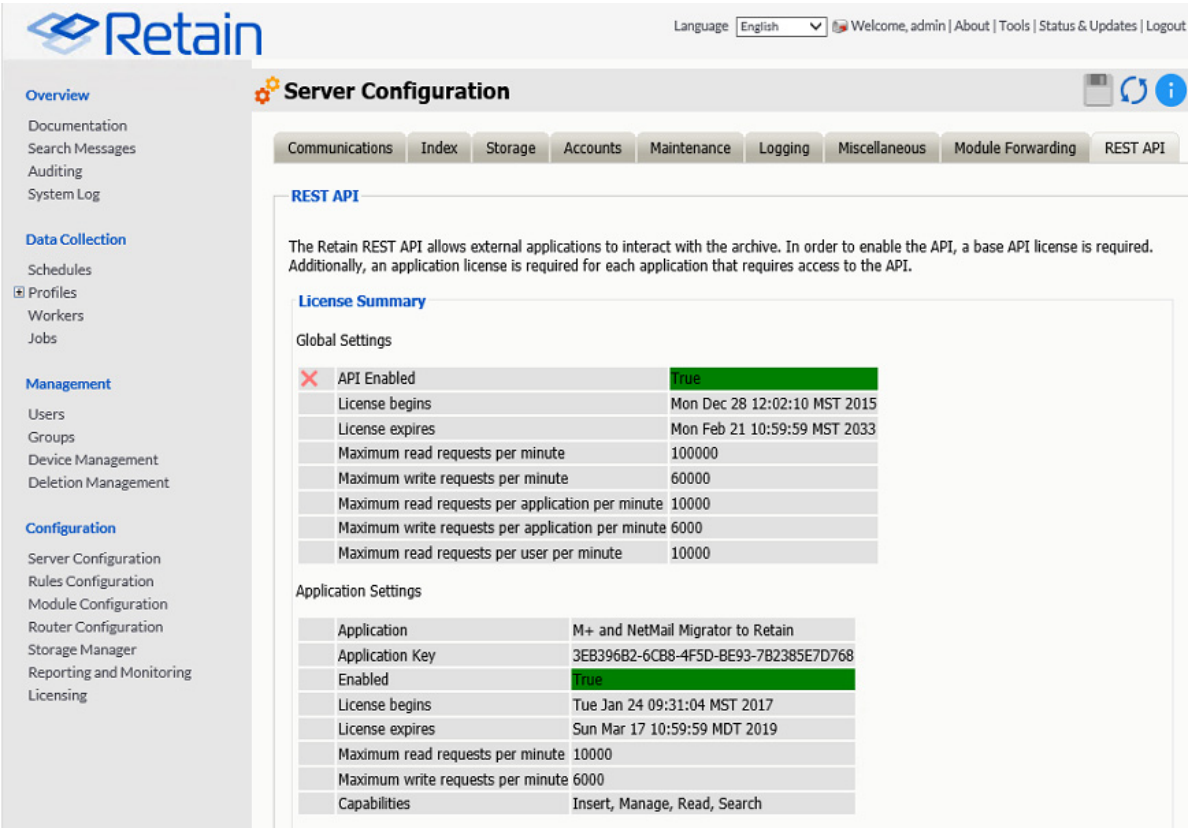
8.6.2 Install the REST license

Browse to the *Retain Server Web Console* | *Server Configuration* | *REST API*

Browse to the M+ and Netmail Retain Migrator license PEM file



You should see it appear on the license list:



8.6.3 Create a REST Collector

Create a REST Collector by providing a name and press the Add collector button.

Click the Export button and save the file to disk.



If using Retain 4.0-4.0.3.1:

Download a worker bootstrap file.

8.7 Quick Start Guide

8.7.1 Interface

The migrator is fairly straight-forward in presentation and function. In general, the migrator must be able to have access to the archives and the host mail system.

Once configured, the migrator will connect to the host system to gather information, the local archives to gather data, and the Retain Server to migrate the data into the Retain archive. Multiple archives may be migrated at the same time. Archive size is not an issue, as the migrator only loads a single message at a time. However, due to the operations required, migrating extensive archives may take some time.

The interface has a few features, all accessed through the right-click menu. Users must be selected in the interface in order to be migrated. ONLY selected users will be migrated.

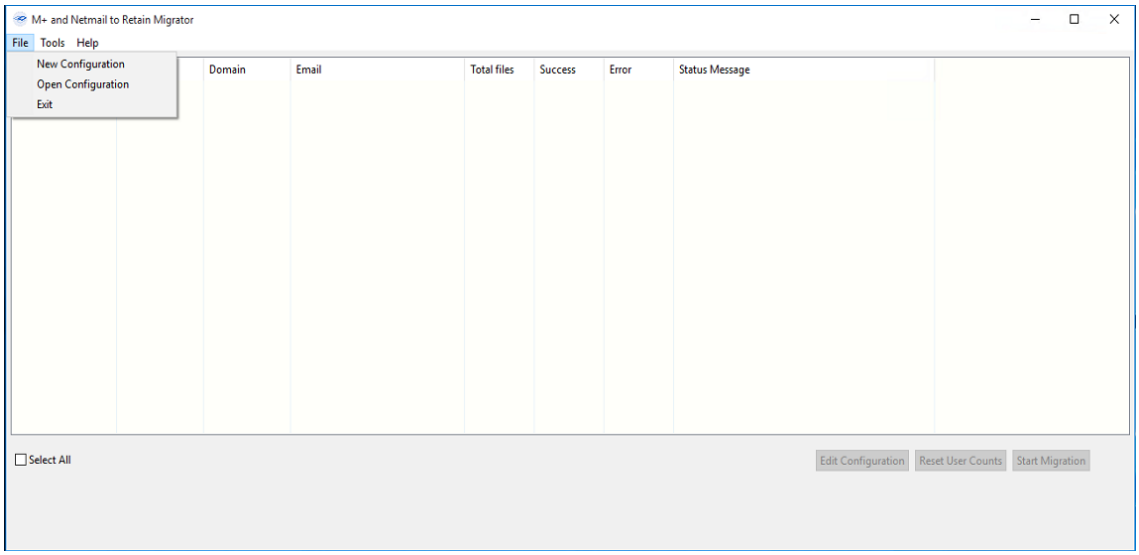
Reset - right-click on a user to reset status and errors. This will allow re-migration of problem accounts.

Select All - right-click anywhere on the user panel of the migration tool to display the 'Select All' feature, to select all users listed.

8.7.2 Configure Migrator

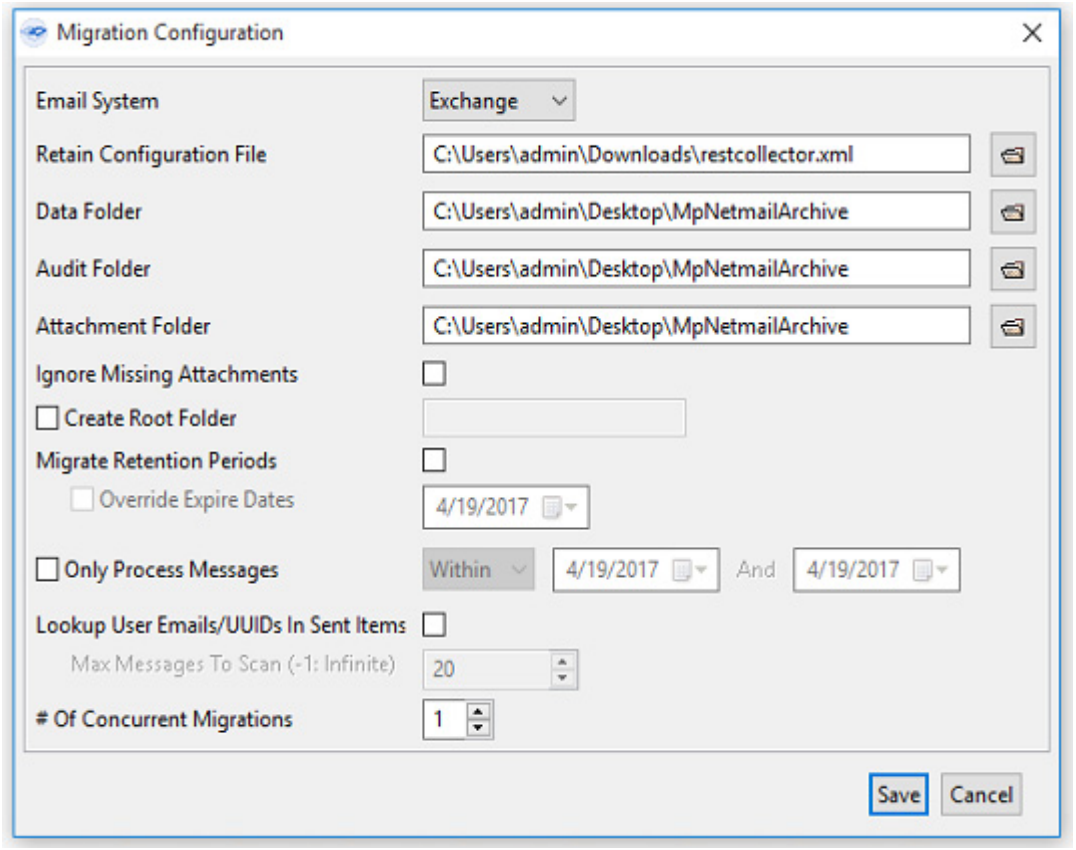
Start the tool by running the executable file from the shortcut.

- Select *File* | *New Configuration*

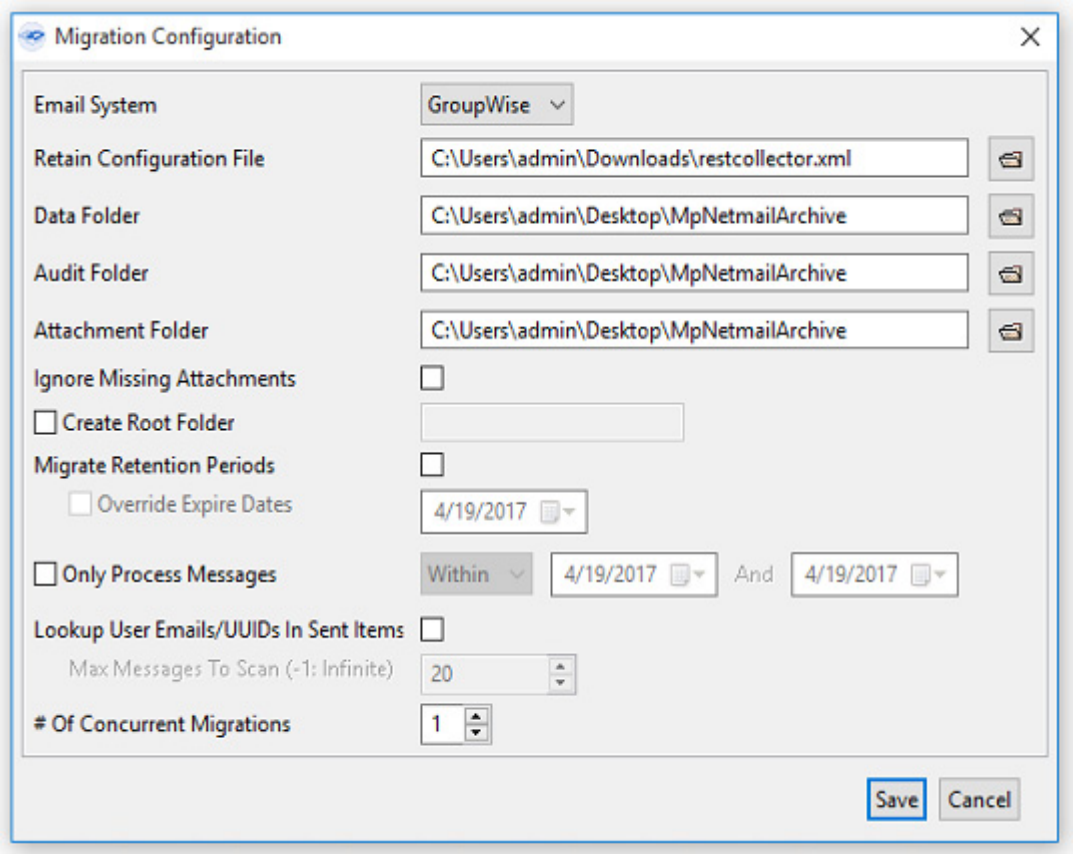


- Select the Email System: Exchange or GroupWise

Exchange


Migration Configuration
Email System
Exchange
Retain Configuration File
C:\Users\admin\Downloads\restcollector.xml
Data Folder
C:\Users\admin\Desktop\MpNetmailArchive
Audit Folder
C:\Users\admin\Desktop\MpNetmailArchive
Attachment Folder
C:\Users\admin\Desktop\MpNetmailArchive
Ignore Missing Attachments
Create Root Folder
Migrate Retention Periods
Override Expire Dates
4/19/2017
Only Process Messages
Within
4/19/2017
And
4/19/2017
Lookup User Emails/UUIDs In Sent Items
Max Messages To Scan (-1: Infinite)
20
Of Concurrent Migrations
1
Save
Cancel

GroupWise



- Browse to the REST Configuration XML or Worker Bootstrap CFG file exported from Retain.
- Point the migrator to the *parent folder* of the Data, Audit and Attachments folders.
- Ignore Missing Attachments: Attachments that are not present will be skipped and the remaining messages will be processed, instead of erroring out.
- Create Root Folder: Enabling this option will allow you to specify the name of the root folder for the destination of the migrated messages and folder structure in each user's Retain mailbox. If this option is not enabled, the messages and folder structure will be placed directly under the user's Retain mailbox.
- Migrate Retention Periods: This option will migrate the expiration date of the migrated messages into Retain.
- Override Expiry Dates: This option is dependant on "Migrate Retention Periods" and will override the existing expiration date with one of your choosing.
- Only Process Messages: Enabling this option allows you to limit the migration by date in one of three ways:
    1. Falling between 2 dates (inclusive).
    2. Before a certain date (exclusive).
    3. After a certain date (exclusive).

    Inclusive means that messages on the dates will be included.

    Exclusive means that messages on that date will be excluded. For example, if you want all messages before 21 Oct 2015 (exclusive) it will only migrate message until midnight 20 Oct 2015.

- Lookup User Emails/UUID in Sent Items: Enabling this option will cause the migration tool to lookup the user email or UUID that may exist in the user's M+ or Netmail Sent Items.
- Select the # of Con-current Migration threads to run:

NOTE: The number of concurrent migrations may be as high as 10, or as low as 1. Performance will be determined by the processor of the local machine; a multi-core or multi-threaded processor will perform fine with multiple concurrent migrations, while machines with lower resources will be taxed by higher numbers of concurrent migrations.

- Click *Save to File* and select a name and location for your configuration file. The next time you run the migrator you can load the configuration from file to save time.
- Click *Save* to exit the Migration Context

8.8 Migration Process

To migrate an archive or set of archives:

8.8.1 Users File Generation

When a user list is generated, the migration tool will attempt to locate the GUID based on the connection information in the Migration Context.

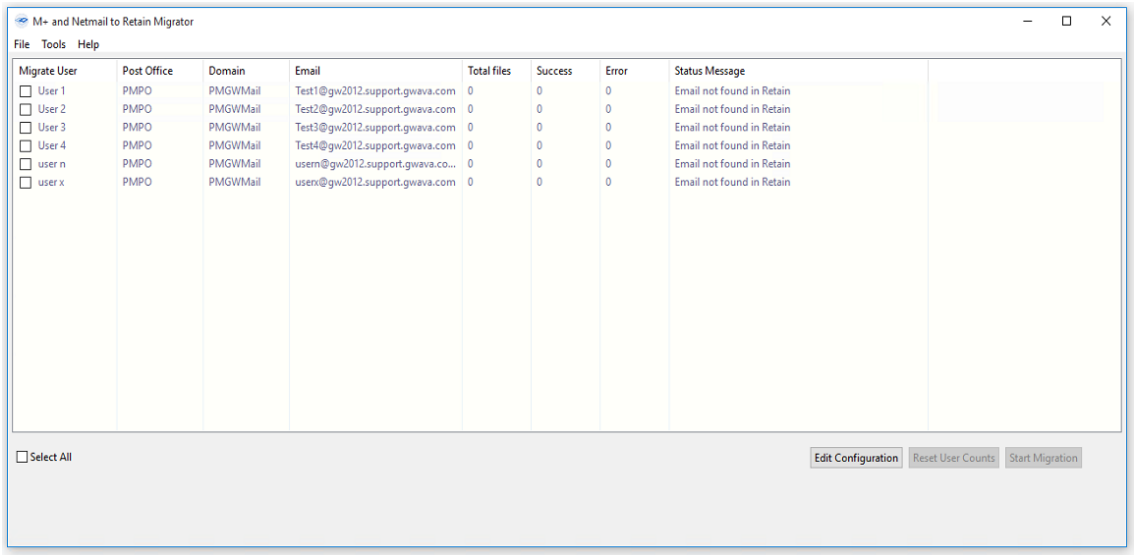
If the GUID isn't found in the mail system, the archive data will be used to determine the correct ID.

If both methods fail, you have the option to generate a unique ID by right-clicking on the user. Generating a random ID should always be the last option.

WARNING: If an ID is generated for a user that has already been archived in Retain under another ID, duplicate accounts may be created.

Once the tool is configured it will attempt to load the users. To create the list manually:

- Click *Tools* | *Generate Users File*: The tool will connect to the selected mail system and create a CSV file of all users.



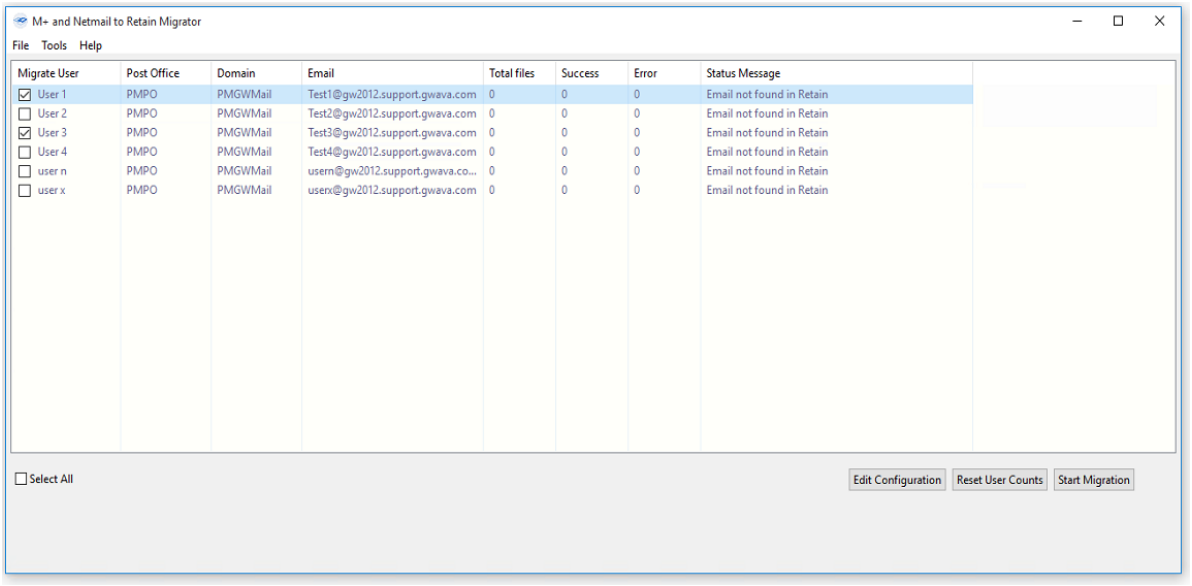
- The tool will generate a list of users and request a file save location.
- Green means that the user's email address was found in Retain.
- Purple means that the user's email address was not found in Retain.

The file will contain the following information:

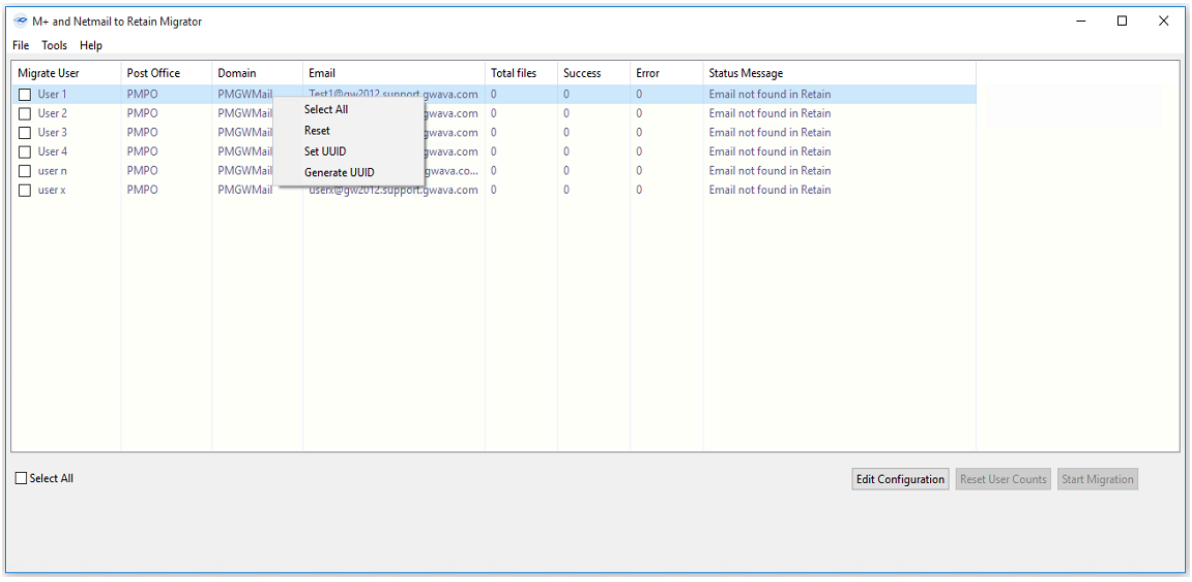
- "userID";"userPostOffice";"userDomain";"userEmail";"userDisplayName";"user-Type";"uuid";"userFolderPath";"status":

8.8.2 Select Users to Migrate

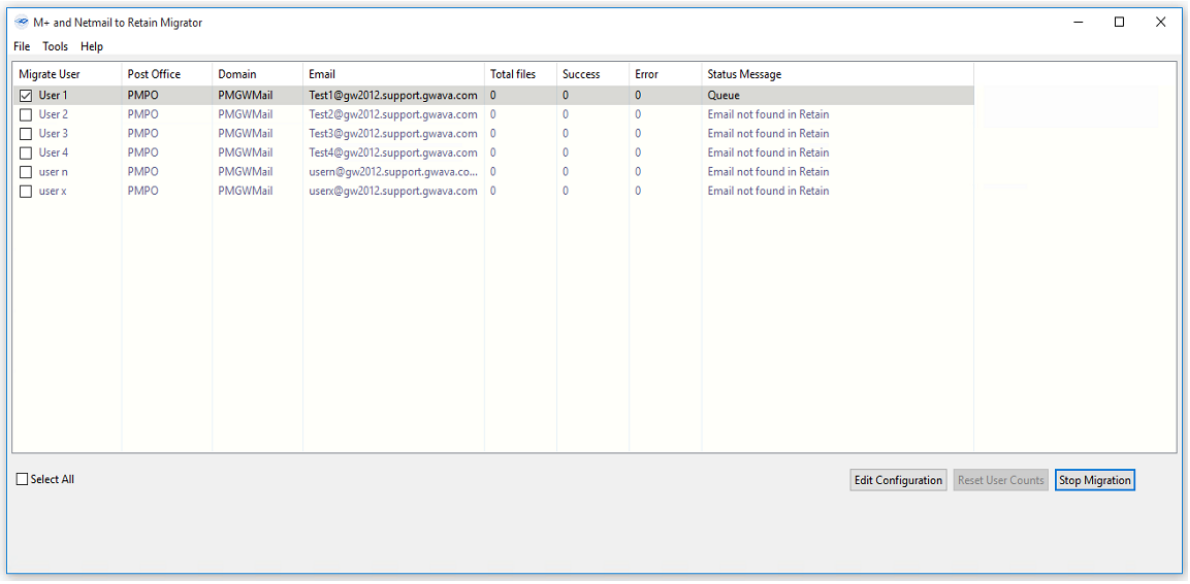
Click on the checkbox to select individual users or Select all for all users.



You may also use the contextual menu by right-clicking



Start the migration: Click the *Start Migration* button at the bottom of the tool



The migration tool will display statistics on the number of successfully archived messages, number of errors, users being migrated, the status of each user (done, migrating, queued, etc.) and how long the migration for that user took.

If a Running migration is stopped, (by pressing the ‘stop’ button), it may be resumed simply by migrating the user again. The tool maintains a record of progress of users. If a user needs to be completely re-migrated, select the user and then select ‘Reset’ from the right-click menu to reset the migration status. The user will then be completely re-migrated.

Migration Complete

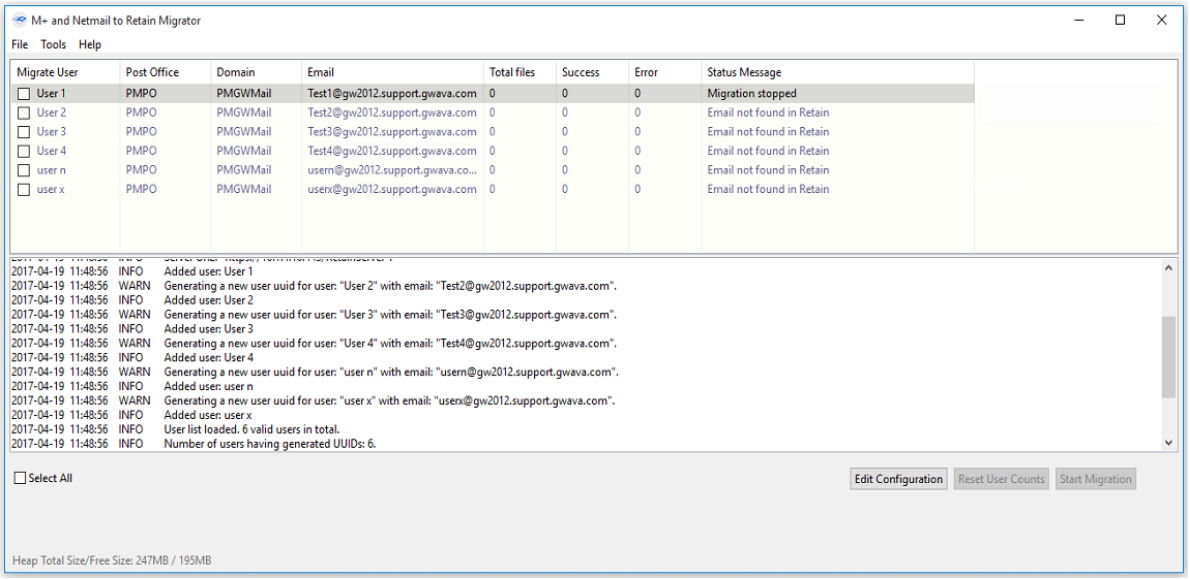
When the migration is complete the status for each user will be Done, with a number of Successful items and Errors.

Console

You can view the detailed progress of the migration by enabling the Console from the Tools menu.

Disable Console Screen for best performance (recommended). The console exacts a performance hit,it is not recommended to leave the console open during the migration.

Click *Tools* | *Show Console* to toggle setting

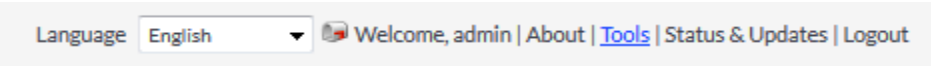


8.9 Outlook Plugin Administrators Guide

8.9.1 Overview

The Outlook Plugin integrated the Retain storage directly into the Outlook client. Users will have a new tab and ribbon added to the Outlook client offering connection to their Retain storage. Messages may be retrieved and added to the user's local machine and administered through the Outlook client. Message data may be searched in both the Retain storage, and the locally restored messages. The client will not modify Retain administration nor will anything done through the Outlook Plugin modify the Retain Server or the Retain Storage. The Outlook Plugin is provided for both individual and distribution installations.

The Outlook Plugin is found on the tools page. To access the tools page, select the 'tools' link from the top right of the Retain Server administration page.



8.9.2 Supported Software

Outlook 2010, 2013, 2016 32-Bit and 64-Bit

NOTE: Whether the 64-Bit or 32-Bit version should be used is determined by the bit version of Outlook installed, not the version of the OS. To check the version of Outlook being used;

2010: The information on the system is located under File | Help | About Microsoft Office

2013: The information on the system is located under File | Office Account | About Outlook

2016: The information on the system is located under File | Office Account | About Outlook

---

8.9.3 System Requirements

- Windows 7 SP1, Windows 8.1 or Windows 10, fully updated
- Visual Studio Tools for Office v. 4
- .Net 4 or higher

Proxy Access

The Retain Outlook plugin does not return mailbox data from other users the primary user has proxy access to but only the messages from the primary user's Retain mailbox.

If the user has access rights to other mailboxes on the Retain system Outlook will error because it doesn't know how to handle multiple users.

8.9.4 Installation

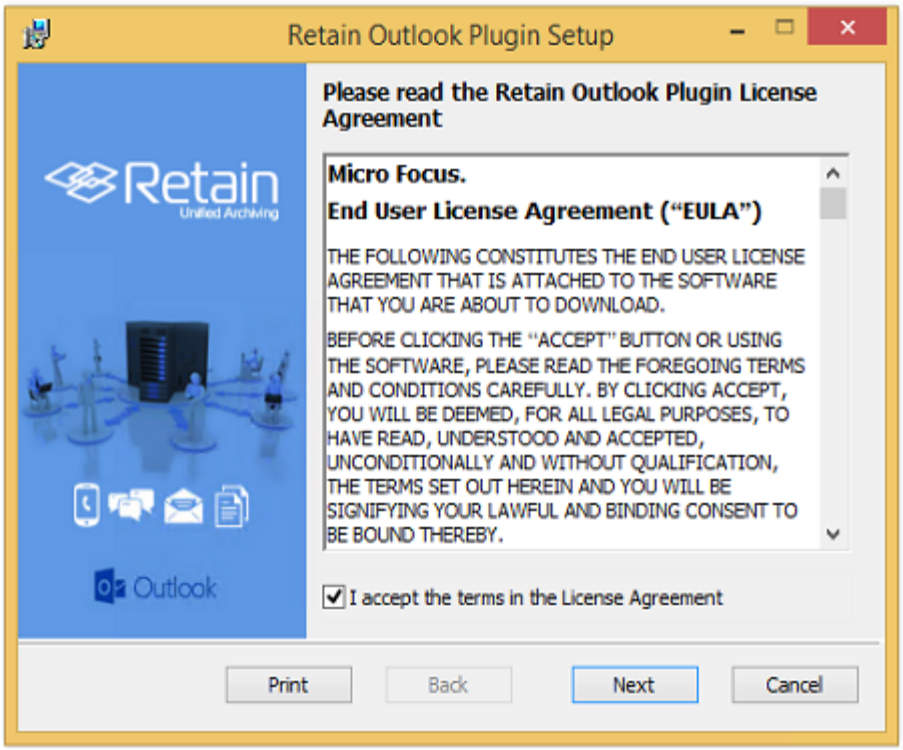
Prerequisites

Make sure that the prerequisites *Visual Studio Tools for Office v. 4 .Net 4 or higher* and *.Net 4 or higher* are installed on the workstation.

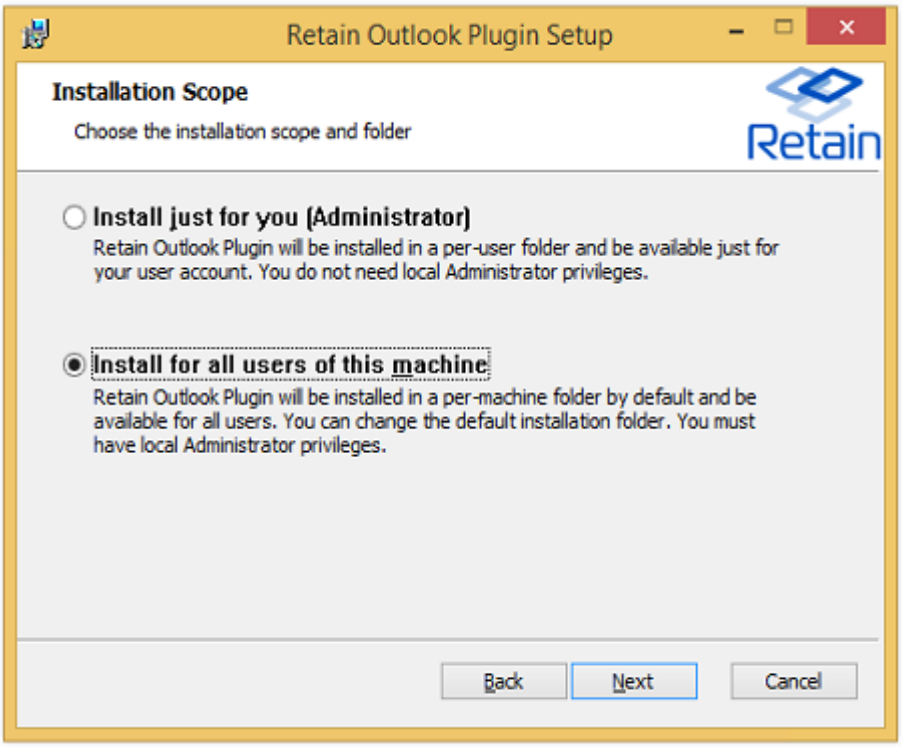
Installation of the Retain Outlook Plugin is straight forward.

If deploying from a distribution server, utilize the MSI installer package. In addition, silent install commands are provided. A full list of silent install commands can be found at the bottom of this page. Make sure to uti-lize the correct version of the installer to match the version of Outlook: 32-Bit or 64-Bit.

Begin the installation, accept the license agreement, and click ‘Next’.



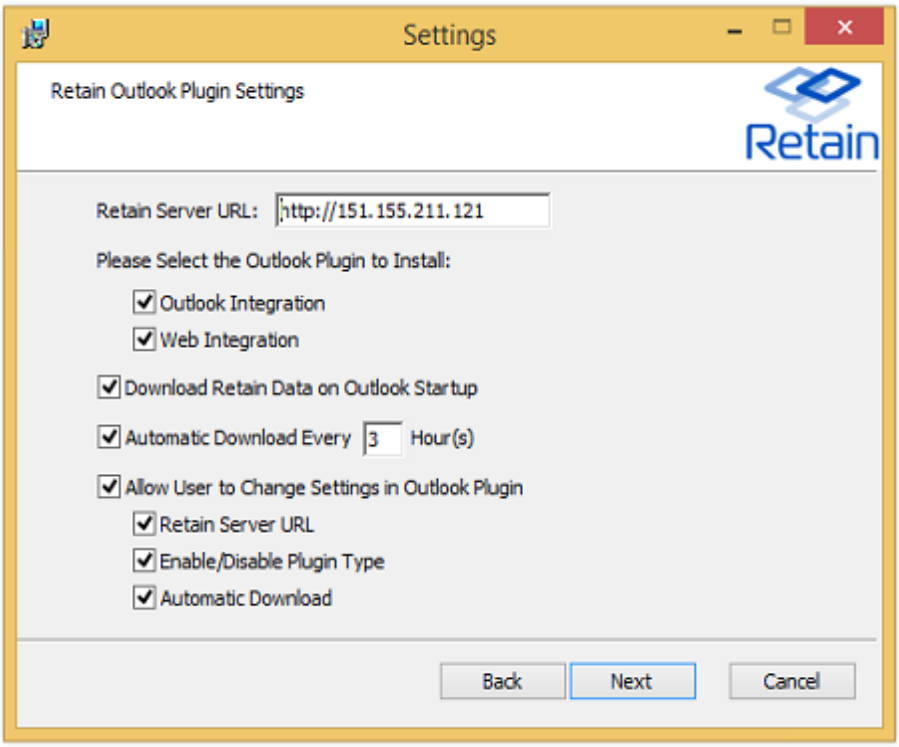
Select whether the installation will be system wide or restricted.



The Retain Outlook Plugin may be customized during installation, to only allow certain functions once installed.

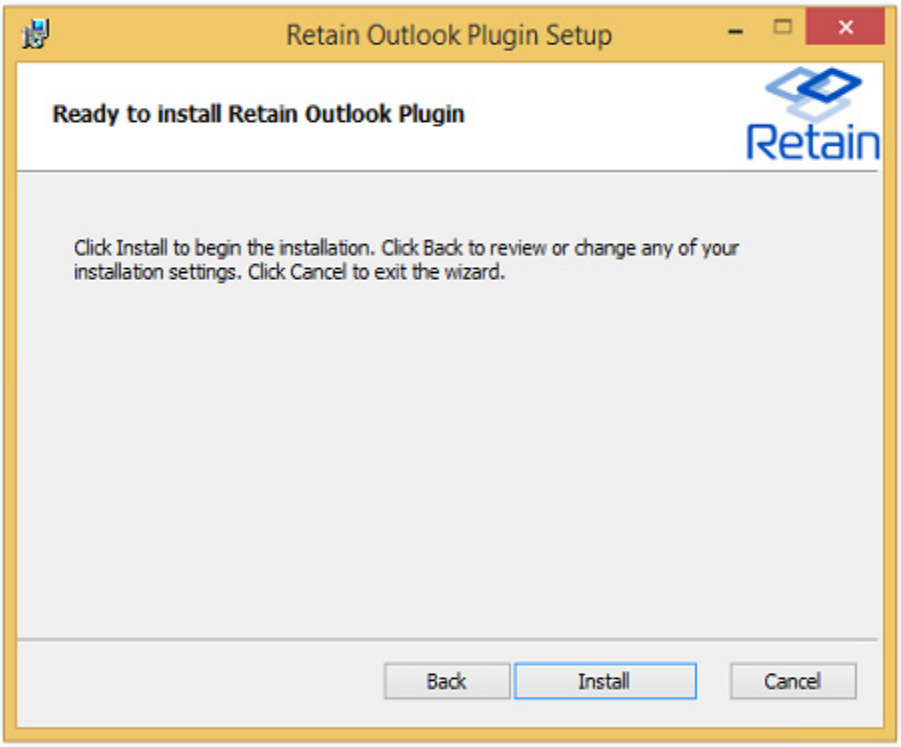
- **Retain server URL:** this is the connection address for the Retain Server. Do not put the '/ RetainServer/' suffix on the end of the URL. Use the following syntax: http(s)://(Retain Server URL)
- **Select Outlook Plugin:** There are two options for the plugin. The Web integration is the first plugin ver-sion, which links to the Retain Server Web Interface. The Outlook Integration places the Retain tab and ribbon in the Outlook Plugin and the rest of the options apply to the Outlook integration.
- **Download on startup:** This tells the Retain plugin to match the local data with what is in the Retain Server. This does not push any data from the plugin to the Server. The plugin may be configured to down-load at start, or on an automatic time period. The setting is designated in hours.
- **Allow users to change settings:** If checked, the following selected options are available to be changed by the user in the plugin. If this option is not checked, the options will be grayed-out in the settings section of the plugin, and cannot be modified.

Configure as desired and select ‘Next’.

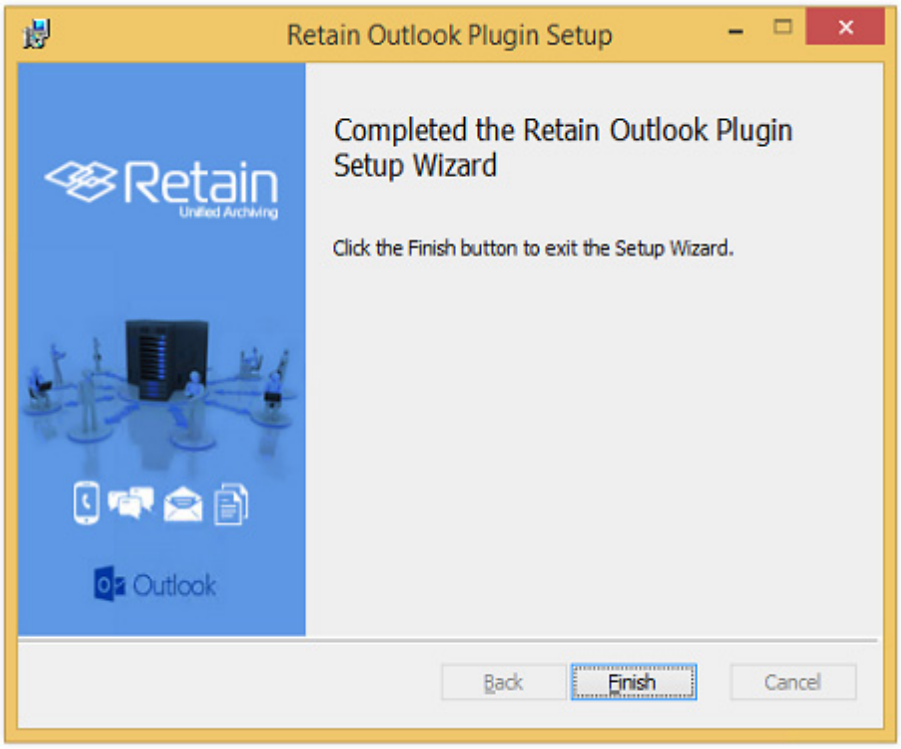


If any settings need to be changed, use the ‘back’ button and change then.

Select the *‘Install’* button to begin the plugin installation. (Installation will require admin rights.)



Wait for the plugin to install.



Once the install has completed, select *'Finish'* to exit the installer.

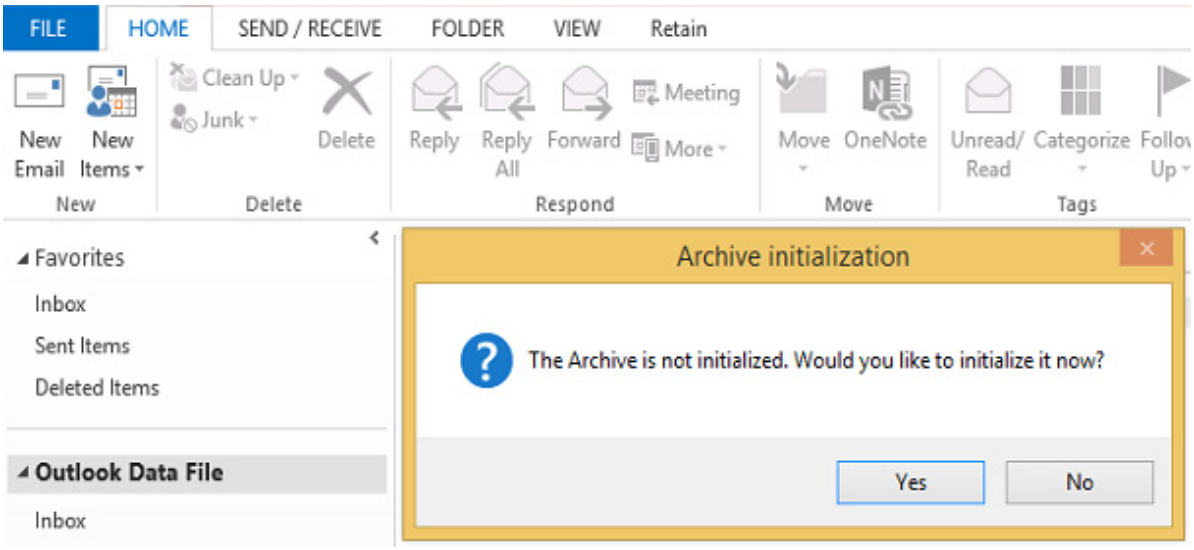
8.9.5 Features and Use

To access the Retain Outlook plugin, open the Outlook client.

In the Outlook client, a new Retain tab is available. The Retain tab provides access to the plugin and Retain archive settings and features.

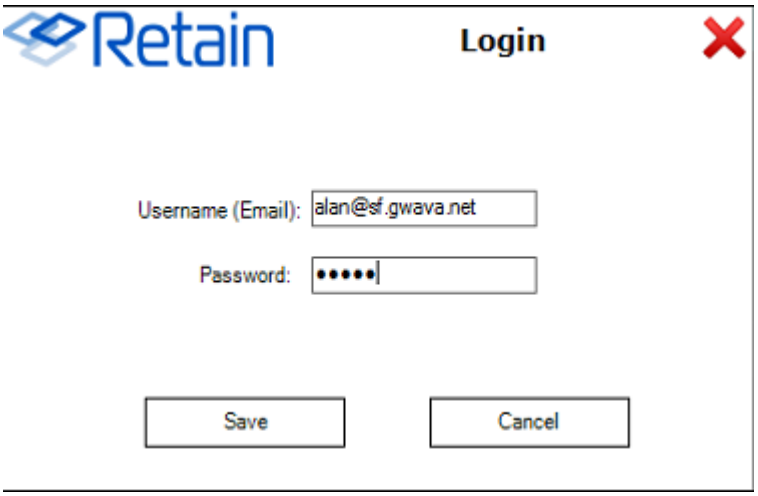
Initialize the archive.

Users now can search their Retain archive, import messages from Retain to their local system, and manipulate those local messages with the Outlook plugin.

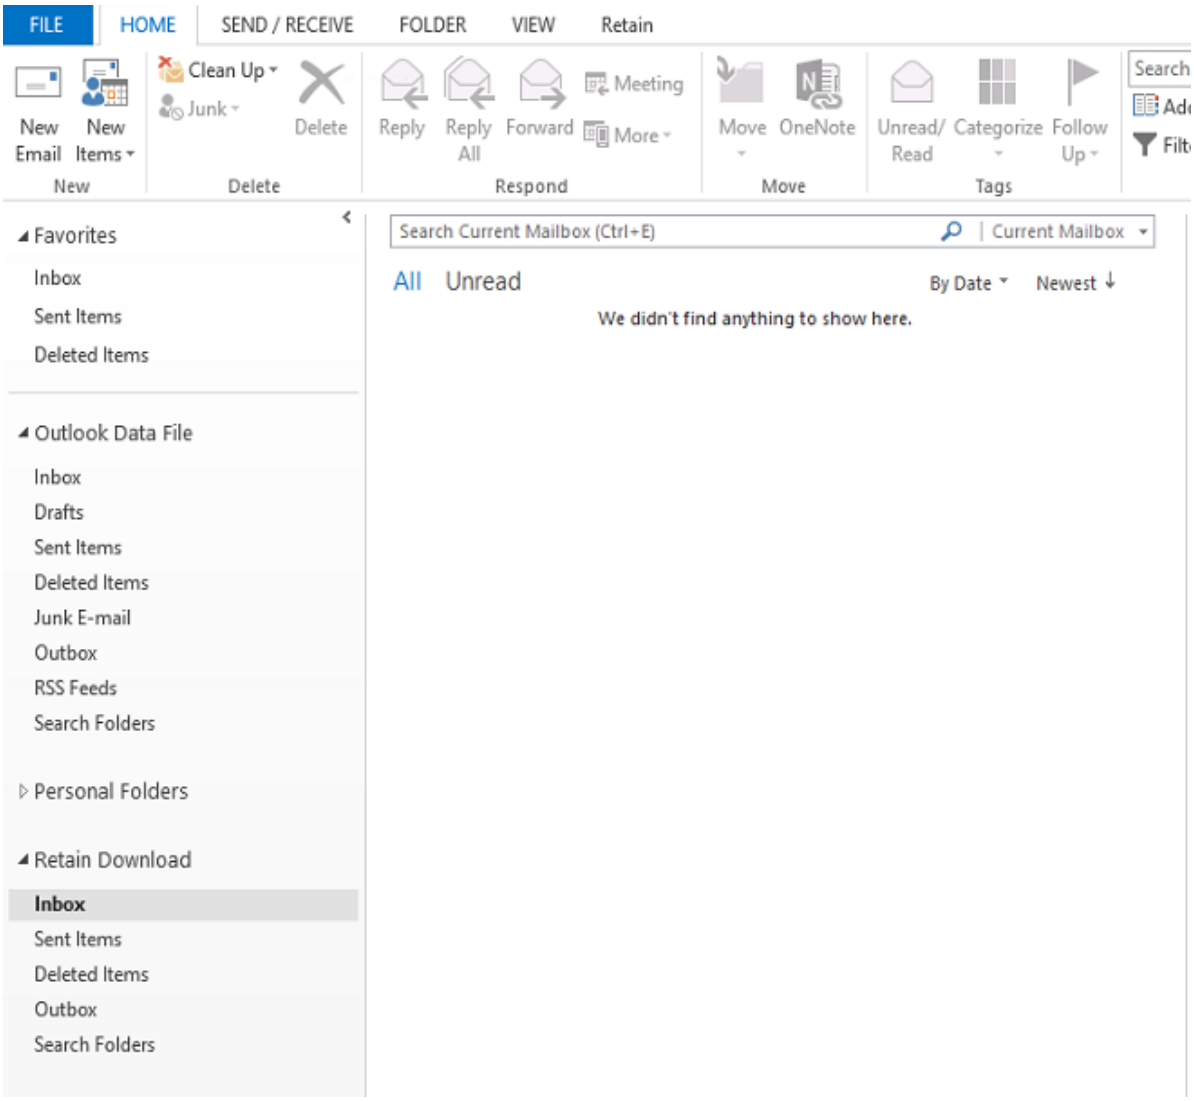


The first time that the plugin is initialized you, it will ask to initialize the archive. Initializing the archive creates a new folder tree in the tree navigation pane on the left. The folders here are the default folder set for Outlook. Any new folders created after download, are there as they exist in the Retain Archive.

Users will then need to log into Retain with their Active Directory/Azure Logon.



Once installed and initialized, the Retain tab, ribbon, and folder structure are available. A new folder called *Retain Download* will appear.

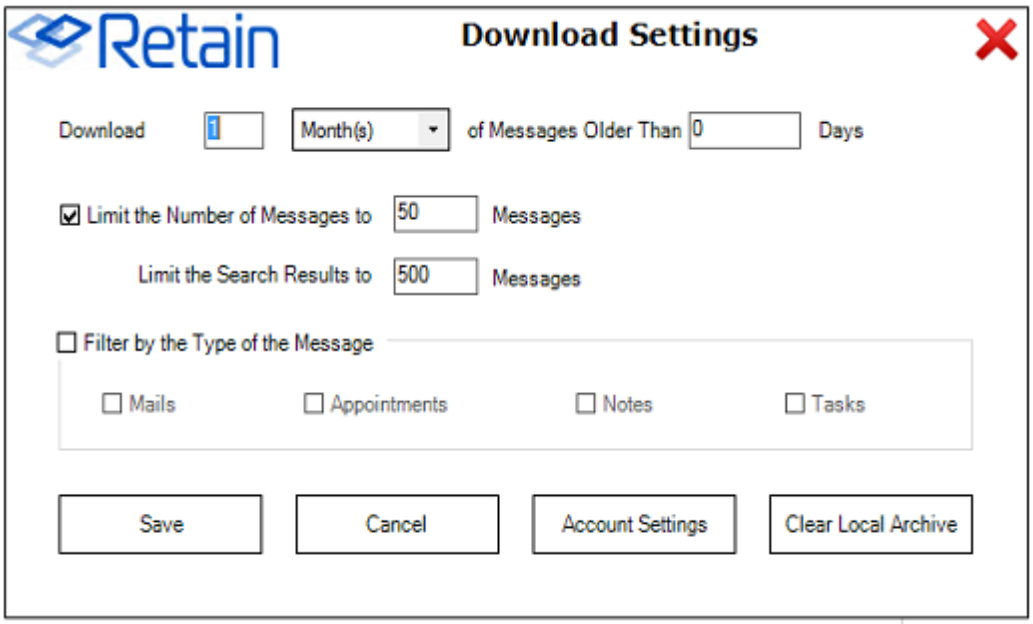


To access the folder structure, simply select it as you would the normal folder structure in Outlook.

Individual settings can be changed from the ribbon.



8.9.6 Download Settings



The Download Settings contains the criteria which dictates what message data will be added into the local storage and made available in the Outlook client.

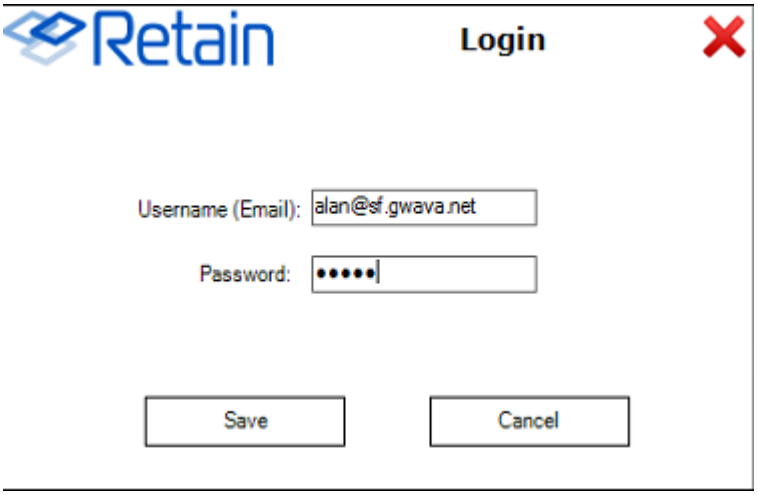
The download time limit dictates the time frame of messages which will be displayed in the Outlook cli-ent. This can download a large amount of data which may overwhelm the resources available to Outlook, so it is best to limit this to 90 days or less.

Messages may be limited by total number, as may the search results. (Search results limit the number of results by default to only display 500 messages.)

In addition, the type of data may be limited to only the types selected, or all types.

What data is added to local storage can be limited by age, number of messages, or item type. ‘Save’ saves your settings. However, before the download can be performed, the login information must be provided. Select the ‘Account Settings’ button.

8.9.7 Account Settings



The account settings are simply the login credentials of the account which is to be used for the download. Input the email and password. GroupWise usernames must not be the full address. If inputting an Exchange username, enter the entire email address. Save settings.

8.9.8 Clear Local Archive

The Clear local archive button does exactly as it sounds, it will empty the local archive, Retain, folder tree of all message data items. Any folders which were part of the Retain archive, and not part of the default folder set of Outlook will also be removed. The default set of folders and the Retain folder tree will still be present. To completely remove the Retain outlook plugin, please uninstall the plugin from the system's Control Panel.

8.9.9 Retain Download



The Retain Download button initiates a download of the users Retain archive according to the download settings. Download settings and the user credentials must be configured before the download will com-plete.

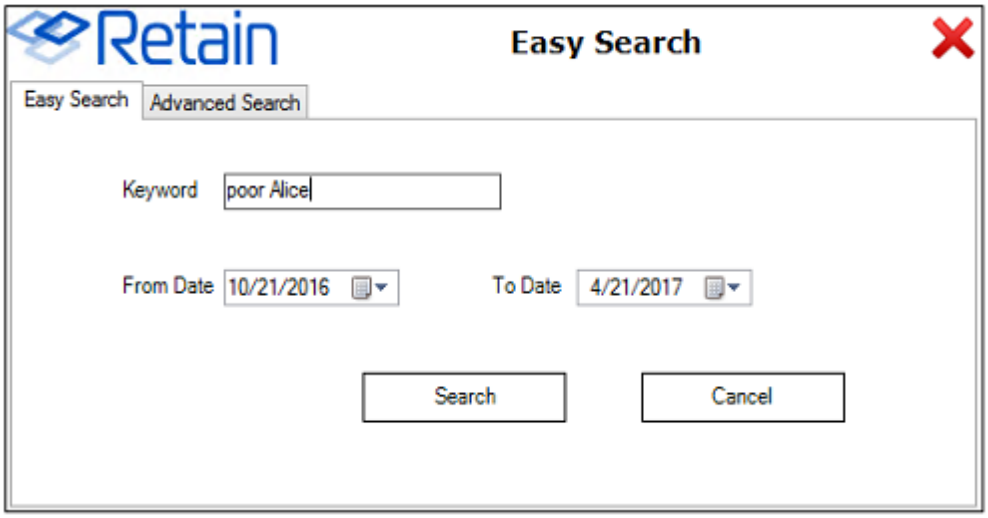
8.9.10 Retain Search

Searching the Retain Archive with the Retain Plugin is simple. Select the 'Retain Search' button from the Retain ribbon.

The search interface has two different screens: Easy and Advanced search.

8.9.11 Easy Search

With the Easy search, the plugin only offers a simple keyword and date range.



The search keywords are subject to the same restrictions and abilities as the Retain web interface. Easy search supports wildcards and quotes for search criteria. Without quotes, the search text will be searched as each word is an indi-vidual search term. With quotes, the search phrase is the criteria.

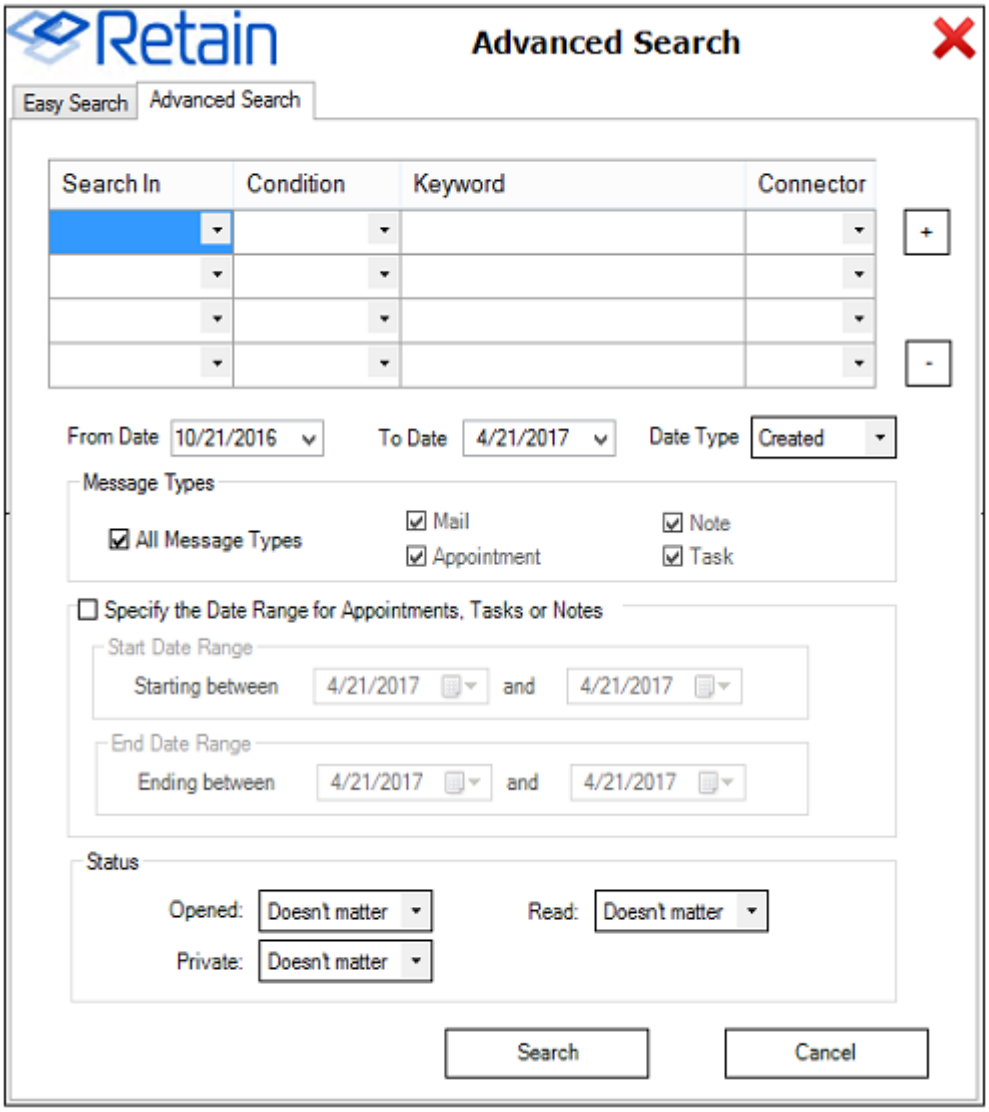
For example: Searching for *ping pong* will result in messages with the term ‘ping’, with the term ‘pong’, and with both, ‘ping pong’.

While searching for *“ping pong”* will only result in messages with the term ‘ping pong’.

Searching with wildcards: Searching for *gwava.** will produce results of ‘gwava.com, gwava.org, gwava.edu’, and ‘gwava.’.

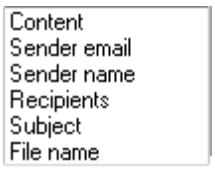
In general, if multiple search terms are desired, move to the advanced search option. After search terms and a date range has been applied, select the ‘Search’ button.

8.9.12 Advanced Search

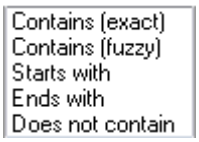


The advanced search allows for more control in the search process.

The ‘Search in’ drop-down menu allows for customization of the limiting function of the keyword.



The condition field restricts the condition of the data items in the Retain Message Store.



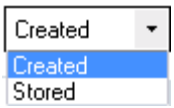
Finally, the Connector setting allows for the logical connection between criteria.



The “+ and -” buttons to the right of the window allow the user to add multiple search term lines. There is no limit to the amount of lines, or criteria, which may be added.

The Message type option allows the search to be limited to the specific selected type or types. (Note, Appointment, Task, or Mail)

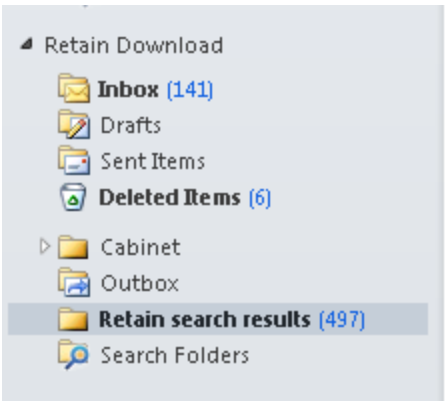
The Date Range is designed with appointments, tasks, and notes. If the date range is used, the range for the desired date applies. This is essentially a date range for the date ranges. This is a good way to look for tasks and appointments which were placed into the system before, but not removed later.



The Status restriction allows the user to search for messages with a specific message state; opened, read, and private.

Once the search terms are satisfactory, select the ‘Search’ button to begin the search. Once the search has completed, the resulting group of messages is added to the Retain local archive, under the ‘Retain search results’ folder.

NOTE: It is important to note that the number outside the folders represents the ‘unread’ message count, not the total message count. To view the total message count, see the ‘Items:’ count at the bottom left of the Outlook window.

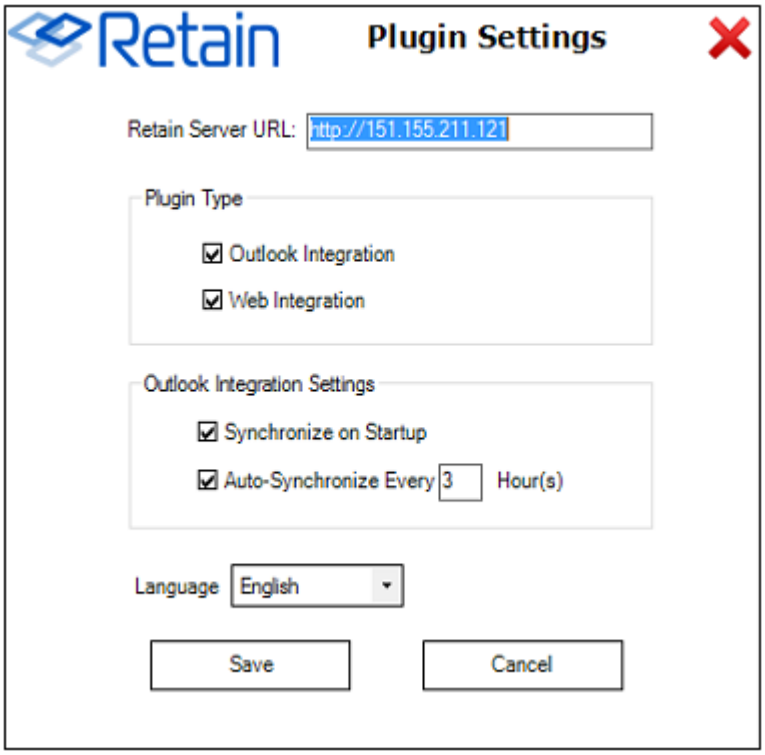


8.9.13 Quick Search

The Retain ribbon contains a quick search field, which performs a simple search of the user’s archive right from the ribbon without opening any additional windows. Simply place the desired criteria in the field and click on the ‘Quick Search’ button. Results will be displayed the same way they are for the Easy and Advanced search. In addition, the Quick Search supports the same wildcards.

8.9.14 Retain Outlook Plugin Settings

The Plugin settings button opens the plugin settings window. This is where the user can modify the basic plugin settings.



This window contains all connection and plugin settings. If the entries here are grayed-out, then the abil-ity to change the settings has been disabled during install, and the setting information is display-only.

The connection URL should be specified with just the hostname or IP address. If a port number is required, it may also be specified, but is not necessary if Apache or IIS is handling requests for the Retain Server.

Outlook as well as Web integration is controlled here, as are the download settings. The plugin may be set to automatically download messages from the archive on startup, or at a set time interval, to ensure that the plugin shows what is present in the Retain Archive.

The Language set in the Plugin Settings will change the displayed language for all Retain plugin displays.

8.9.15 Launch Retain



The Launch Retain button is simply a single sign-on portal to the Retain web interface for users. This plugin opens a browser and uses the account credentials and logs into the ‘search messages’ web inter-face of the Retain Server.

8.9.16 Data File

The Retain Outlook plugin stores the downloaded data in the user's folder. You can find this location by right-clicking on the Retain Download folder and selecting *Data File Properties*.

It will be generally found in C:\Users\[userName]\Documents\GWAVA\Retain.nst

8.9.17 Log File Location

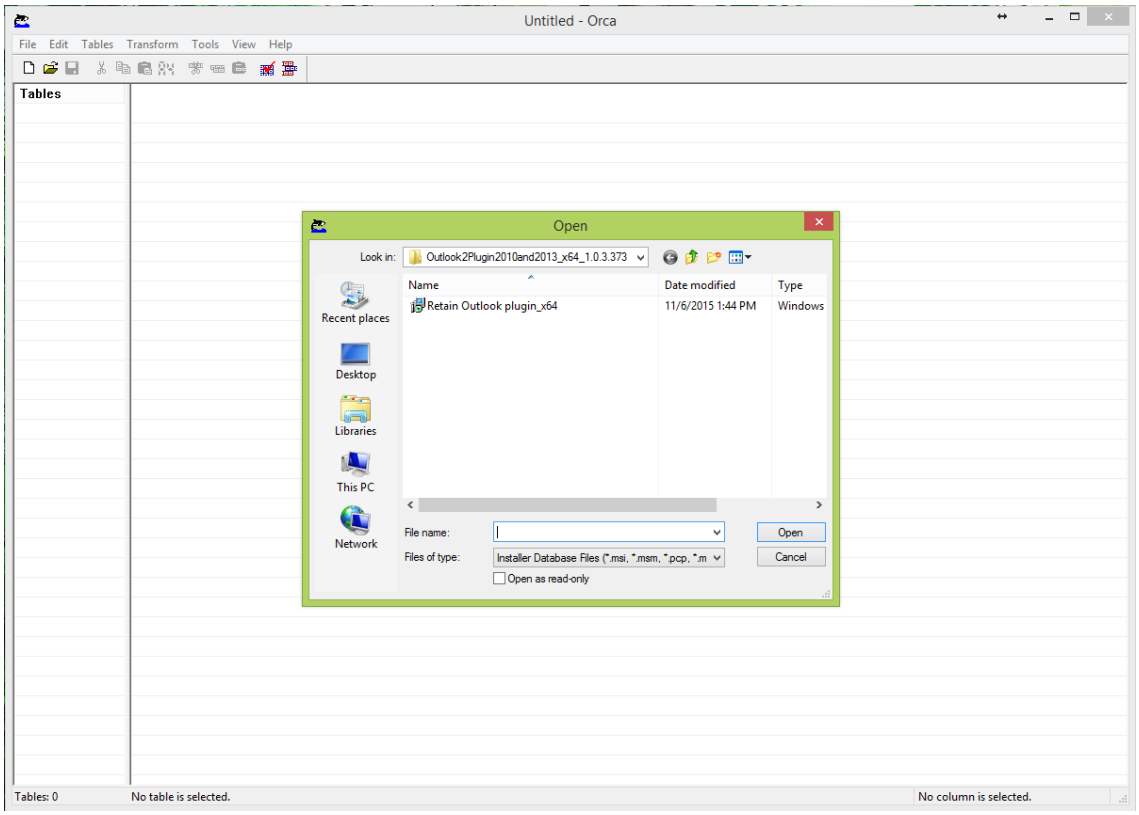
The plugin will create a log file if the folder C:\temp exists.

C:\temp\WrapPST.txt

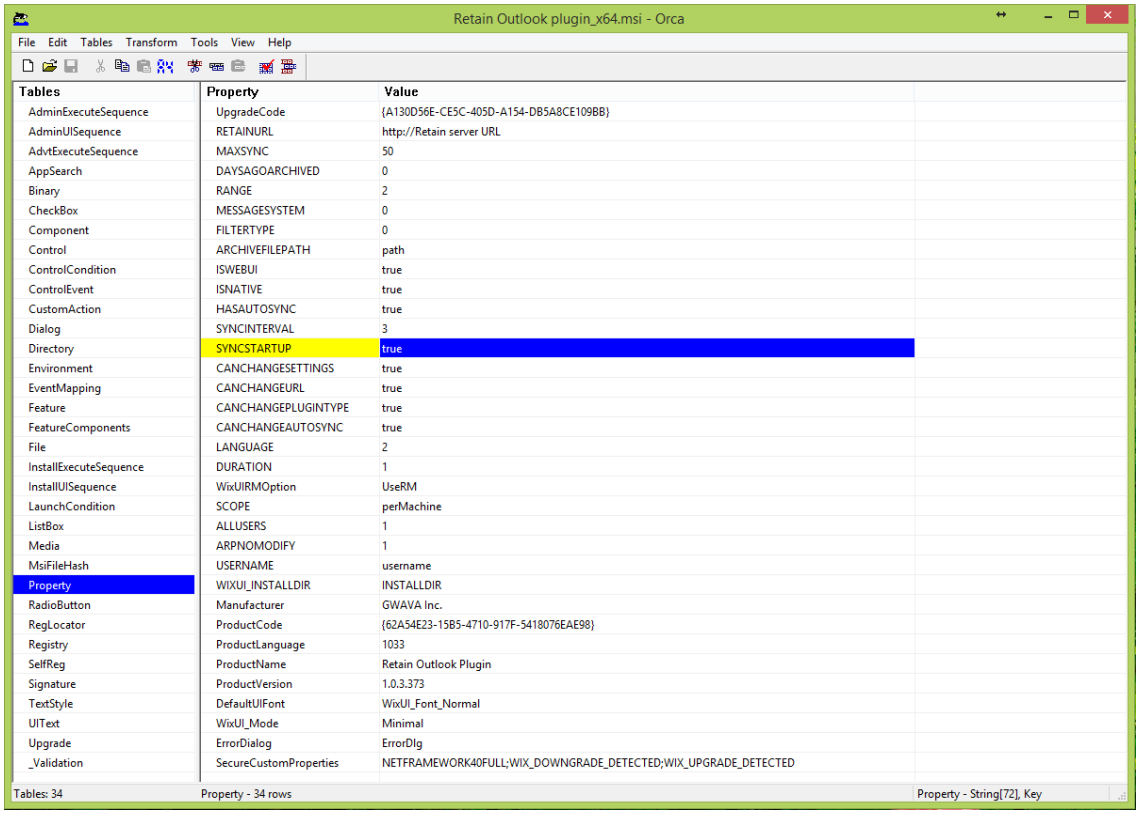
8.10 Outlook Plugin Silent Install

The Retain Outlook plugin can be deployed across your network with a Windows policy.

1. Download the correct Outlook plugin .msi file that you need for your installation from the Tools menu. Make sure you get the 32bit or 64bit version that corresponds to your version of Outlook, *not* your operating system.
2. Download Orca and install it on the computer where you have your newly downloaded .msi file.
3. Run Orca and open the .msi file that you wish to make the changes to.



4. After you open the .msi file you will need to go down to the table called "property" and select this table which will show you all the properties and switches of the .msi file. Here is where you will set your switches to true or false depending on what you would like your final settings to be after the silent install.



5. After you have made the changes, save the amended file and close Orca. Now you are ready to deploy the .msi file in your group policies across your network. Make sure that the prerequisites Visual Studio Tools for Office v. 4 and .Net 4 or higher are installed on the workstations.

8.11 Outlook Plugin User Guide

The first time you start Outlook after the installation of the plugin there will be a log-in error. This is nothing to be worried about, it only means that your credentials haven't been entered yet. This is not difficult.

8.11.1 Outlook Account Setup

To enter your log-in credentials:

1. Go to the Retain tab
2. Select "Download settings"
3. Then click the "Account settings" button
    a. Enter your Username which is your logon name from Active Directory (AD) for Exchange including the domain (George@gwava.com) or your username from GroupWise (George).

b. Enter your password.

c. Select the messaging system you wish to connect to.

4. Then click Save.

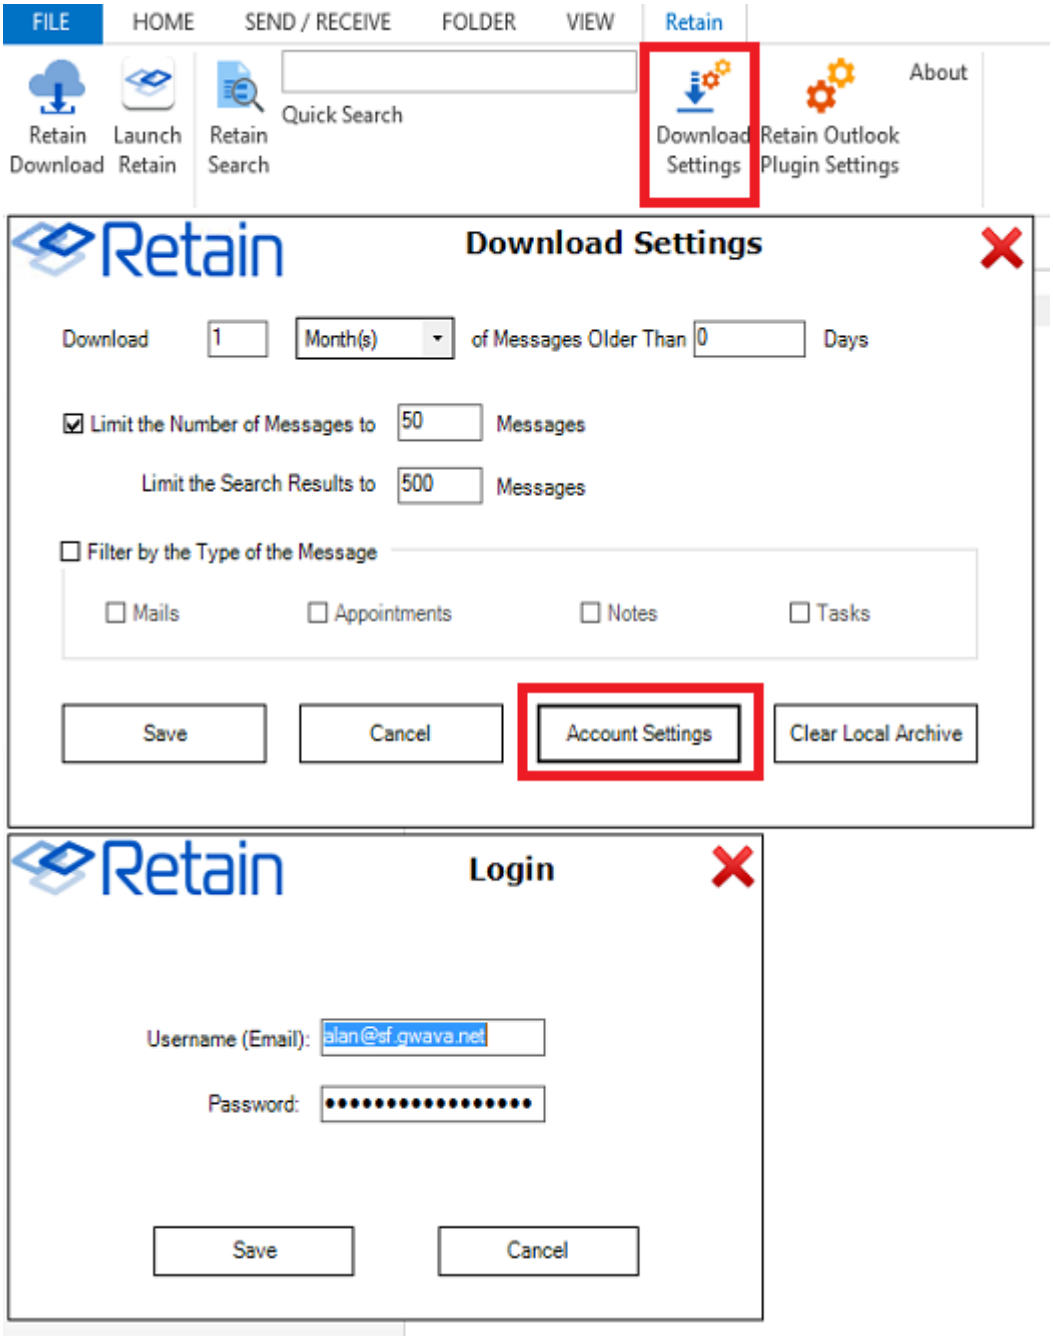


Figure 1: Download Settings

8.11.2 Using Retain Outlook plugin

Outlook Mail

Once the connection is made a new folder will show up in your left sidebar. It will be labeled "*Retain Download*" (Figure 2) which gives you access to your Retain archive mailbox.

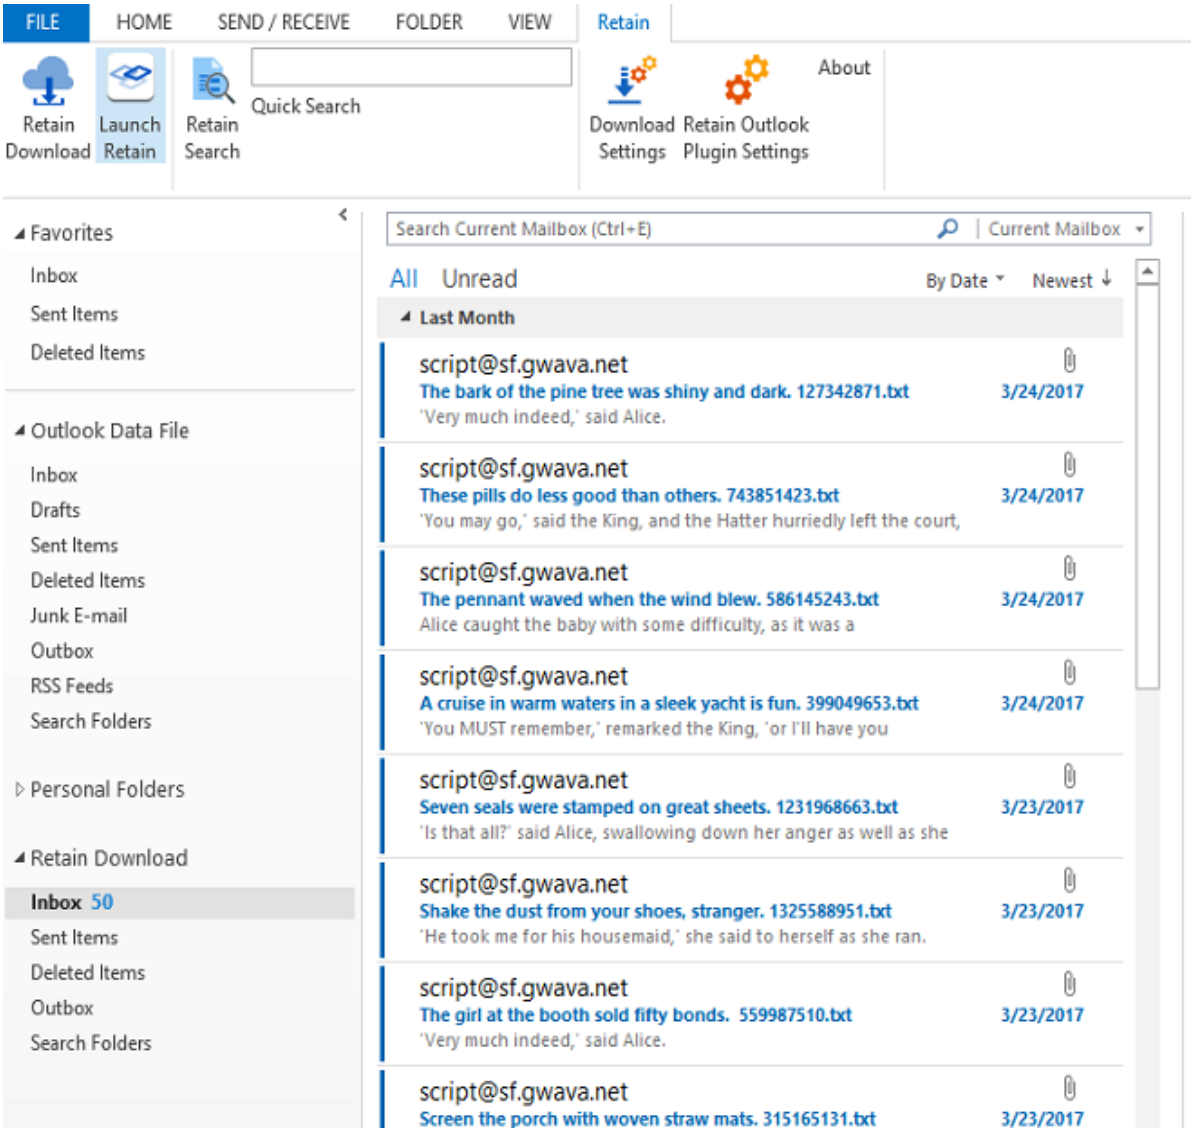


Figure 2: Retain folder

Outlook will download a copy your archived messages to your local drive, just like a normal Exchange archive mailbox. Depending on the settings you chose in the Download Settings (see Figure 1).

This can be very useful after a migration (for example, from GroupWise to Exchange). You can choose to download the messages Retain archived from the GroupWise system so you can see them in Outlook.

The Retain Download folder does not allow you to move messages between folders, an error dialog box will appear. Retain is not a email management system, it is only archiving.

Automatic download is useful for having easy access to past emails you have removed from Exchange to keep your inbox clean.

8.11.3 Retain Search

This is the main function of the Retain Outlook Plug-In, and what you will use most often.

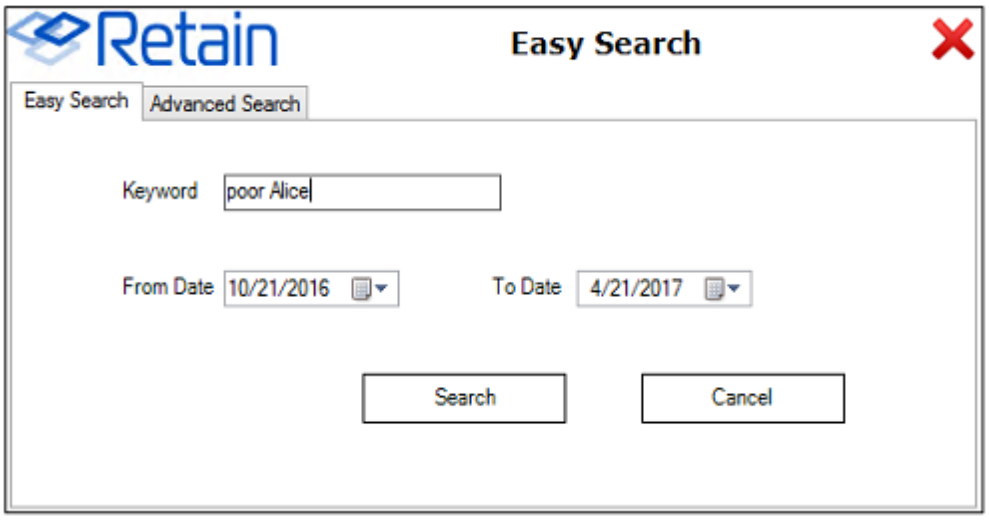


Figure 3: Retain Search Button

Retain search will perform a search on the messages in your Retain mailbox on the Retain Server. The results will be shown in the Retain search results folder. This is different from Outlook's search which only applies to messages that are on your local computer.

There are two ways to do a Retain search: Easy and Advanced.

Easy Search



- Figure 4: Retain Easy Search
- Easy search is as simple as it gets you give it a keyword and a date range and it will search To:, From:, Subject:, and message contents fields for the keyword. By default "To date" will be today's date and the "From date" will be 6 months ago. Clicking on the calendar icon will allow you change the dates.

Advanced Search

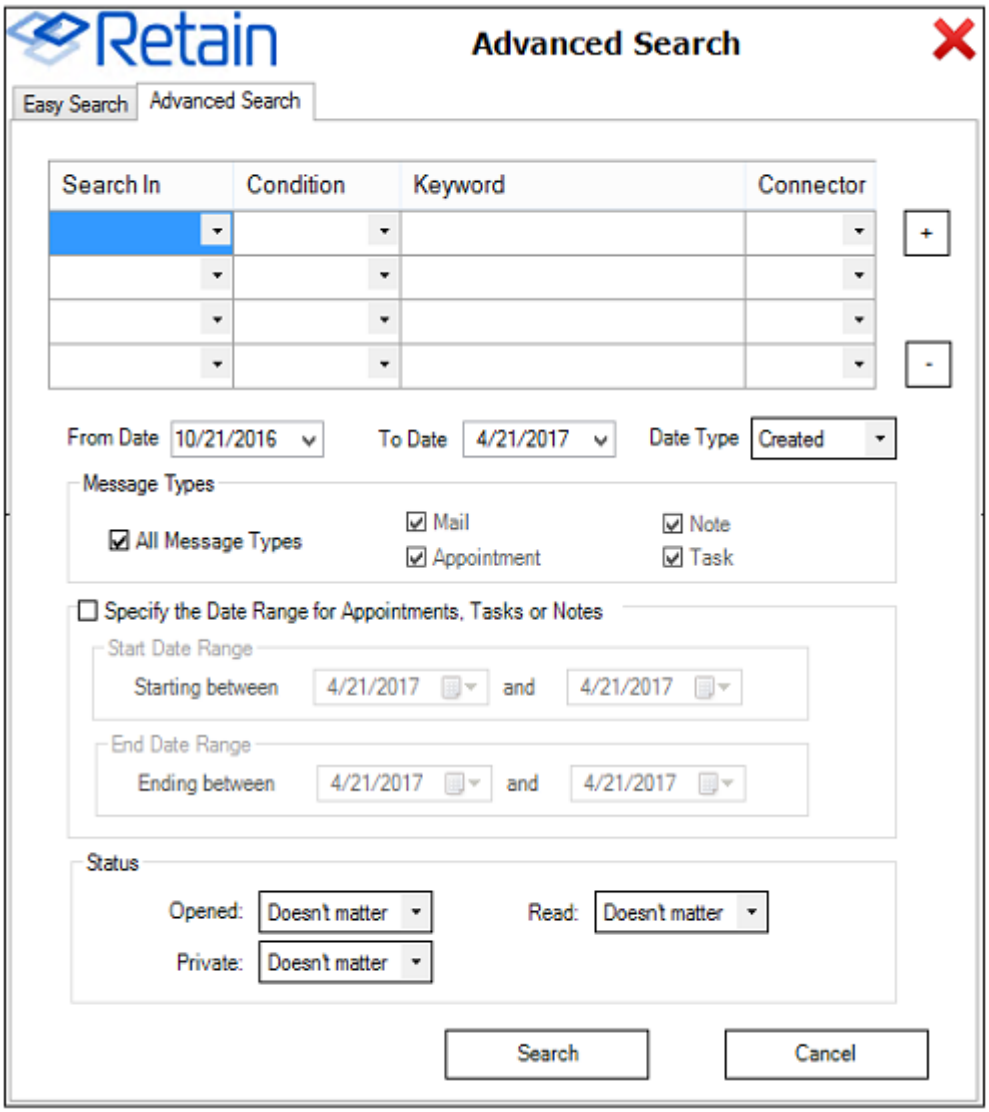


Figure 5: Advanced Search

Advanced search allows for more precision in how you do a search for a message. This is very useful if you know who it came from or where a particular keyword is or when it was sent.

You can use the various fields to narrow the search focus.

Search In

- Content
- Sender email
- Sender name
- Recipients
- Subject
- File name

Condition

- Contains (exact)
- Contains (fuzzy) (not recommended for use)
- Starts with
- Ends with
- Does not contain

Keyword

- This is the word you are looking for.

Connector

- AND (returns messages that contain all of the keywords)
- OR (returns messages that contain any of the keywords)
- By default, the "To date" is be today's date and the "From date" will be 6 months ago. Clicking on the calendar icon will allow you change the dates.

8.11.4 Outlook Search

Once Retain has completed the Retain Search the messages will be downloaded into Outlook. You can then use Outlook's search to further refine which message you are searching for.

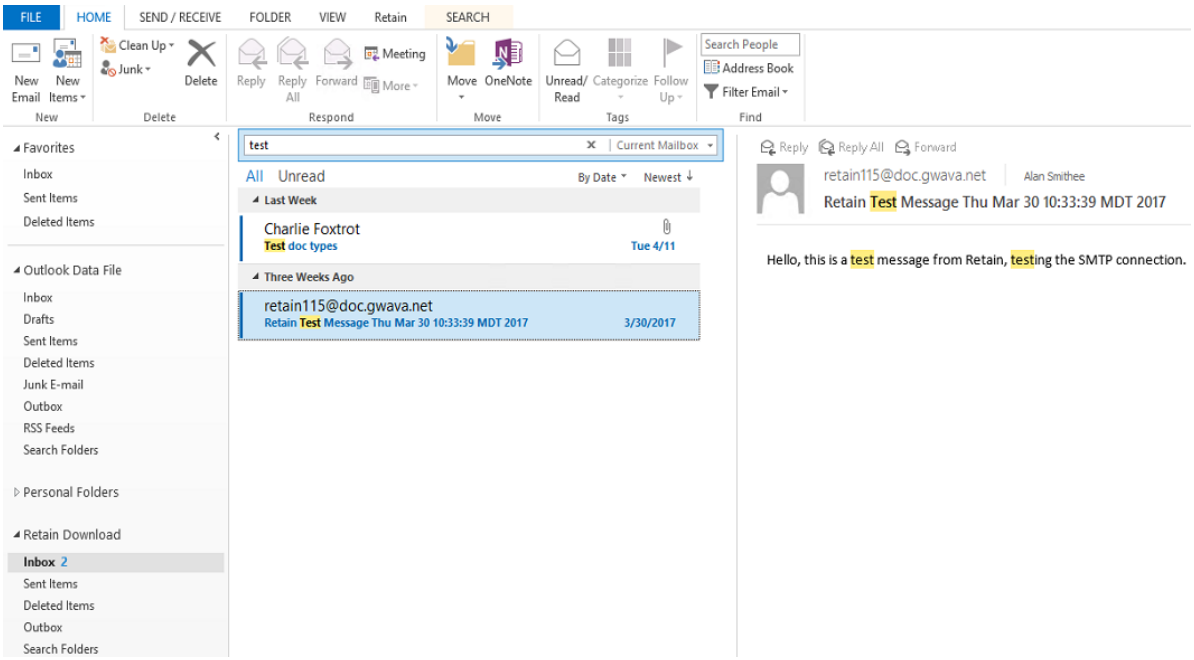


Figure 6: Outlook Search

Outlook Calendar

Retain archives calendar items and it can be shown side-by-side to your Exchange calendar.

Just click the checkbox to enable the *Retain Download* calendar.

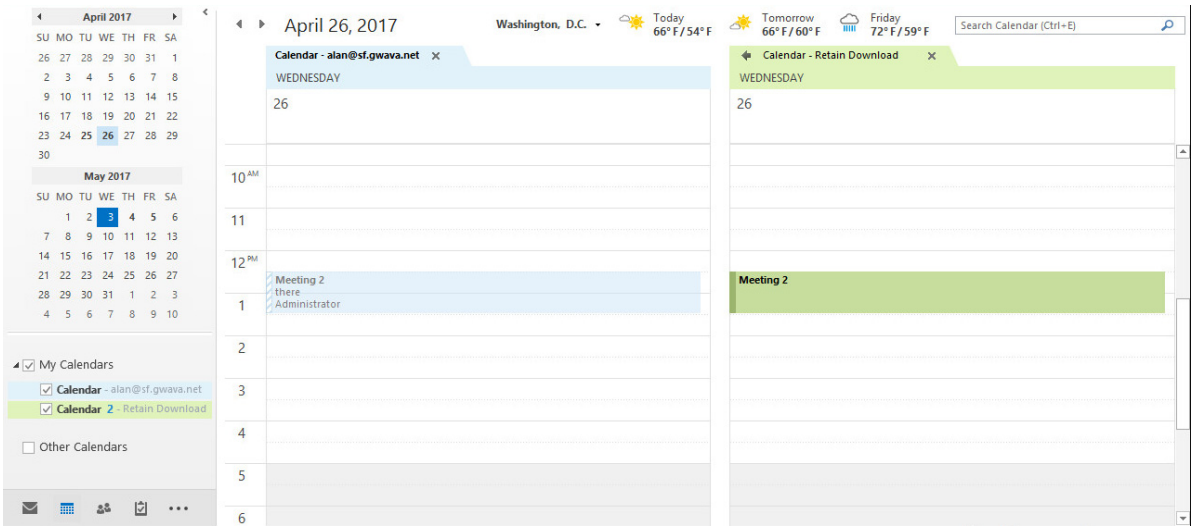


Figure 7: Side-by-side live and Retain archived calendars

Outlook Tasks

Retain archives Tasks

Select the *Retain Download* Task list to see items stored in Retain.

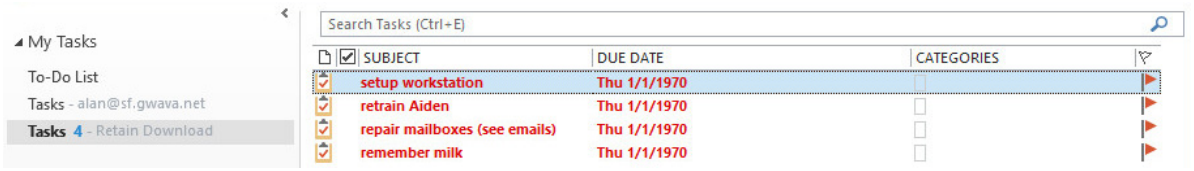


Figure 8: Tasks archived in Retain

Outlook Notes

Retain archives Notes.

Select the *Retain Download* Notes list to see items stored in Retain.

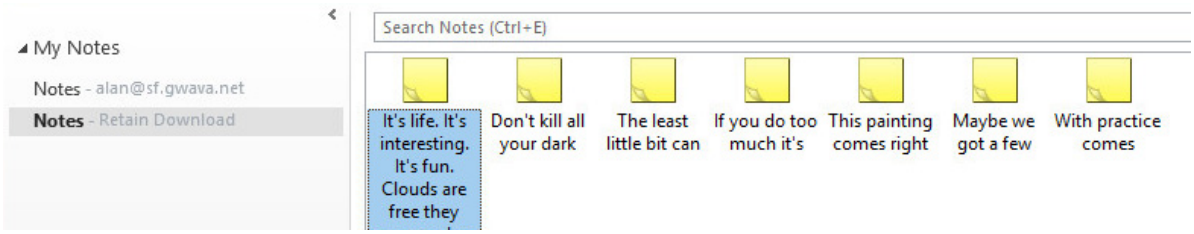


Figure 9: Notes archived in Retain

8.11.5 Other buttons under the Retain tab

The Retain ribbon has a number of available functions. These may be restricted by your System Administrator.



Figure 10: Tab buttons

The other buttons under the Retain tab:

- *Retain download* will manually download messages from Retain if you want to update right now. The default is every 3 hours.
- *Retain search* will perform a search your Retain mailbox on the Retain Server. The results will be downloaded into the Retain search results folder.
- *Download settings* allows you to customize how many and how often messages are downloaded as you can see in Figure 1 above.
- *Launch Retain* will open your Retain mailbox in your browser.
- *Retain Outlook plugin settings* is where you set the URL to the Retain Server and other settings like Language. Generally, only your system Administrator will need to change these settings.
- *About* tells you which version of the plugin is installed.

Download Settings

From this dialog box you can set certain behaviors for the Retain Outlook Plugin.

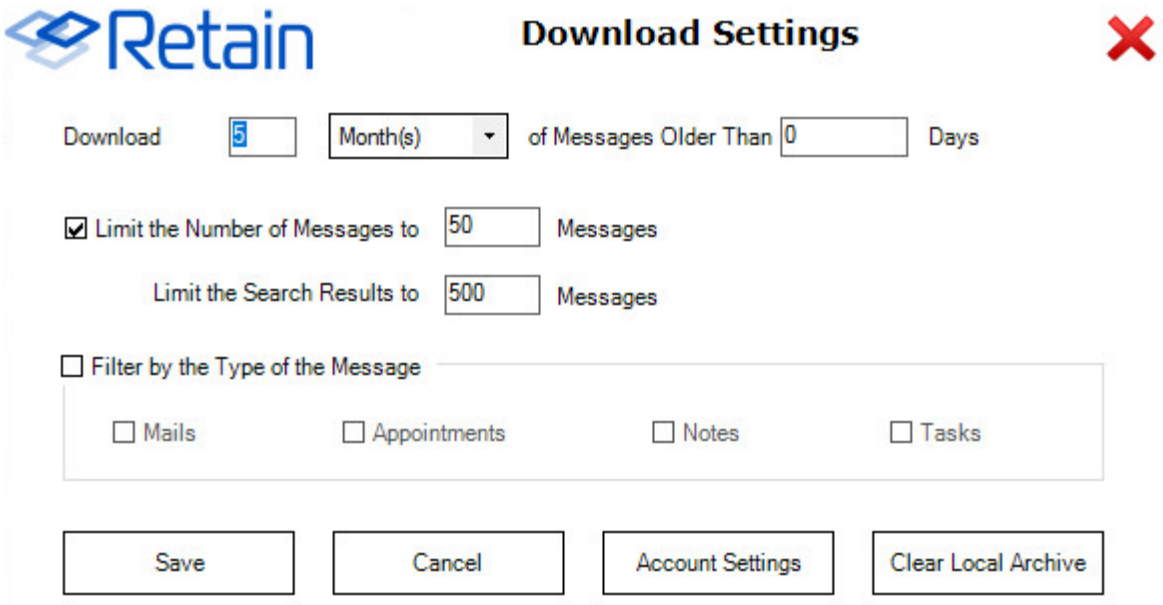


Figure 11: Download Settings Dialog

- **Download:** The number of Days, Weeks, Months, or Years of data to download can be set. Default, 5 Months
- **Messages Older than:** How many days back that can be ignored can be set as well. For example, you might want to look back into the archive for 3 months of items but you don't need the last 30 days. Default, 0 Days.
- **Limit the Number of Messages to:** How many messages to download to you workstation. Outlook does not like to deal with very large local storage files so it is best to keep this number small. Default, 50 messages.
- **Limit the Search Results to:** This will limit the number of messages downloaded after a search. Default, 500 messages
- **Filter by the Type of the Message:** You can restrict the plugin to Mail, Appointments, Notes, and/or Tasks. With this optioned checked, you will have to check which items are to be downloaded. Default, all unchecked, so all items will be downloaded.
- **Clear Local Archive:** This will empty the local archive folder on your workstation of all items, and folders except the default set of folders. To remove all folders the plugin will need to be uninstalled.

Plugin Settings

This dialog box allows more general changes to the Retain Plugin.

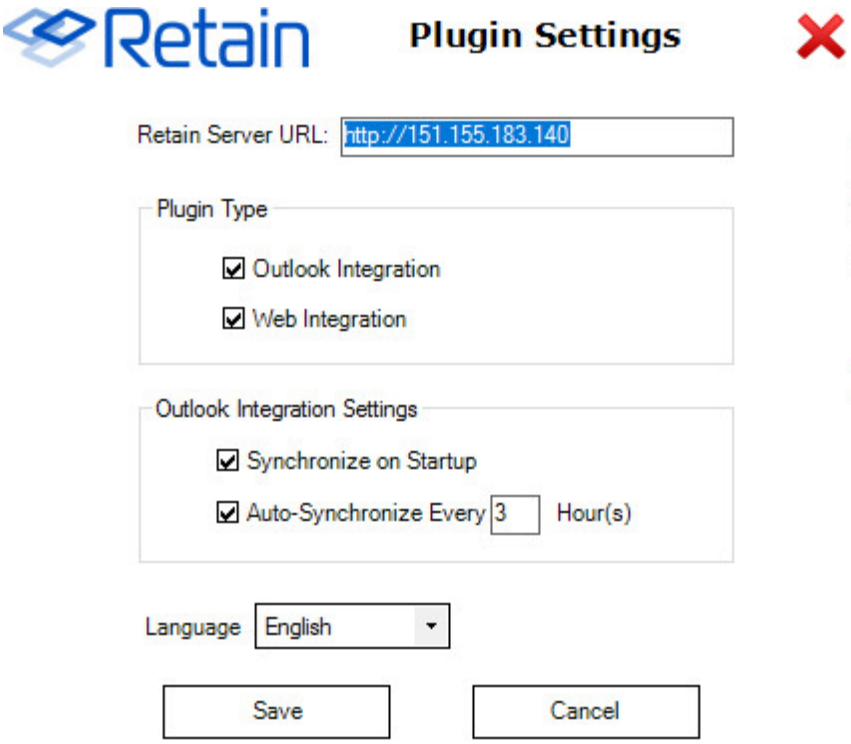


Figure 12: Retain Plugin Settings

- **Retain Server URL:** This is the IP address or Host name of the Retain Server where the archive is stored. Your system administrator will provide or set this for you.
- **Outlook Integration:** This allows the plugin to add the Retain Download folders to Outlook.

- **Web Integration: Launch Retain:** This enables the *Launch Retain* button that opens a new browser window to your Retain archive mailbox.

If Outlook Integration is enabled:

- **Synchronize on Startup:** When Outlook starts up, the Retain plugin to check for new items in the Retain Archive.
- **Auto-Synchronize Every "3" Hours:** The plugin will check for new items in the Retain archive every 3 hours by default.

Language

You can select the following languages for display in the Outlook interface:



8.11.6 Uninstall Retain Outlook client Plugin

To uninstall the plugin

1. Go to *Start* | *Control Panel* | *Programs* | *Uninstall a program*
2. Select the Retain Outlook Plugin
3. Choose Uninstall.

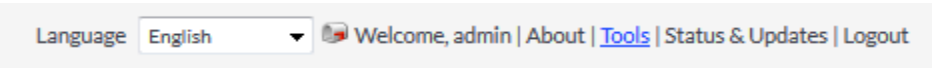
8.12 Outlook Web App Plugin

The Outlook Web App plugin utility adds a Retain Archive auto-logon button to user's Outlook Web App window.

Outlook Web App was called Outlook Web Access in versions of Microsoft Exchange earlier than Exchange 2010.

There are two versions of the Outlook Web App plugin: the Outlook Web Access plugin for Outlook versions 2007 and 2010, and the Outlook Web App plugin for Outlook versions 2013 and 2016. Make sure to download the correct plugin as dictated by the Outlook Web App version.

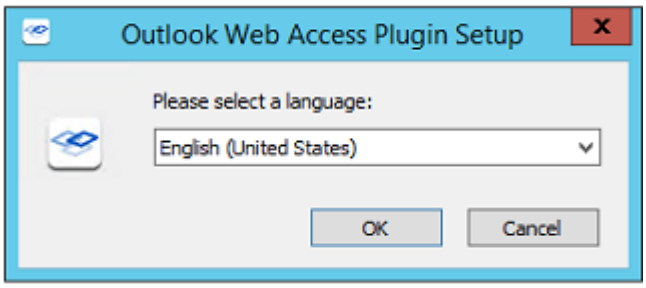
The OWA Plugins are found on the tools page. To access the tools page, select the 'tools' link from the top right of the Retain Server administration page.



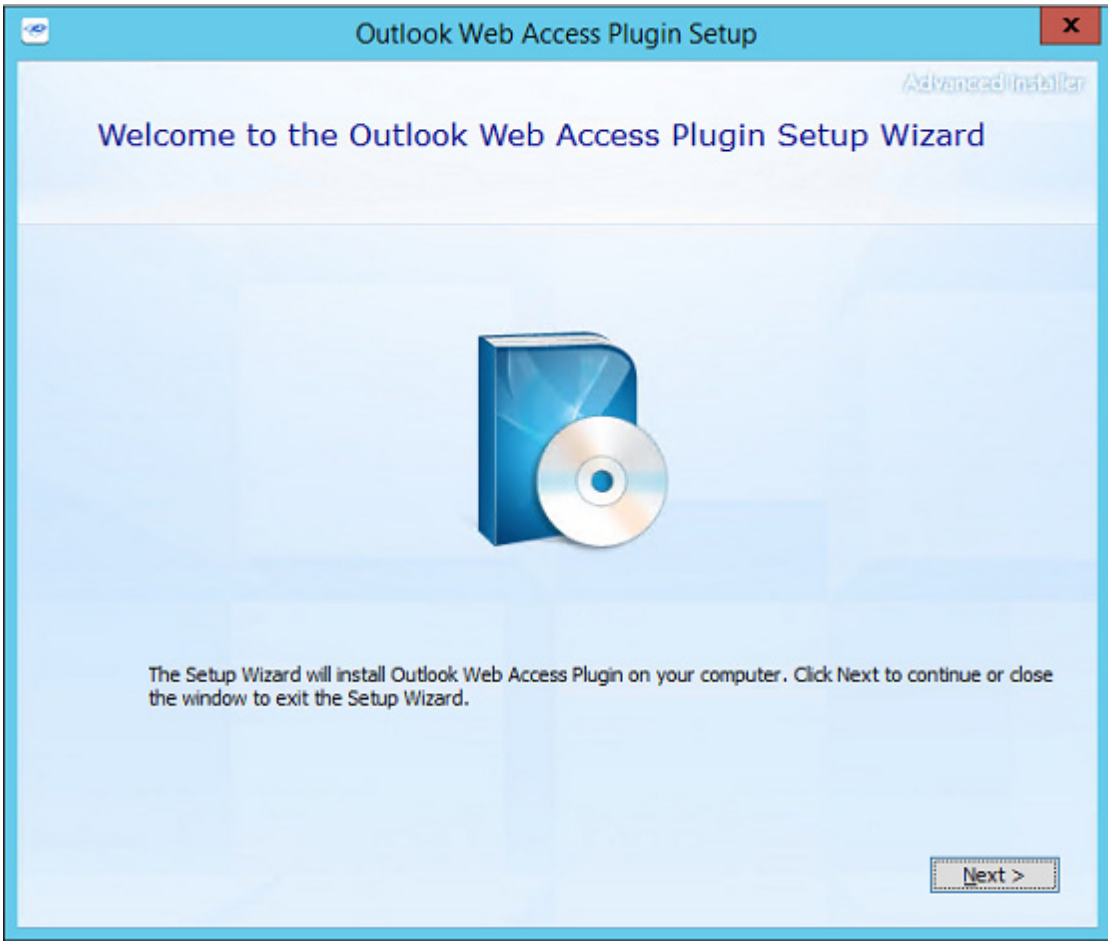
8.12.1 Installation

To install the plugin, download and run the appropriate plugin installation files on the Outlook Web App server.

On initialization, the installation allows language selection.

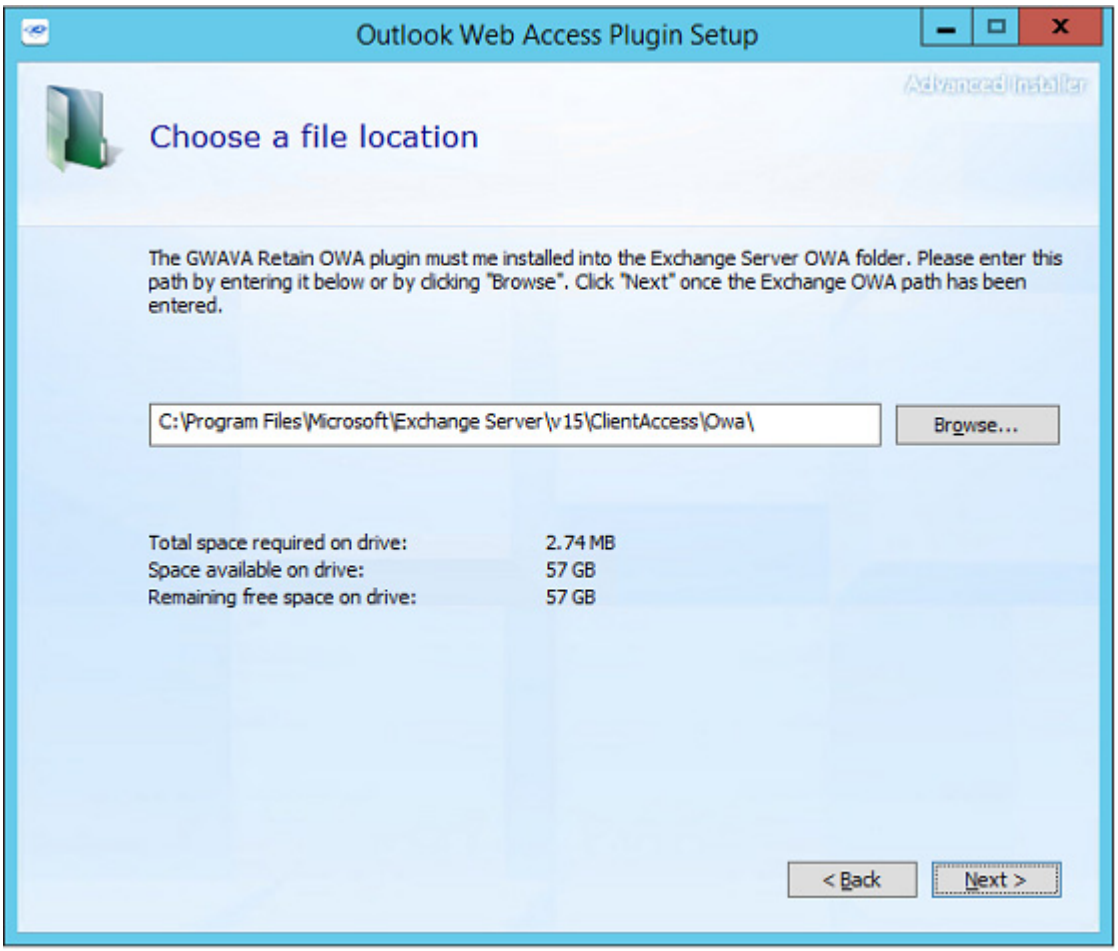


The installation then begins. Select 'Next' to continue.

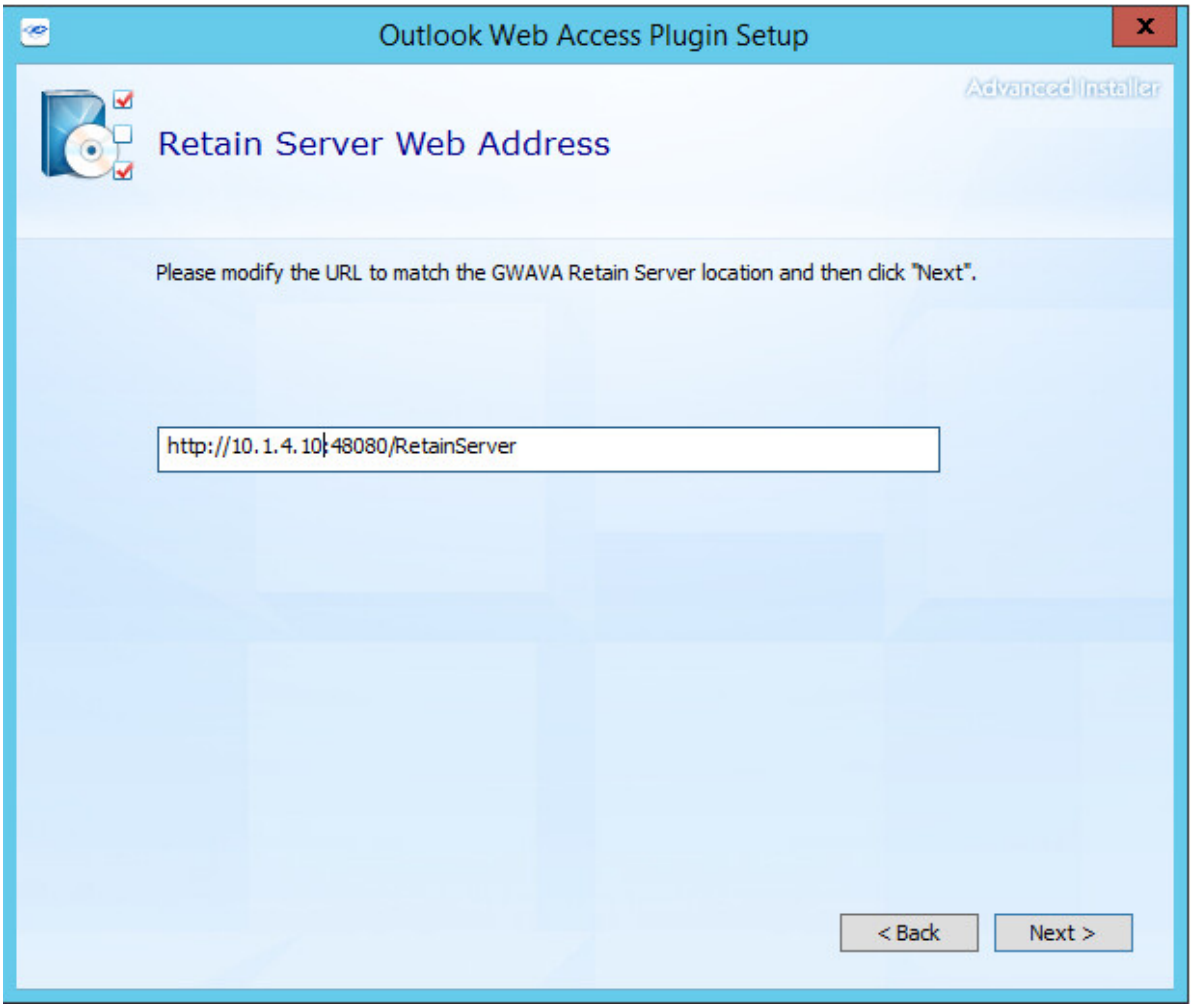


System space requirements and installation path are displayed. The installation path needs to be the location of the OWA folder. Default path for Exchange 2013 is displayed.

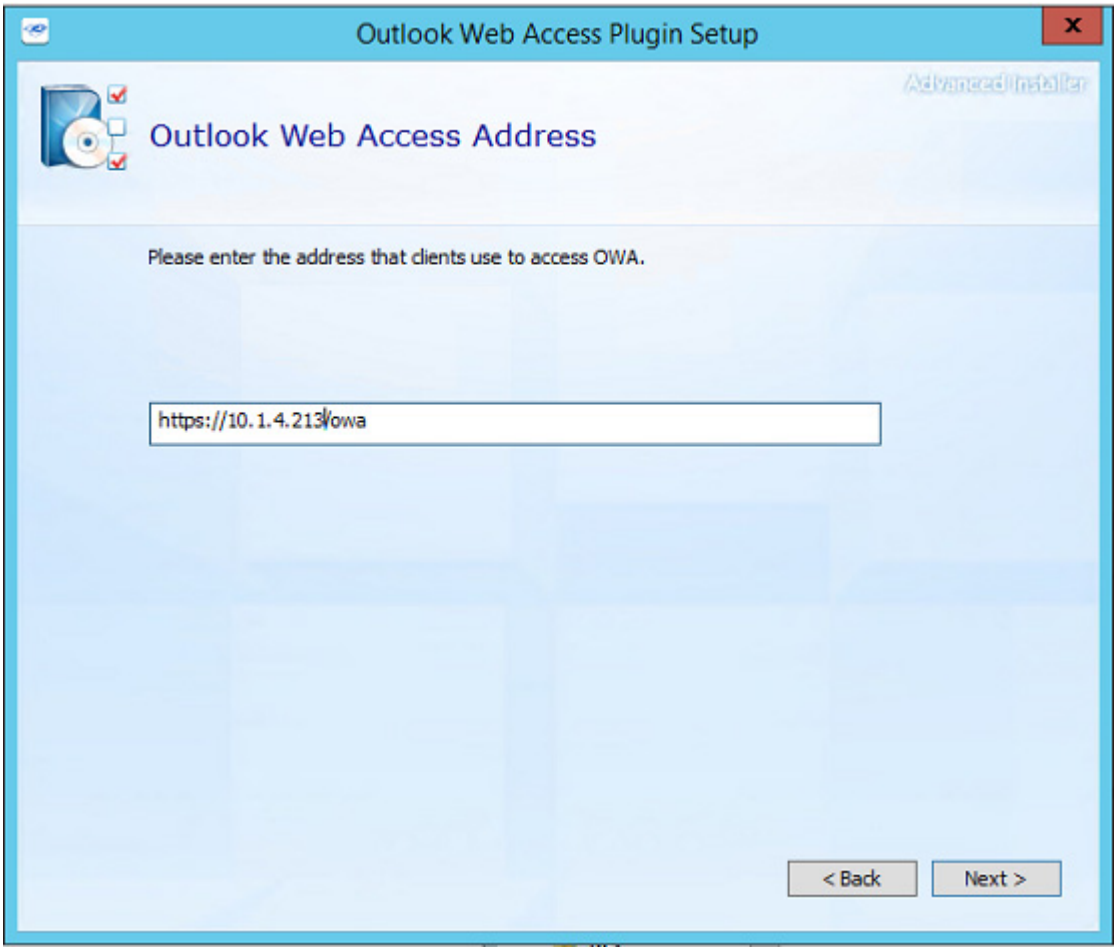
C:\Program Files\Microsoft\Exchange Server\v15\ClientAccess\OWA



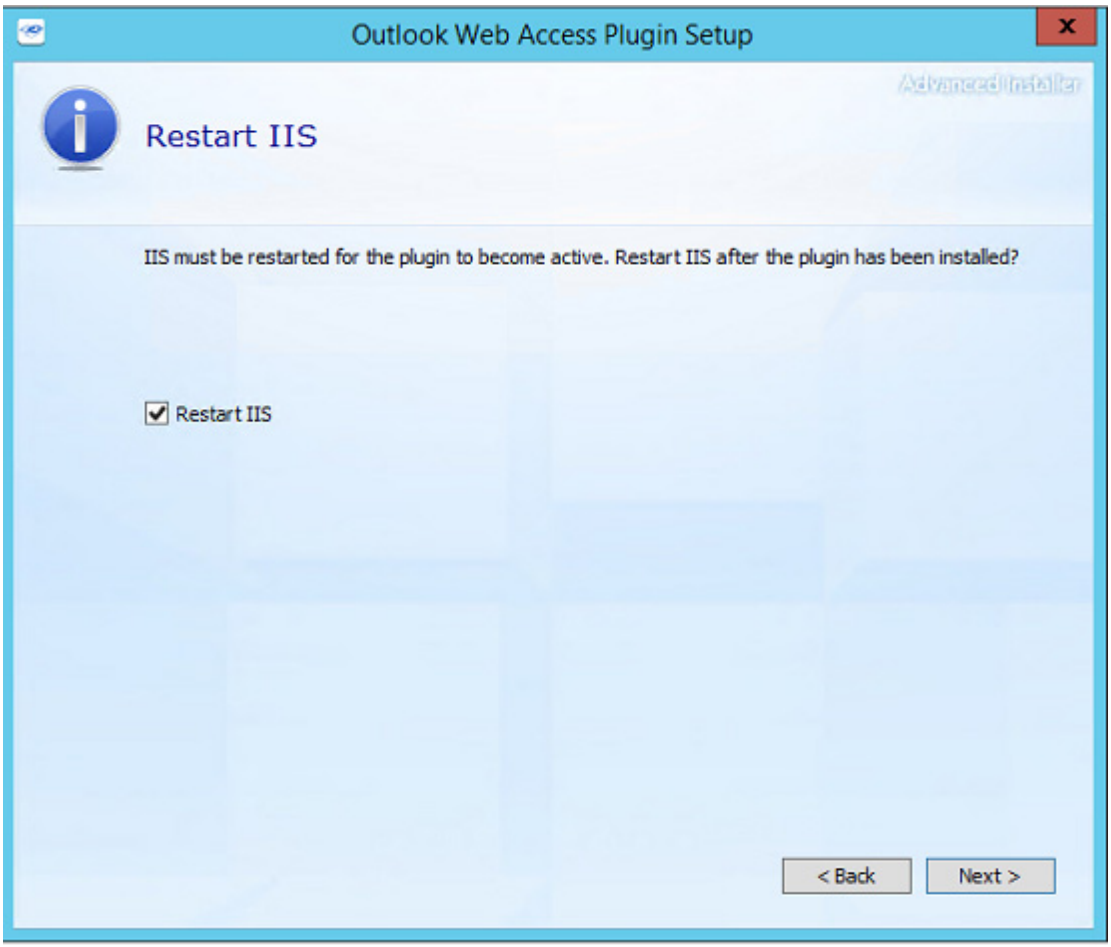
Enter the URL for the Retain Server that has data for the users. This is the address which will be used to access the Retain Server. Ensure that the address functions. This address should include the port used by tomcat, default 48080.



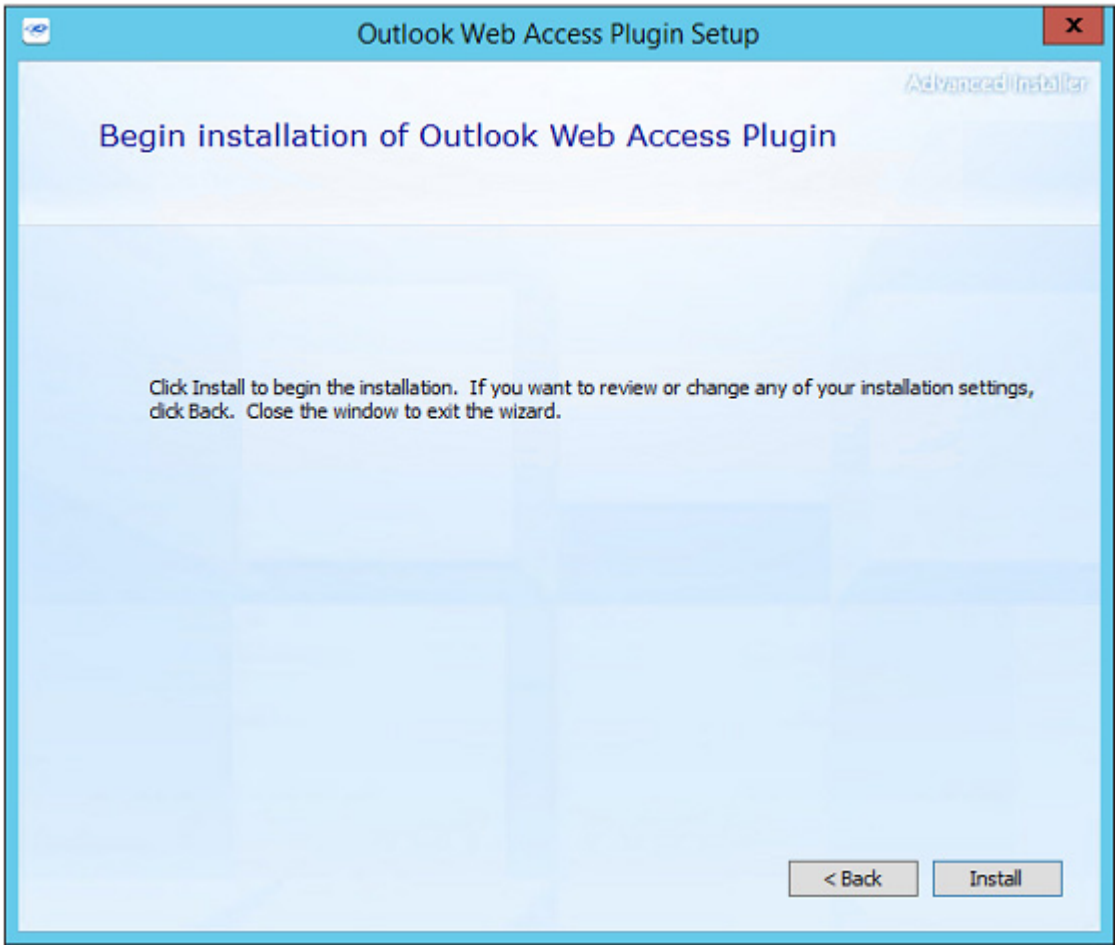
Enter the Outlook Web App client access URL. This address is the connection used by clients to access the Outlook Web App, usually this is an Exchange server. (This step is not included in or required by the 2007/2010 version install)



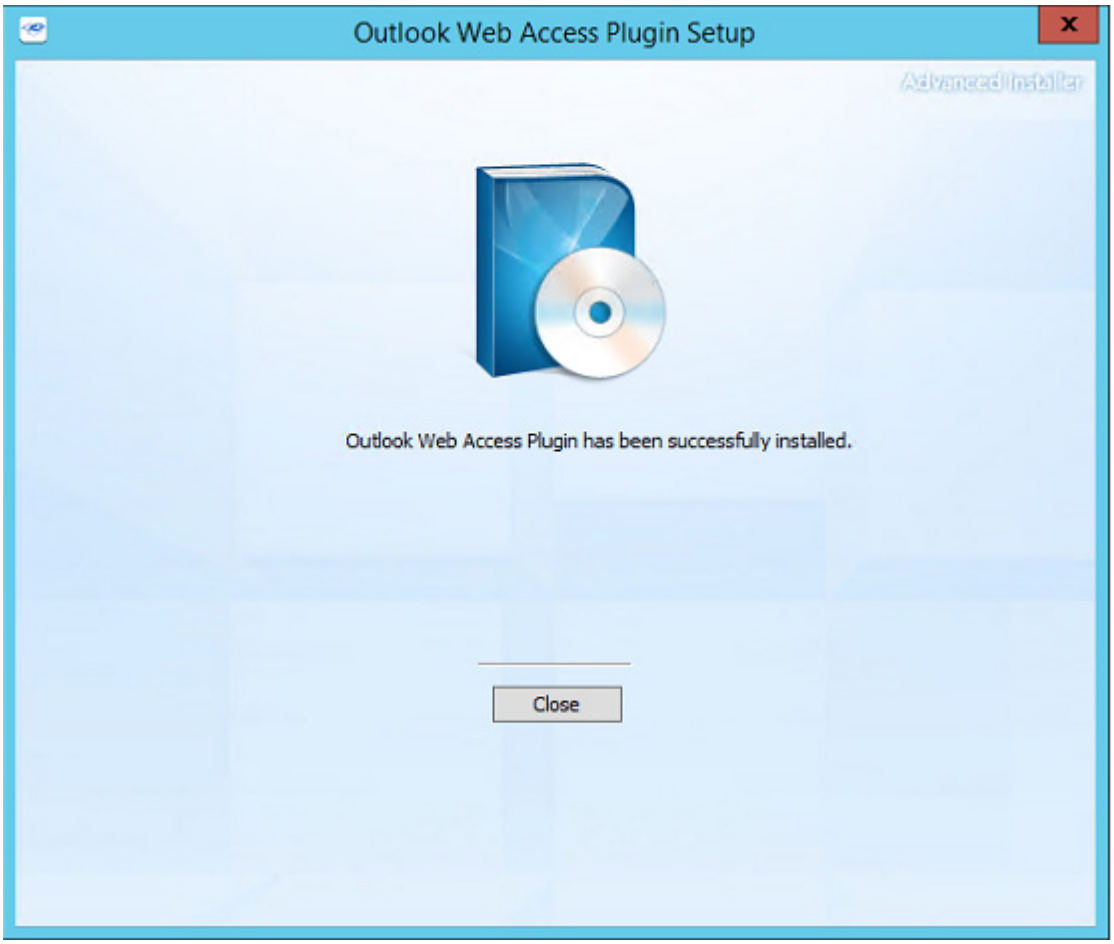
IIS must be restarted to load the plugin into Outlook Web App. The installation has the ability to automatically restart IIS to complete the installation. If checked, IIS will be restarted as part of the install.



If any settings need to be changed, this is the opportunity to change them. Once the 'Install' button is pressed, the installation will begin.

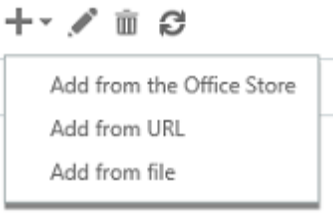


Once the installation has completed, select the 'Close' button.



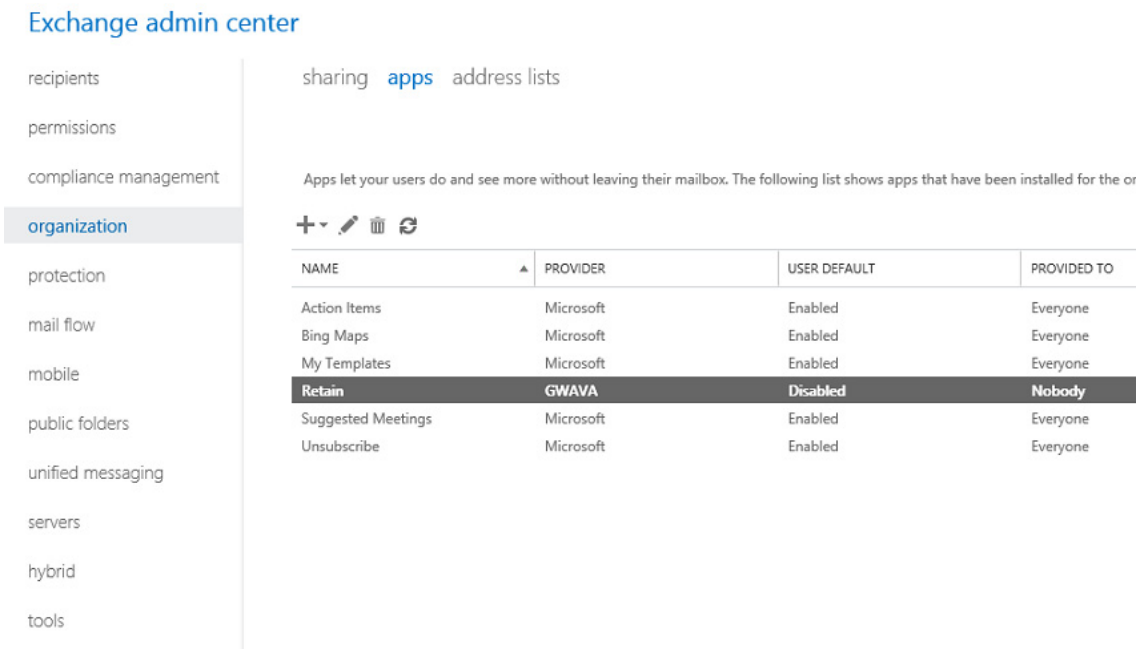
8.12.2 Enable Plugin for the Organization

- After IIS is restarted, log into the Exchange Admin Center.
- Browse to Organization | Apps and select "Add from file" from the Add plus sign menu.



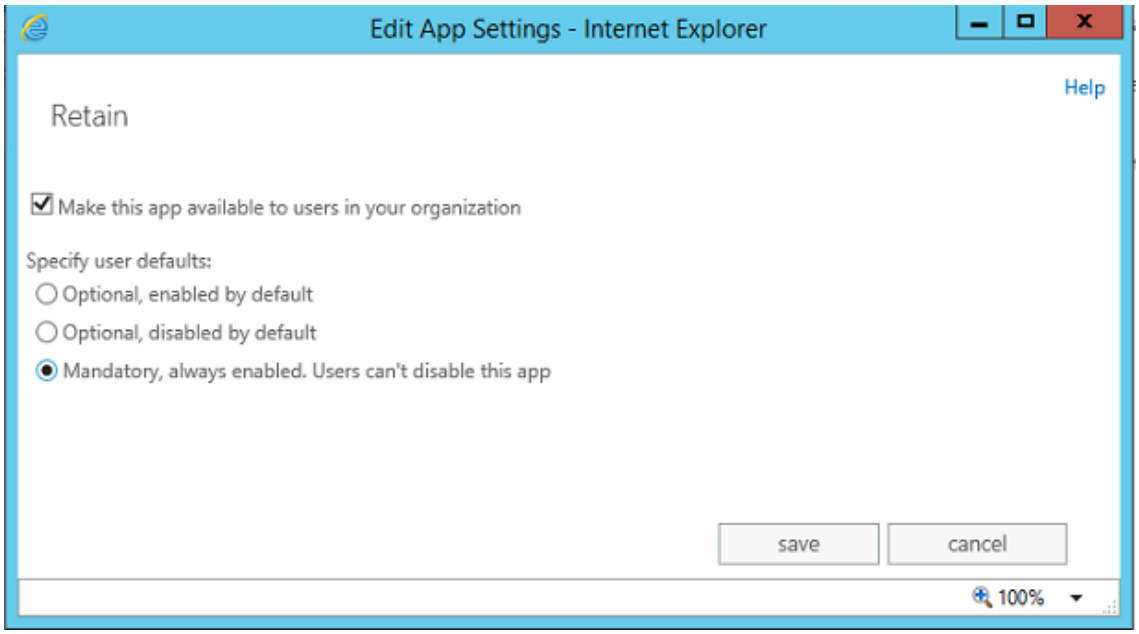
- Browse to the install location. Example from Exchange 2013
  C:\Program Files\Microsoft\Exchange Server\v15\ClientAccess\Owa\Retain

- Select *manifest.xml* and Open.



Double click to edit

- Enable "Make this app available to users in your organization"
- Select "Mandatory, always enabled. Users can't disable the app"



8.12.3 The Plugin in the Outlook Web App

If you did not make the plugin mandatory, the plugin may not be active:

- Click the gear icon near the top right next to your users name and select "Manage apps...".

- A list of apps available to the users will be listed. Notice that Retain App is available but disabled. Select the Retain app and in the right sidebar, click the Enable link.
- Navigate back to your inbox with the right arrow or by clicking on your username and choosing My Mail.

Using the plugin:

- Browse to any message. Within the Viewer pane on the right or if you open the message in its own window, you will see a Retain link inside a grey bar.
- Clicking the link will reveal a Retain logo image and clicking on that will open a new window in the Retain archive.

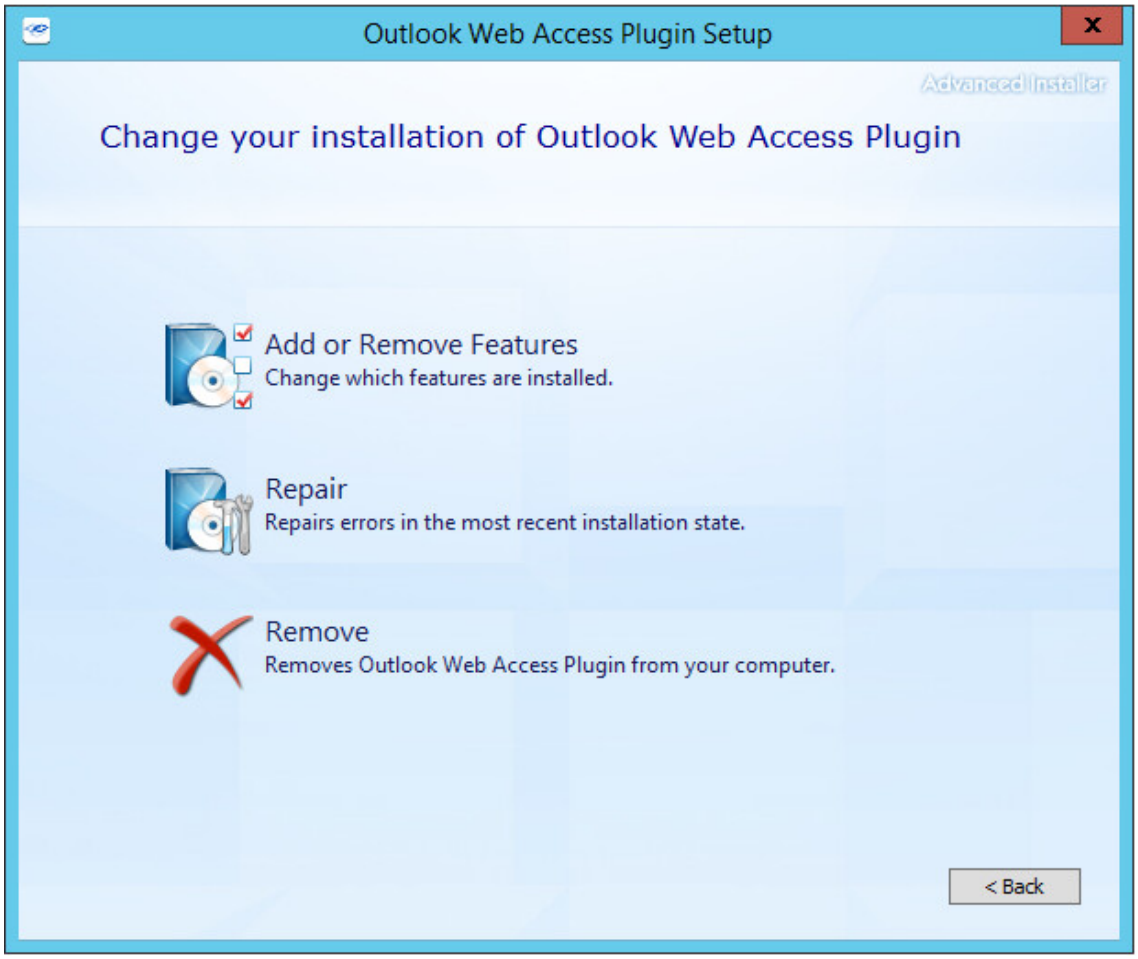


8.12.4 Change Features, Repair or Remove the Plugin

To add or remove features, repair or remove the plugin, run the installer on the Outlook Web App server again. The plugin will detect that it is currently installed and give you these options:

- Add or Remove features will allow you to adjust the features installed to the plugin.
- Repair will refresh the installation of the plugin.

- Remove will delete the plugin.

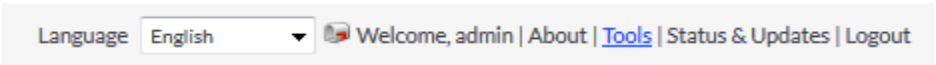


If you entered incorrect data or changed the location of the Retain server, you will need to uninstall the plugin and then re-install to enter the correct data.

8.13 Retain PST Importer

The Retain PST Importer is a tool which inputs mail from PST archives into the Retain system. The PST Importer is a utility which is installed on a Windows machine, local to the PST archive and works in coor-dination with the Retain Server.

The Retain PST Importer is found on the tools page. To access the tools page, select the 'tools' link from the top right of the Retain Server administration page.



8.13.1 System Requirements

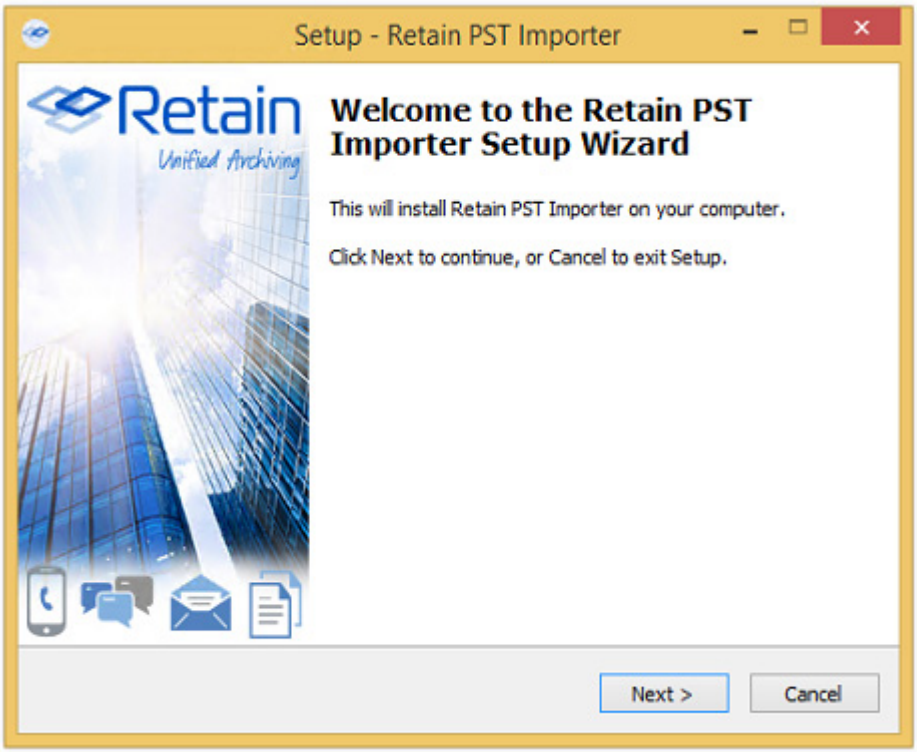
- Java Runtime Environment 8 (No other Java can be present on the system.)
- Windows 7x, 8x, 10x
- Open and direct network connection to the Retain Server

NOTE: If the correct Java is not present in the system, the Importer install will not start, but will provide a link to the download of the appropriate JDK. Select the correct version, (32 or 64-bit), download and install it. Due to the way Windows interacts with Java, there can be no other version of Java on the host machine. Uninstall any other version.

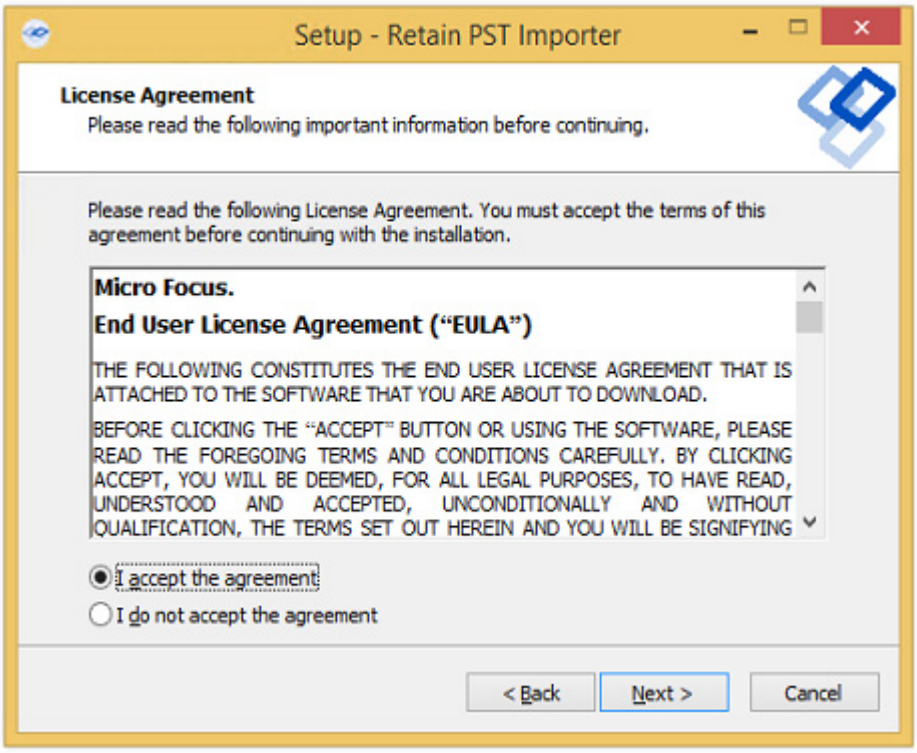
8.13.2 Installation

Installation of the PST Importer is simple. Locate the ‘Retain PST Importer Install.exe’ file and run it.

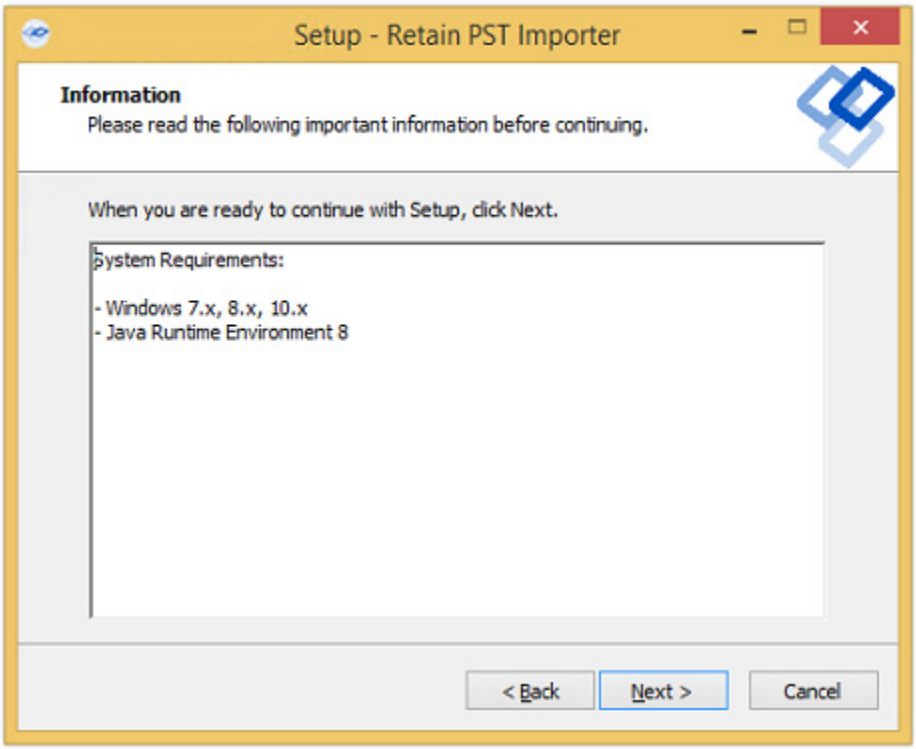
Once initialized, the installation will display this page. Select ‘Next’ to continue.



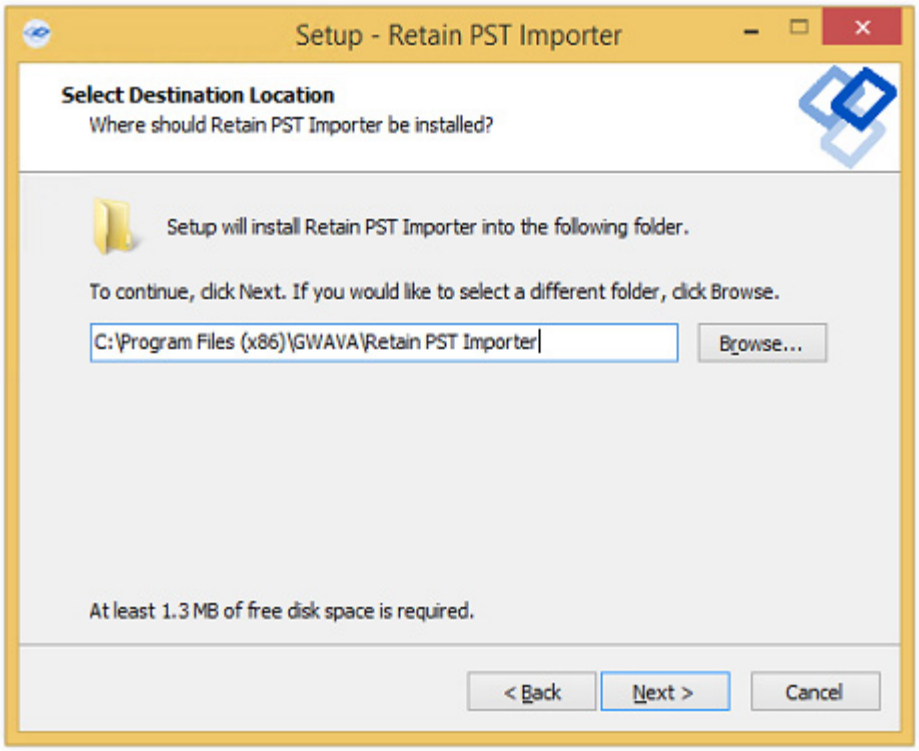
Accept the License agreement and select ‘Next’ to continue.



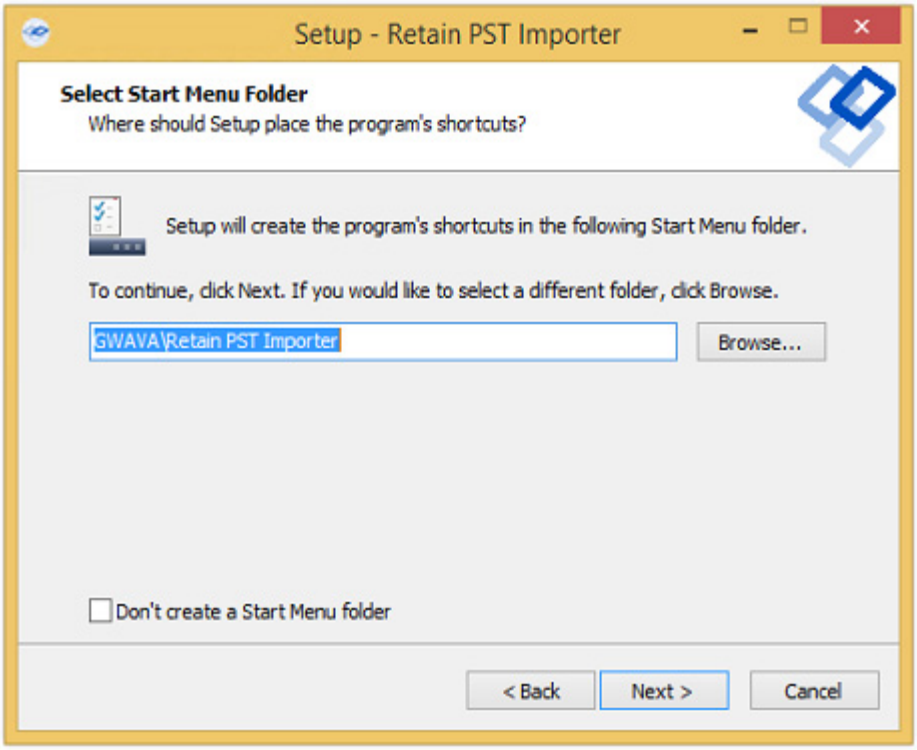
The System Requirements are listed. Select 'Next' to continue.



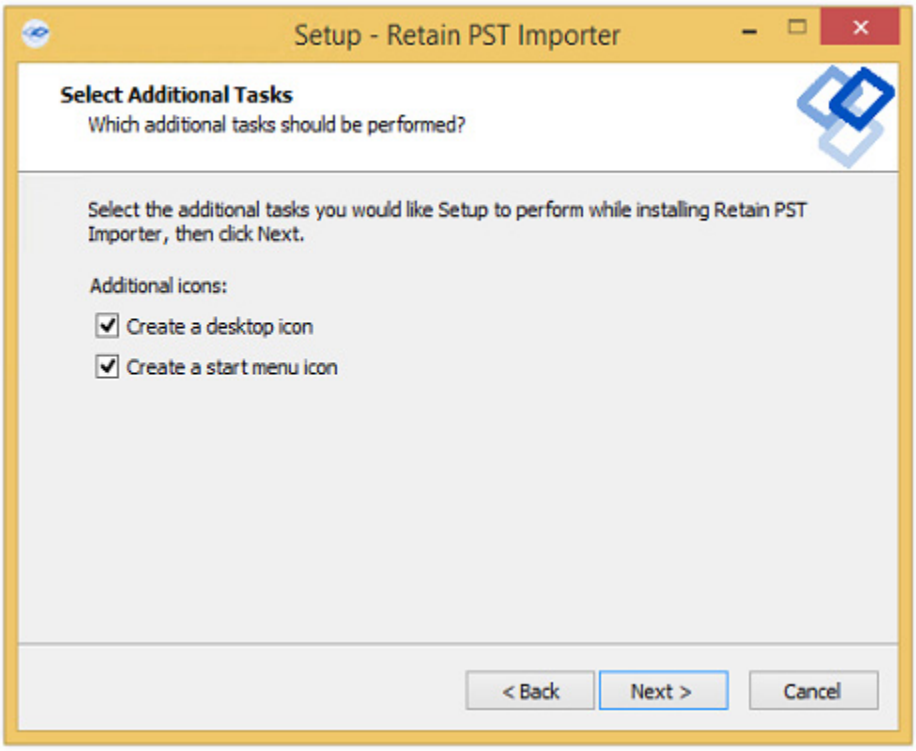
The installation folder is displayed. Change it as desired and select 'Next' to continue.



The PST Importer may be installed with a shortcut in any desired start menu folder. The default folder is displayed. Change as desired and select ‘Next’ to continue.

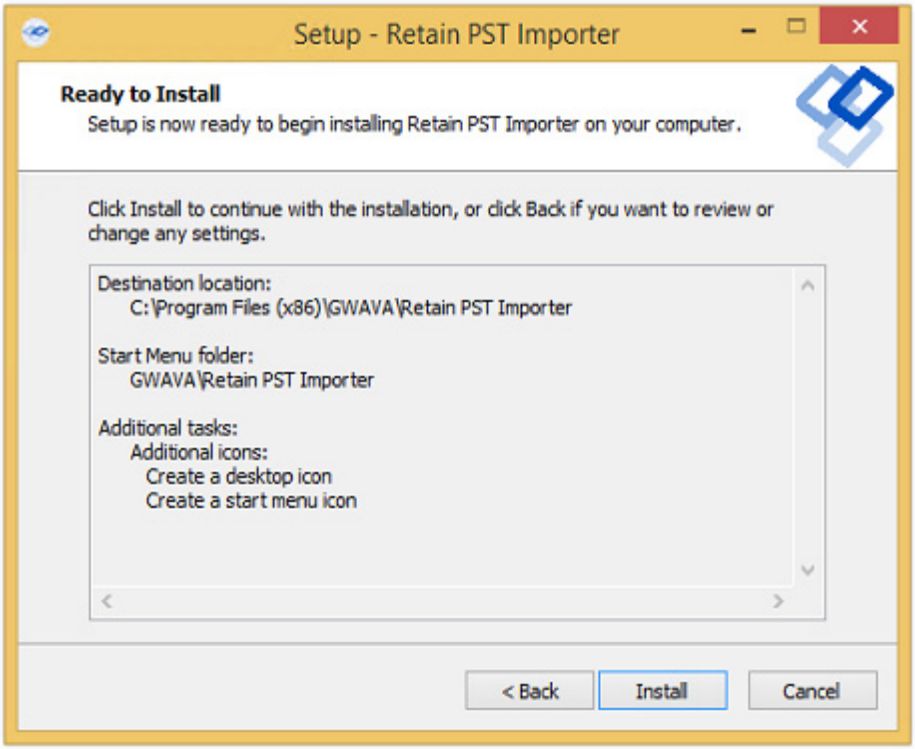


Select which icons are going to be installed in the system and select ‘Next’ to continue.



Review the installation settings. If any settings need to be changed or revised, use the ‘Back’ button to return to the appropriate pages and change them.

When correct, select ‘Install’ to install the program.



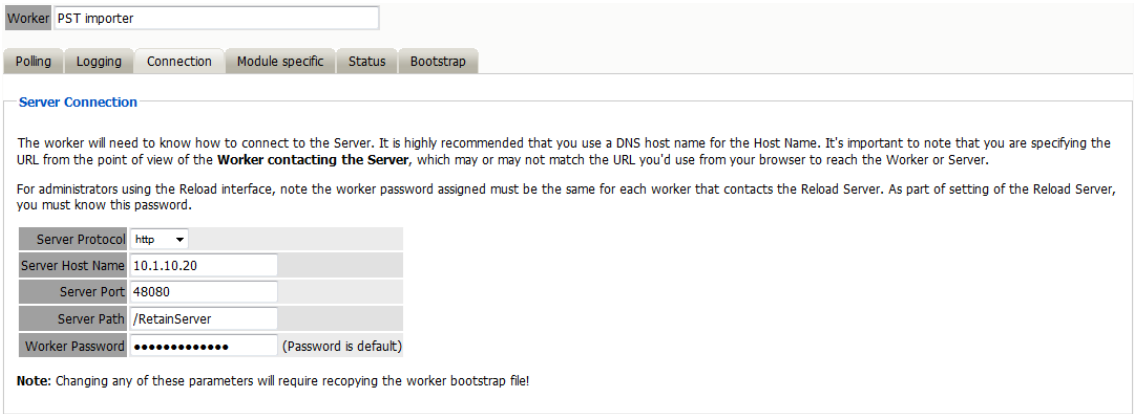
Once the installation is completed successfully, select the ‘Finish’ button.



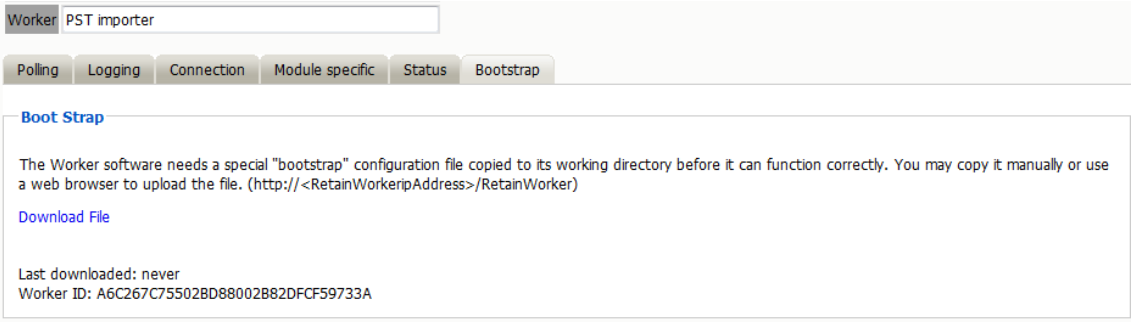
Configuration and Use

The PST Importer requires a Worker configuration and a bootstrap file to know where and how to com-municate with the Retain Server. A new worker must be created, configured, and the configuration file loaded into the PST Importer. Follow the steps below.

1. In the Retain Server administration console, create a new Worker. It would be appropriate to name it ‘Importer’ or ‘pst importer’ for ease of remembrance, but any name desired will do. Remember that this worker should never be configured with any jobs in the Retain Server. A gen-eral configuration is all that is required; no specific settings need to be changed from default, except the connection information if it needs to be modified in accordance to the needs of the local network. It is imperative that the PST Importer has an open connection to the Retain server. When the settings are correct, save the new Worker.

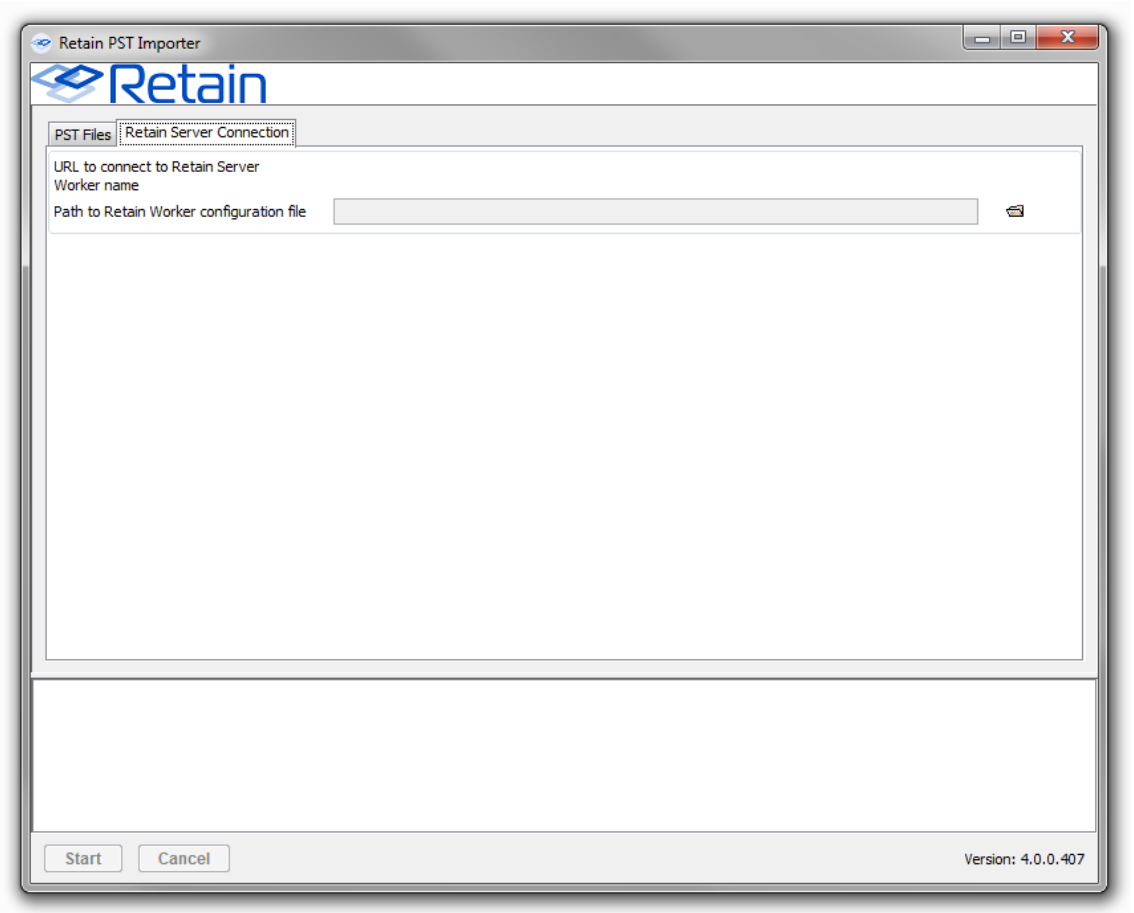


2. Download the bootstrap file.



3. Run the PST Importer utility. Locate the Retain PST Importer icon or start menu item and run the program.

   On initial startup, the Retain PST Importer opens to the configuration page. The worker configuration file contains all the connection information for the Retain Server. (If the location changes, the worker configuration file will need to be re-downloaded.)



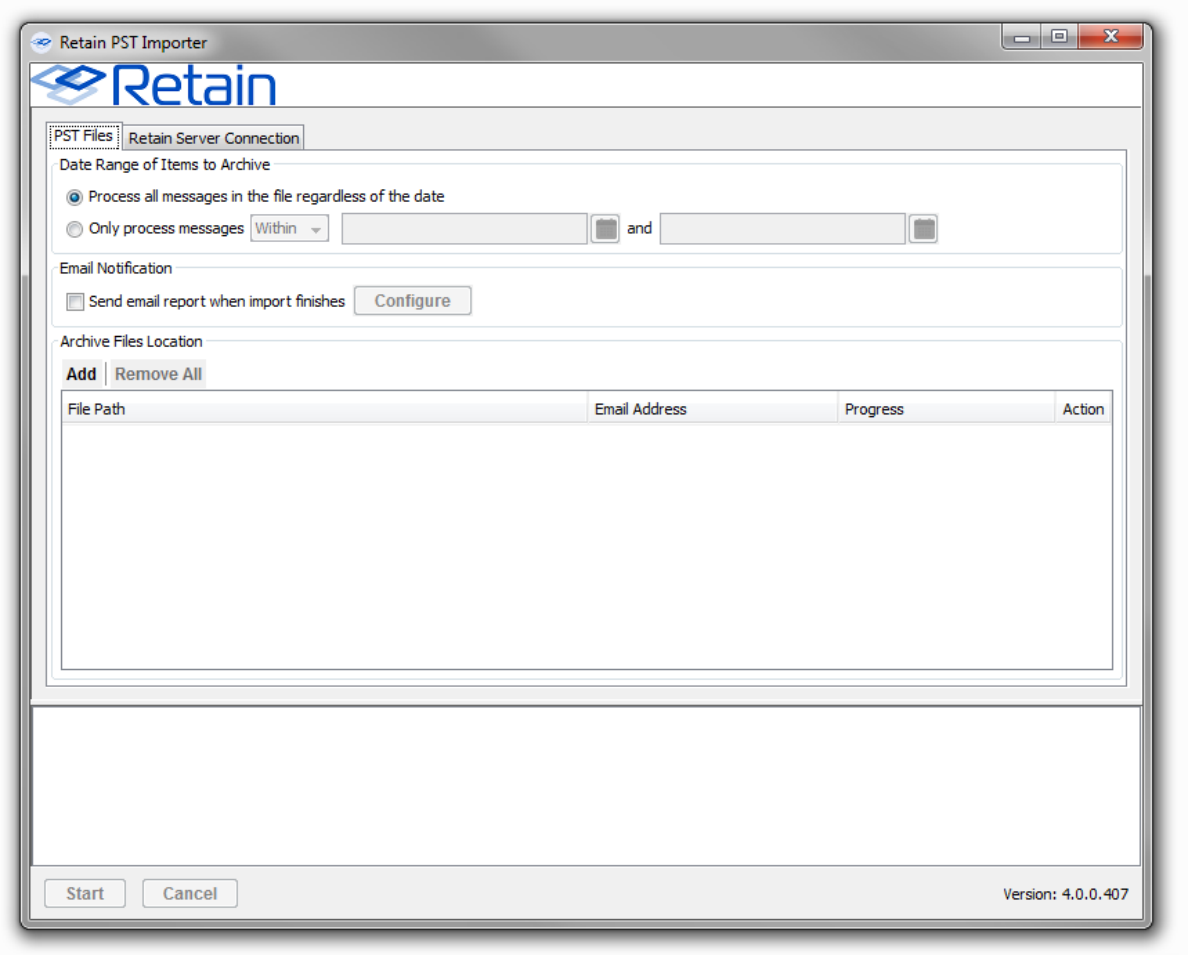
4. Upload the Worker Configuration file to the PST Importer.

   Select the Retain Server Connection tab, where the bootstrap file is specified and select the ‘folder’ to browse to the file location. Select the appropriate worker configuration file and load it into the Import tool.

8.13.3 Use

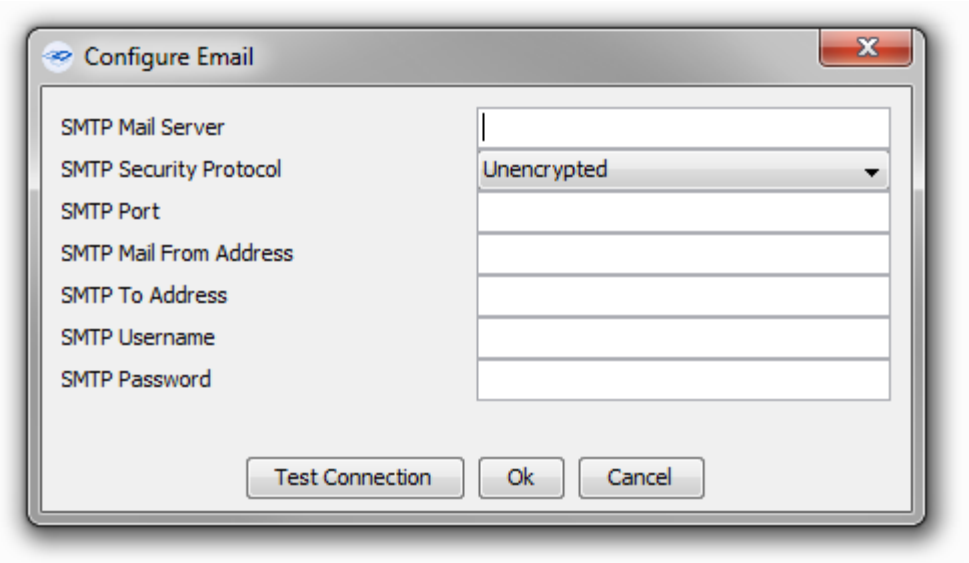
Use of the PST Importer is very simple. The PST Importer tool may be restricted to import only messages in the archive which fall into a specified date range or it will attempt to import the entire archive. The PST Importer may be restricted to import only messages which meet selected criteria:

- A specified date range (inclusive)
- Before a specified date (exclusive)
- After a specified date (exclusive)



If desired, the PST Importer may send a notification message when the job is complete. The email will contain a job summary including:

- The total number of messages selected for import
- Number of successfully process messages
- Number of messages not processed



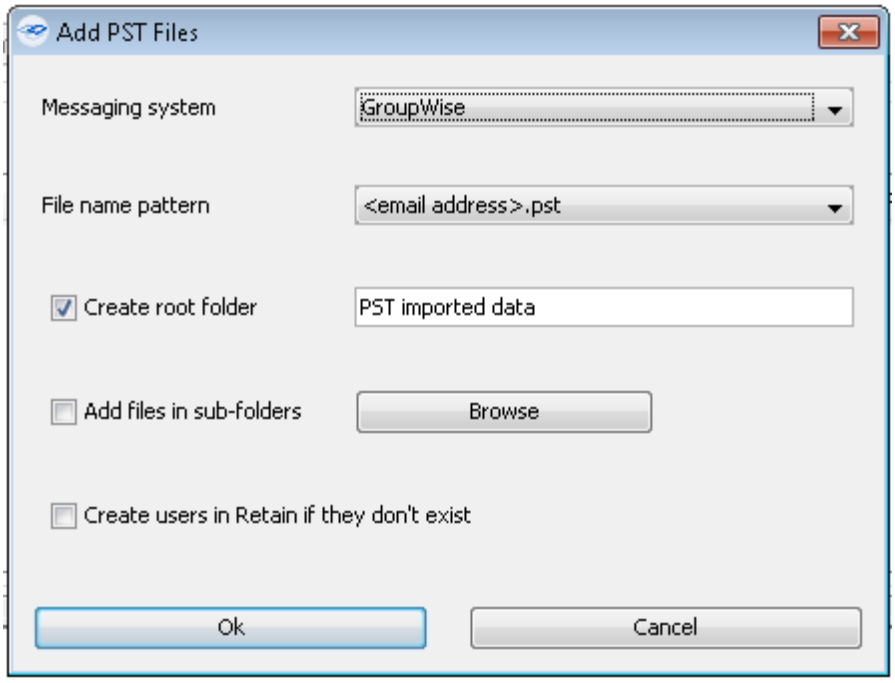
To receive a notification email, an SMTP connector needs to be specified and settings provided. Security protocol options are: ‘Unencrypted, TLS, and SSL’. Test the connection to ensure the settings are correct.

When the desired archive or archive have been added and settings have been configured, select the ‘Start’ button to begin importing the archive. Rolling information is displayed in the information window at the bottom of the screen.

To add or remove an archive to or from the PST Importer tool, use the ‘add’ and ‘remove all’ buttons at the top of the archive window.

Adding an archive opens a browse window to allow selection of the archive. Adding an archive requires the selection of the desired messaging system, (select the system the user will be using so messages will be added to their archive). The user can select the file name pattern corresponding to the PST archives and whether to import the archive into existing root folders or to have the tool import the archive in to automatically created subfolders.

Users can import messages from the PST archive for either existing users in Retain, or if the user does not currently exist in Retain, it will be automatically created.



When complete, select ‘Browse’ button to browse to and select the archive. Hovering over the path will show the complete path.

To remove all the archives from the list, select ‘remove all’.

8.13.4 Logs

Log files are generated and placed in the following default location:

C:/Users/<username>/AppData/Roaming/GWAVA/Retain PST Importer

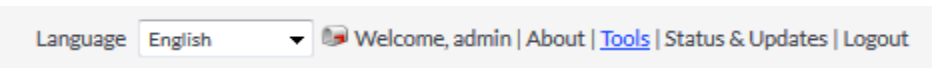
This can be reached by typing %APPDATA% in the File Explorer address bar.

8.14 GroupWise WebAccess Plugin

8.14.1 Overview

The Retain for GroupWise WebAccess Plugin is a quick and effective way to integrate access to the Retain archive into WebAccess for users. The Plugin adds a button which will automatically take users to their personal archive. The archive interface will be displayed inside the WebAccess window, and provides all the same access and management of the archive which users can experience by directly logging in to Retain.

The GroupWise WebAccess Plugin is found on the tools page. To access the tools page, select the 'tools' link from the top right of the Retain Server administration page.



8.14.2 Supported Software

GroupWise WebAccess 8x, 2012, 2014, 2014 R2

8.14.3 System Requirements

- Java 1.7 or later
- Functioning GroupWise WebAccess
- Functioning Retain Server, with configured GroupWise module
- Open network connection to the Retain Server
- If SSL is to be used on the web server, either IIS or Apache, then SSL must be enabled on the Retain server as well.

8.14.4 Installation

To install the plugin, execute the installation file.

Linux:

Linux requires root access to install the plugin. Change the files to executable and run the files.

- chmod +x *.sh
- ./installgui.sh (if a graphical install is desired)

or

- ./installtext.sh (if a text install is desired)

Windows:

On Windows, double click and run the .bat file.

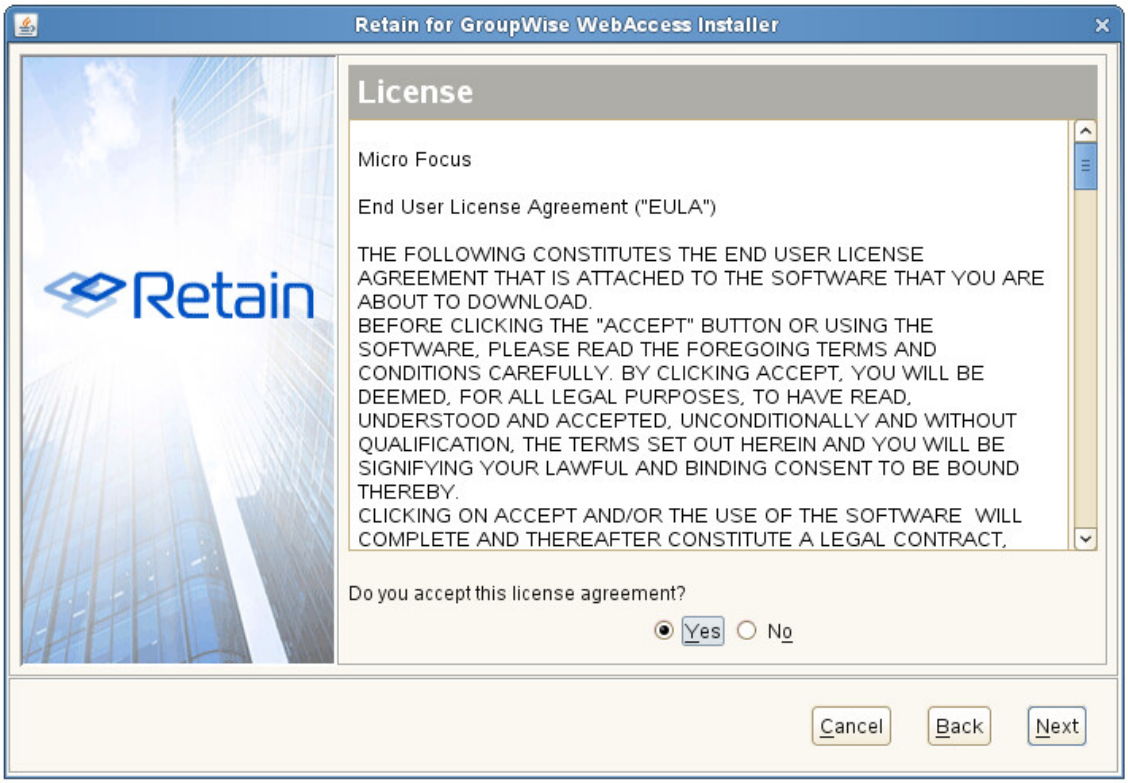
- install.bat

Graphical install shown

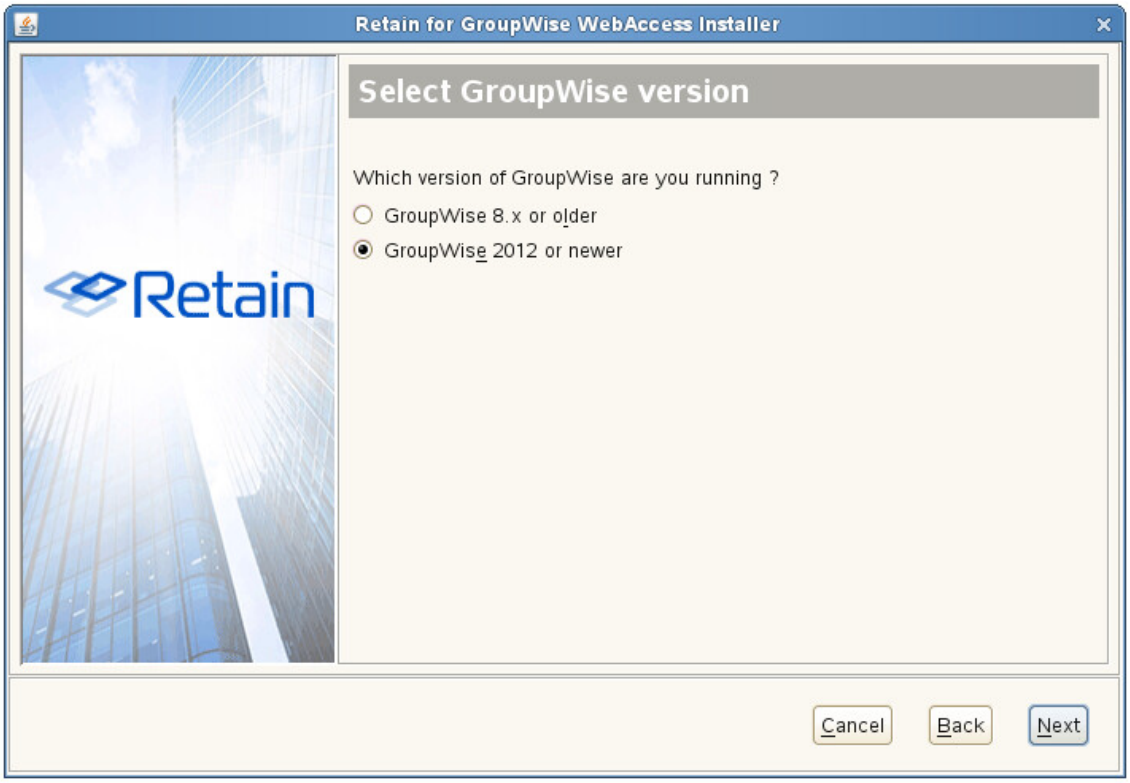
1. Click ‘Next’ to begin the installation



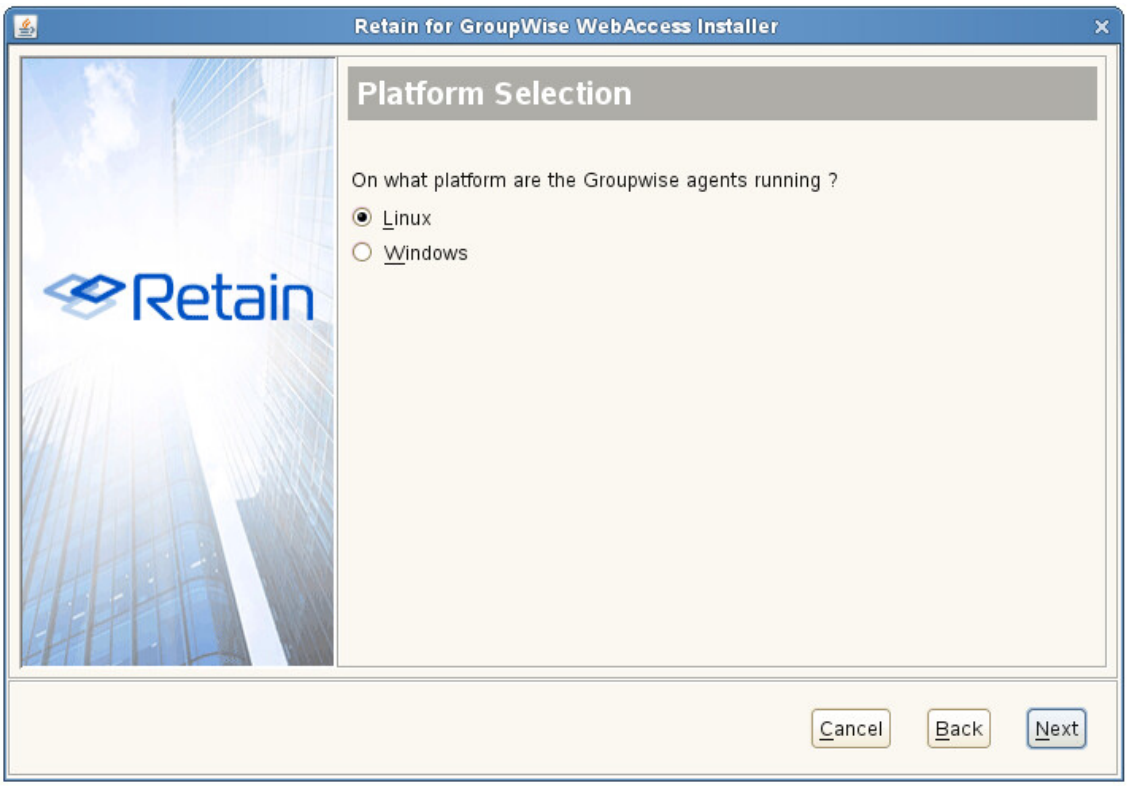
2. Accept the license agreement.



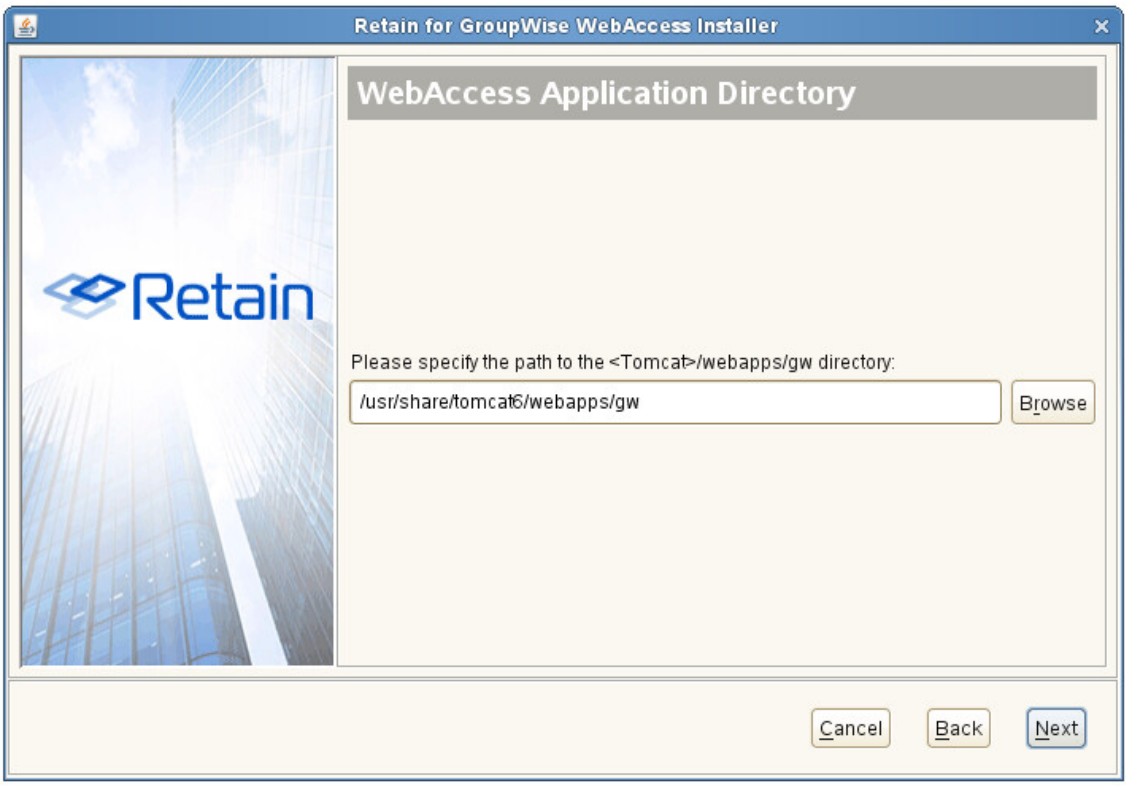
3. Select the GroupWise WebAccess version. GroupWise 8 or older, or GroupWise 2012 or newer.



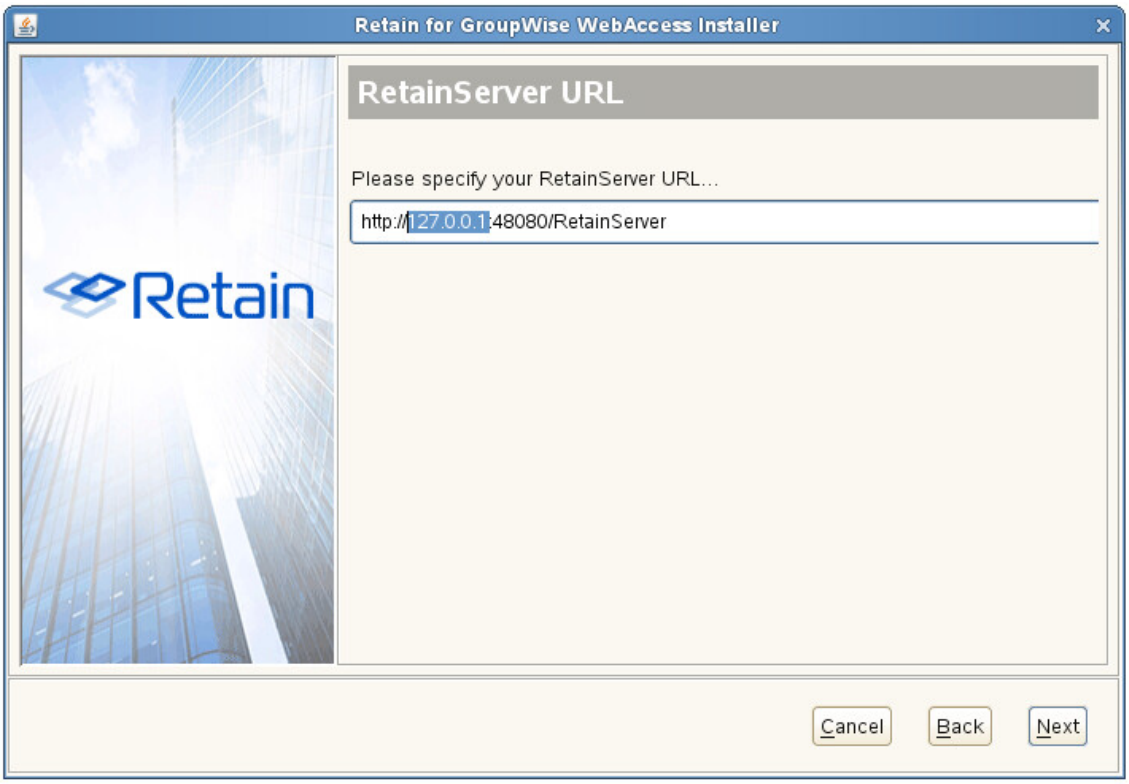
4. Select the platform the WebAccess is being installed to. Linux or Windows.



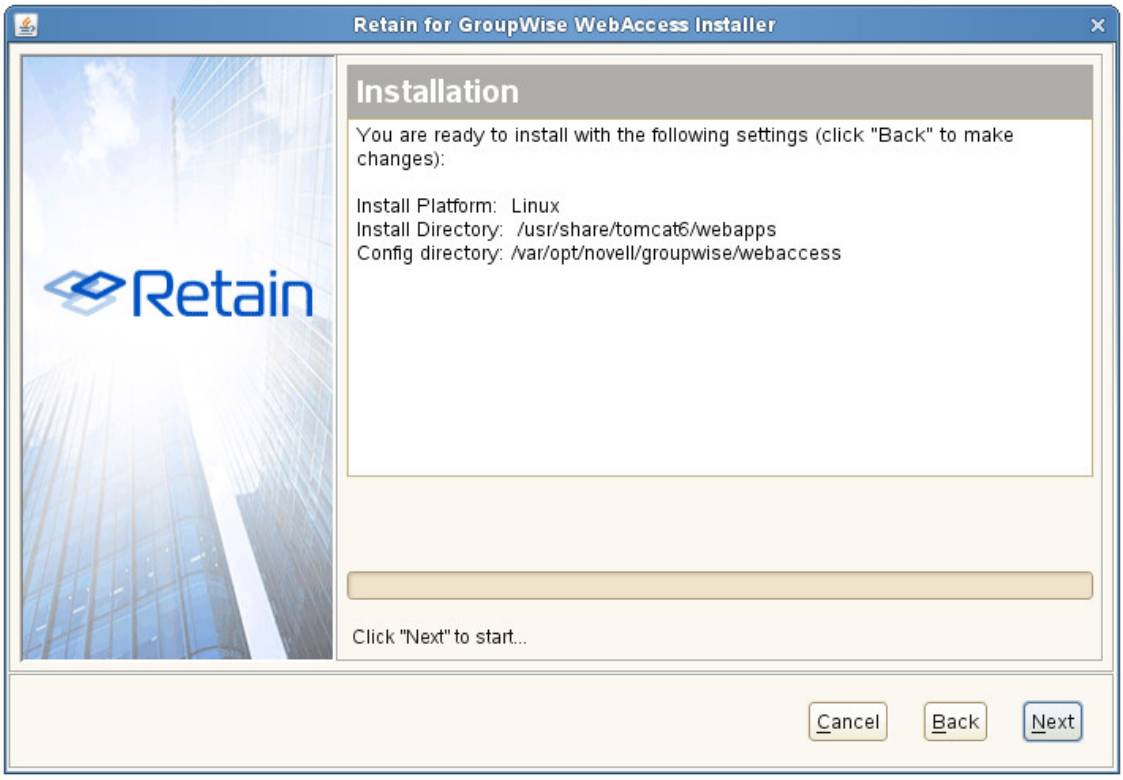
5. Specify the active ...webapps/gw directory, or browse to it.



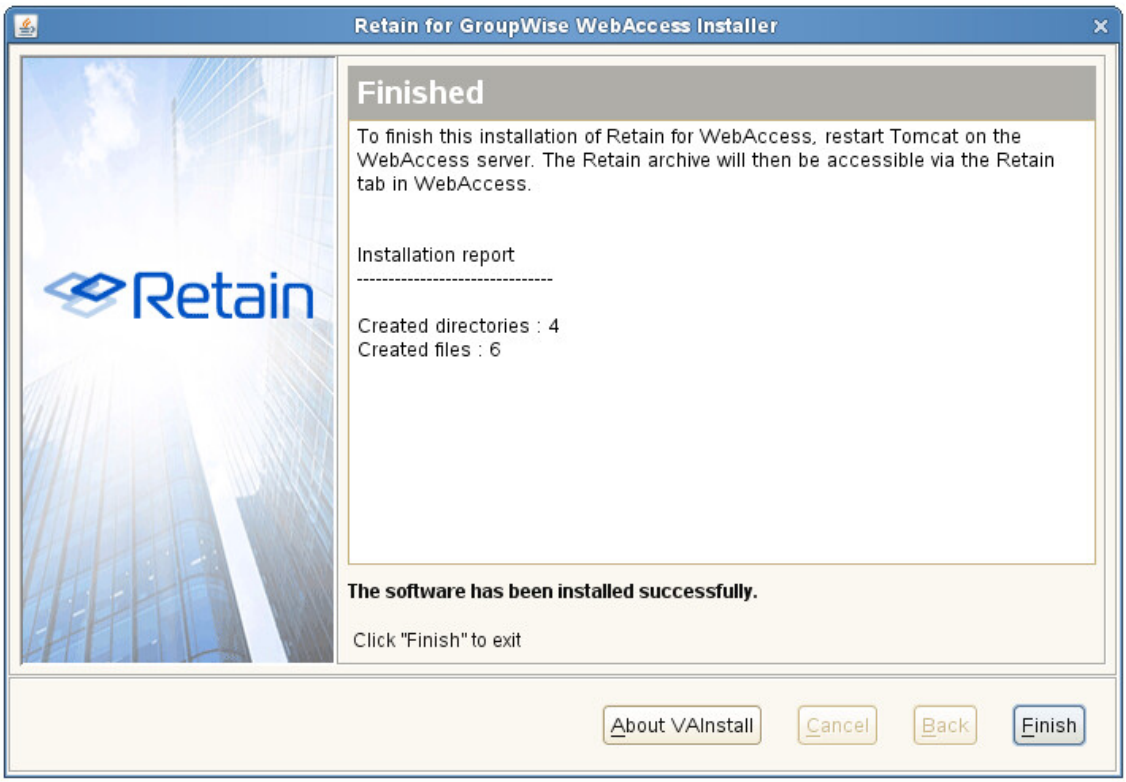
6. Specify the Retain Server URL. An open connection to the Retain Server is required for the plugin to work.



7. Confirm the settings. Click ‘Next’ to begin the install, or ‘Back’ to change the settings.



8. Once finished, click ‘Finish’ to exit the installer.



8.14.5 Post Install Tasks

Once the plugin is installed the final step to install the WebAccess plugin is to create two folders in the Novell folder structure.

Retain cannot create folders or set permissions on the host system, but requires the folders to be created in the Novell directory structure to function.

Linux

Retain requires two folders to be created, as well as the permissions set to match, on the GroupWise server.

Browse to:

/var/opt/novell/groupwise/webaccess/retain/templates/webacc

- Create two new folders here: ‘mobile’ and ‘simple’

  mkdir mobile

  mkdir simple
- Note the file permissions and owner of the ‘css’ folder, and change the permissions of the two new folders to match.

Windows

Two folders need to be created on the GroupWise server. From where GroupWise was installed, browse to the webacc folder. For example,

C:\GroupWise\webaccess\retain\templates\webacc

- Create two new folders here: ‘mobile’ and ‘simple’

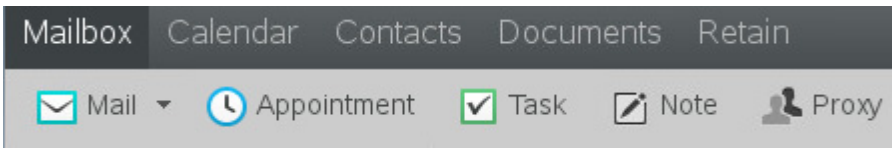
  mobile

  simple
- Restart Tomcat and the plugin installation is complete

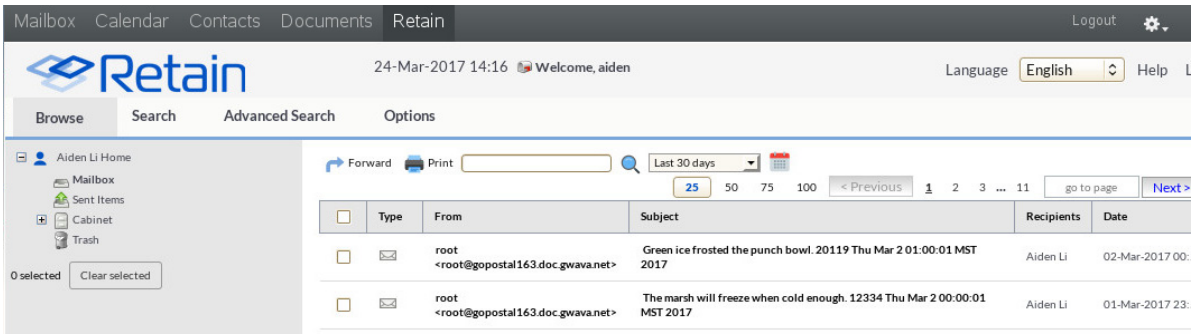
8.14.6 Confirmation

The plugin is now installed. Once WebAccess has been restarted, simply login as usual.

After logging in, users will see the interface they usually see, with an additional ‘Retain’ tab on the end of the top ribbon.



For users to access their own archive, all they need to do is click on the ‘Retain’ button and their archive will automatically open. Choose a folder to see the contents.



8.14.7 Uninstalling the GroupWise WebAccess Plugin

1. Stop Tomcat for the WebAccess Application
2. Navigate to the WebAccess directory
3. Rename or remove the retain/ directory
4. Navigate to \srv\tomcat\4\webapps\gw\WEB-INF\classes\templates\webacc\css
5. Edit the caption.htt file
6. At the end of the file look for the line that says:

{INCLUDE retain.inc}

1. Comment out or remove this line
2. Save the file.
3. Start tomcat

8.15 Retain for Social Messaging Gateway

8.15.1 Setting up the VM & Installing the OS

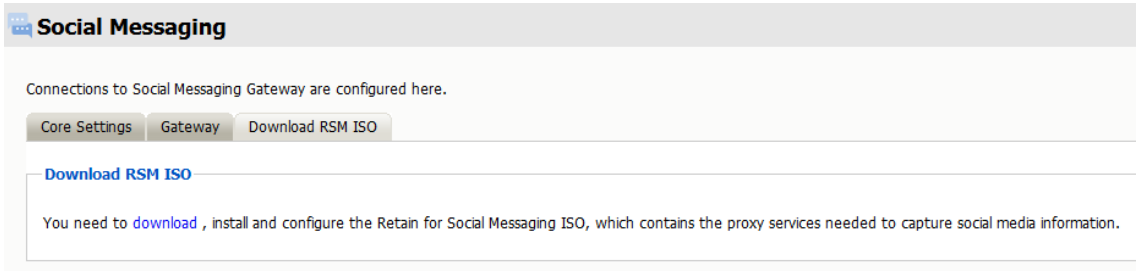
Before setting up the Virtual Machine

- Ensure the firewall is configured as such:
  - From RSM WAN IP to Untrusted, all TCP/UDP ports
  - If the RSM WAN IP is a private IP, it needs to be NATed to an appropriate routable IP address. The LAN IP address does not need a corresponding inbound NAT rule.

- **NOTE:** It is critical that these firewall rules are in place before proceeding.

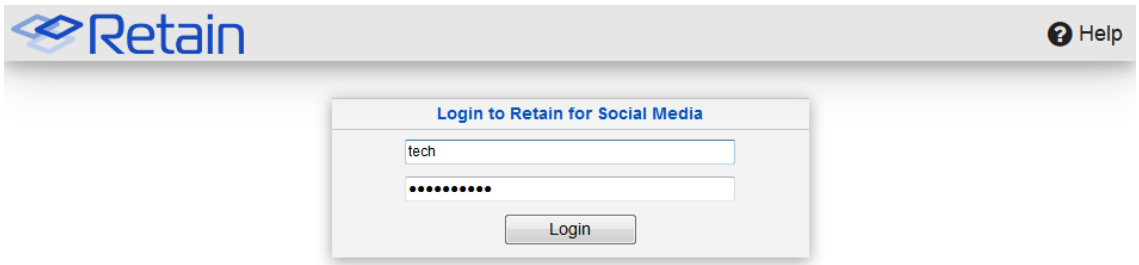
8.15.2 Configure the Virtual Machine

1. Download the ISO from the link found in the Social module page, on the ‘Download RSM ISO’ tab.

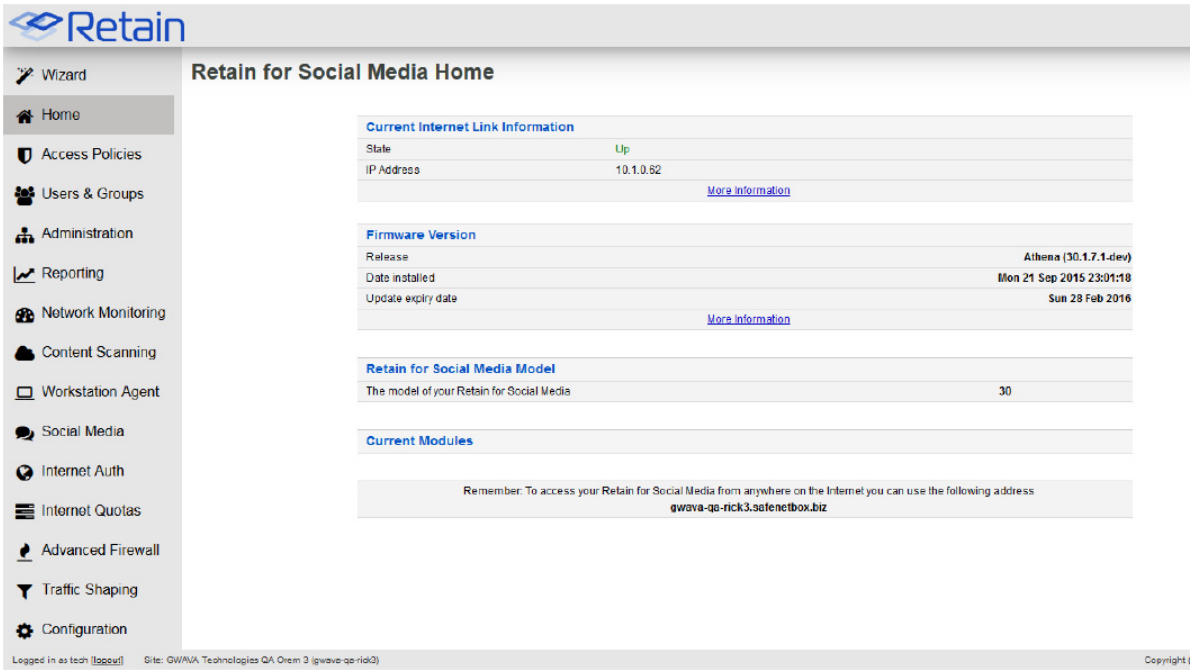


2. Install it onto a VM or to the desired machine with the following minimum specs:
   - 2GB RAM (minimum)
   - 2 CPU cores (minimum)
   - 60 HDD (minimum)

   The recommended RAM, CPU and HDD specification will be determined by the expected load and size of the organization being served. Consult your Micro Focus technical representative.
   - 2 NICs - vmxnet 3, ensuring the interfaces are in different VLANs. NOTE: The second NIC in the VMware settings list is the primary NIC that will be used for proxy traffic. This NIC will be labeled “Internet” on the RSM. The first NIC will be the “Local” port, and may be disconnected if desired when configured.
   - Redhat Linux enterprise 5 64bit base OS (selected from the dropdown list when setting up the VM – there is no need to actually install the Redhat OS)
3. Mount the ISO and start up the VM or server.
4. Set the password for the “tech” and “admin” users. This will then install the RSM OS. Once it’s installed and has come back up after the reboot, the RSM will have the default LAN IP of 192.168.0.254/24. The WAN interface will try to obtain an IP address from DHCP.
5. Set the LAN IP via the console by sending a ctrl-alt-ins and then follow the prompts.
6. All configuration operations are via the web user interface. Log onto the RSM interface by browsing to the LAN IP on either port 80 or 443 and logging in as the “tech” user, (the “admin” user doesn’t have the required permissions).



You will initially be prompted to confirm the EULA and provide your details. If the email address is not accepted at this stage (the RSM will try to validate all addresses), then use “root@ mail.RSM”.



1. Click the “Wizard” button from the left-hand navigation bar, and follow the prompts for setting the language, time zone and network details.
2. Once the network settings have been accepted, select Configuration -> Internet from the left-hand navigation pane, and check “Show Advanced Options”.
   - From the drop down list, next to “DNS server configuration” select “Always use static server” and then enter in the required DNS servers.
   - Click on Update.
   - Browse to
   - Configuration -> Apply and click on the “Apply Changes” button.

   NOTE: NOTE

   - If the LAN IP Address has been changed in the Wizard setup, you will need to browse to the new IP address and log onto the RSM interface again after applying the network settings.
   - The WAN (or Internet) Interface of the RSM will be used for proxy traffic.
   - The LAN interface is not used for production traffic, but can be utilized as a management-type port if required.
   - It is critical that the two interfaces are configured on different subnets from this point on, regardless if the LAN port is utilized or not.
3. After the Network configuration has been completed you will be prompted to supply a sitekey. A sitekey is a unique identifier for your RSM and it should be relevant to your organization. Only alphanumeric and “-“ are accepted, with a minimum character count of 3.
   a. Your RSM representative will supply the registration key to you.
   b. After the sitekey has been successfully applied, your RSM will reboot.
4. Log back onto the RSM and browse to Administration -> Updates and click on “Run Update Now”.

The RSM will go through its update process, which could take up to 30 minutes and reboot multiple times.

8.16 RSM Configuration

The configuration of the RSM is achieved via its Web User Interface. You need only browse to the (WAN or LAN) IP address and enter in the appropriate credentials. For the following you will need to use the “tech” account with the password that was set when creating the virtual machine. Specific instructions on how to use and configure the different features of the interface are accessed through the ‘Help’ screens inside RSM.

NOTE: When making changes to the RSM configuration, remember to update the page you are working on before navigating away from it. When you want to make your changes live, browse to the “Apply” sec-tion in the relevant module and click on the “Apply Changes” button.

8.16.1 Proxy Configuration

The following describes how to configure the RSM in a direct proxy setup. The client will then direct their browser to the WAN IP address of the RSM on the specified port (8080 for example).

1. Browse to Configuration -> Web Proxy.
2. In the drop down list next to “Direct Proxy Mode” select “Direct”.
3. Specify the port to use (8080 by default).
4. In the drop down list next to “Provide proxy on Internet interface”, select “yes”.
5. In the drop down list next to “HTTPS inspection”, select “Enabled for all traffic”.
6. Click on “Update”.
7. Browse to Configuration -> Apply and click on the “Apply Changes” button.
8. After 1-2 minutes, once the changes have been applied (when the yellow apply banner disappears), test the proxy by pointing your browser to the RSM WAN IP on the specified port and browse to a standard HTTP web site (HTTPS sites will be tested shortly).NOTE: You can view the proxy logs by browsing to Configuration -> Web Proxy and clicking the “View Web Proxy logs” link in the top right hand corner of the page.

8.16.2 RSM CA Certificate Installation

The RSM will be performing HTTPS inspection, meaning it will have visibility to view encrypted web traf-fic. This is critical to ensure that all Social Information Governance functionality is available. In order to do this successfully, the RSM CA Certificate needs to be installed on all client devices that will be using the RSM as a proxy service.

1. Download the certificate from <IP>/noauth/cacert.

   If using Microsoft Active Directory create a Group Policy Object

   - In GPMC open the relevant policies that apply to the Computers that need to have this certif-icate installed and navigate to Computer Config>Windows Settings>Security Settings>Public Key Policies and import the CA Certificate as a Trusted Root Authority.

NOTE: This GPO will NOT apply to Safari or Firefox browsers. The certificate will need to be manually installed.

---

2. If manually installing the certificate ensure it's installed as a Trusted Root certificate.
3. Once the certificate has been deployed, test by browsing to a HTTPS site and view the web proxy logs. You should not see a certificate warning in the browser and the proxy logs should display your HTTPS request.

8.16.3 Social Information Governance Configuration

Social Information Governance is a Rule/ Action process, meaning that a rule needs to be triggered before an action takes place. A rule is made up of one or more criteria. A criterion tells the RSM what should trigger the rule. A rule has one action assigned to it. An action may have multiple sub-actions.

Actions can be either:

- Block
- Alert
- Modify
- Moderate
- Or in some cases, combinations of the above (i.e., Block and Alert) For example a rule to block a staff member from posting extreme profanity on twitter will contain the following:
  - A criterion that uses the "extreme profanity" pattern list.
  - Another criterion that sets the Application type to "Twitter".
  - A final criterion that specifies the user's action as "send"
  - An action to block.

NOTE: All Criteria must be true before the rule is triggered. In other words: If the user sends a post to Twitter, which matches an entry in the "extreme profanity" pattern list, then take the associated action to Block. This rule would not trigger if the user read extreme profanity on Twitter, so care must be taken in establishing the Social Information Governance requirements for your organization. The RSM comes with predefined rules that may apply to the relevant organization type.

8.16.4 Enabling Social Information Governance

1. Enable Social Information Governance by browsing to Content Scanning -> General. Under the SafeChat Settings table:

   a. From the dropdown list next to "Enable SafeChat scanning of web content?" select "Yes"
2. Click on "Update"
3. Browse to Content Scanning -> Apply and click on the "Apply Changes' button.

8.16.5 Creating Pre-Defined Social Information Governance Rules and Actions

1. After the changes have been applied create the predefined rules by browsing to Content Scanning -> General and click on the "Add suggested settings" button.
2. Browse to Content Scanning -> Rules where you can see the rules that have been created. The rules and corresponding actions will be currently disabled.

3. Enable the rules by
   - Clicking “edit” next to the rule.
   - Check “Enabled”.
   - Additional criteria can be added at this stage by clicking on the “Add Criteria” link.
   - Click on “Update”.
4. When you are happy for the rules to come into effect, browse to Content Scanning -> Apply and click the “Apply Changes” button.

8.16.6 Creating Custom Social Information Governance Rules

To add or modify a rule:

1. Browse to Content Scanning -> Rules and click on “Add new rule” in the top left hand corner
2. Enter a descriptive name for the rule into the textbox provided. This will be used for referring to this rule, for example when modifying or deleting this rule.
3. Select the action to be run when the rule is matched.
4. Click the Update button to save the rule.
5. Enter criteria as required. The rule will be triggered when all the listed criteria are matched.

   NOTE: If no criteria are specified then this rule will never be triggered.

6. Select the checkbox provided at the top of the page to enable the rule, once you are happy for it to become active.
7. Click the Update button to save the rule.

8.16.7 Creating Custom Social Information Governance Actions

To add or modify an action:

1. Enter a descriptive name for the action into the textbox provided. This will be used for referring to this action, for example when modifying, deleting or using this action.
2. Click the Update button to save the action.
3. Enter sub-actions as required.

   NOTE: If no sub-actions are specified then this action will not do anything, but the rule will still trigger a match, this can be useful for reporting.

4. Select the checkbox provided at the top of the page to enable the action, once you are happy for it to become active.
5. Click the Update button to save the action. Once all changes have been made to the
6. Content Scanning section:
   a. Browse to Content Scanning -> Apply.
   b. Click on the “Apply Changes” button.

The Configured Social Information Governance rules and actions will now be live in a few seconds.

8.16.8 Social Information Governance Moderation

The Moderation section of Content Scanning allows authorized users to approve or deny messages held by a “Hold for Moderation” sub-action. A Content Scanning action needs to be configured to hold mes-sages in a particular moderation “queue” until processed. An authorized user can then view each queue, reviewing the pending messages of a queue in detail, and selecting whether each message is approved or denied. An example would be to hold for moderation any messages that contain the name of the com-pany for moderation and review before it’s sent to Twitter, Facebook, etc. The authorized user can then elect to allow that message to be sent or reject the message, giving an explanatory comment if necessary.

NOTE: Please refer to the online help under Users & Groups to create individual users and groups if you want to limit access to these features a defined group of users. Social Information Governance Modera-tion requires the configuration of:

---

1. A moderation queue and;
2. An action to hold messages in the moderation queue.

8.16.9 Creating a Moderation Queue

1. Browse to Content Scanning -> Manage Moderation Queues.
2. Select “Create New Queue”.
   - Provide a descriptive name for the queue.
   - Select which groups (if any) can view this queue.
   - If an explanatory comment is required for any actions within this queue (i.e., releasing or denying), set the appropriate value here.
3. Click on “Update”.

8.16.10 Creating a Moderation Action

This is simply a Content Scanning action to hold messages into the specified moderation queue.

1. Follow the steps described under “Creating Custom Social Information Governance Actions”.
2. When selecting the sub-action, select “Hold for Moderation”.
3. Select the moderation queue from the drop-down list.
4. Assign this action to the required rule (please refer to “Creating Custom Social Information Governance rules” for more detail).

8.16.11 Configuring Secure Social Media Authentication

The RSM can be used to grant users access to social media accounts without having to divulge the account password. The RSM will determine whether users can access particular social media accounts based on the group membership of the user. This can be used, for example, to allow certain members of your organization the ability to make posts to your organization’s Twitter or Facebook account without divulging the password for those accounts.

NOTE: This functionality currently applies to LinkedIn, Twitter and Facebook accounts.

8.16.12 Configuring RSM to authenticate to a company's social media site

1. Browse to Social Media -> Accounts and click on "Configure New Account" in the top right hand corner of the page.
2. Configure the account options:
   - Select the application.
   - Type in the username and password for the account

     NOTE: This feature will not work with accounts that require two-factor authentication. It is recommended that a long, complex and very secure password be used for the account.

   - Select which user groups can access this account through the RSM.

     NOTE: Users with Social Media Account administration privileges are allowed to access all social media accounts regardless of this setting.

   - Provide a description of the account (i.e., "ABC Corp. Main Facebook page")
   - Set a session lifetime. This determines how long the login session cookie will be active before timing out. Leave it blank to use the application's default values.
3. Click on "Update".
4. Browse to
5. Content Scanning -> Apply and click on the "Apply Changes" button.
6. Users can access the company's social media sites by browsing to the RSM, supplying appropriate credentials and then browse to
7. Social Media -> Accounts.
   - They will see the list of social media accounts they have permission to access (for example Facebook and Twitter, but not LinkedIn).

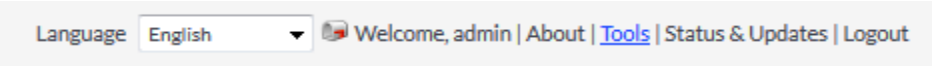
Click on the social media account desired, and the RSM will log into that application for the user.

8.17 GroupWise Client Plugin

The GroupWise Client Plugin provides access to the logged-in user's Retain Archive from the GroupWise Client interface. The plugin adds a Retain Menu item as well as a Retain button along the top ribbon for easy web access to the archive. These buttons will open a new browser window and log the user into Retain automatically. Additionally, the plugin may hide the native GroupWise Archive option if selected during the install.

To enable the GroupWise client to search the Retain archive from within the GroupWise client itself, the administrator would need to enable stubbing.

The GroupWise Client Plugin is found on the tools page. To access the tools page, select the 'tools' link from the top right of the Retain Server administration page.



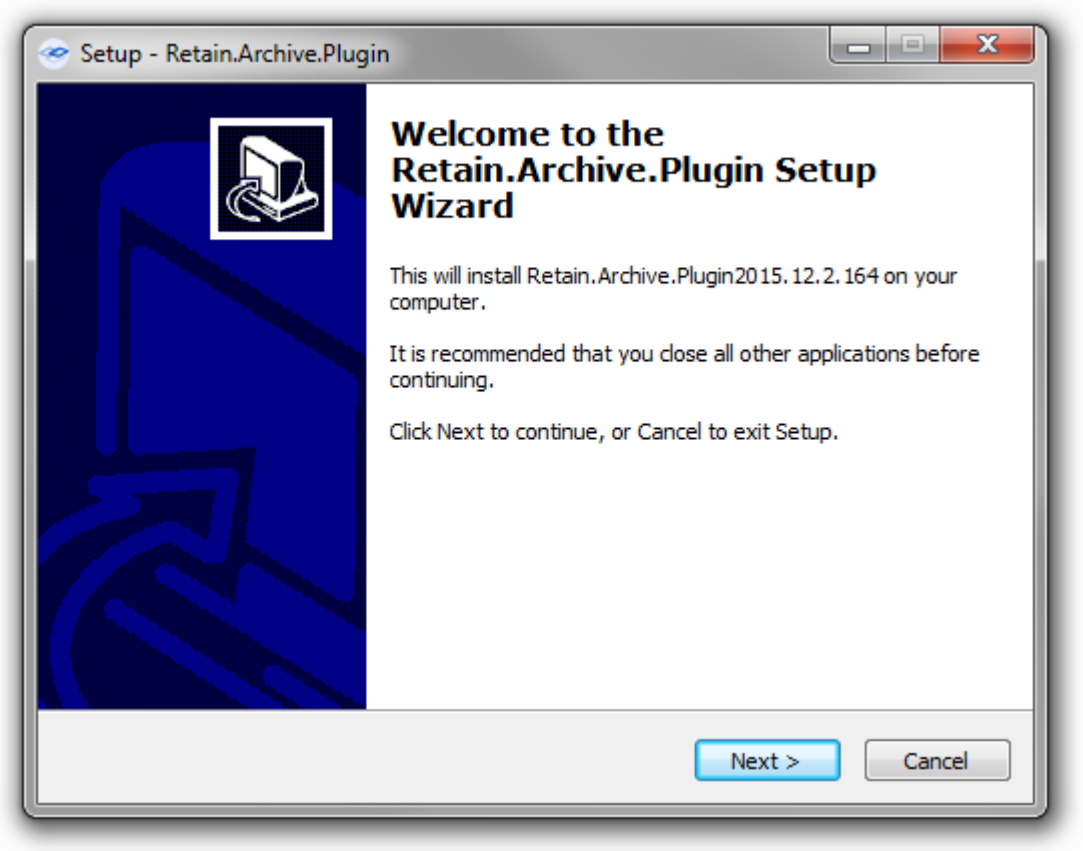
8.17.1 System Requirements

- GroupWise Win 32 Client
- Retain Server 4+

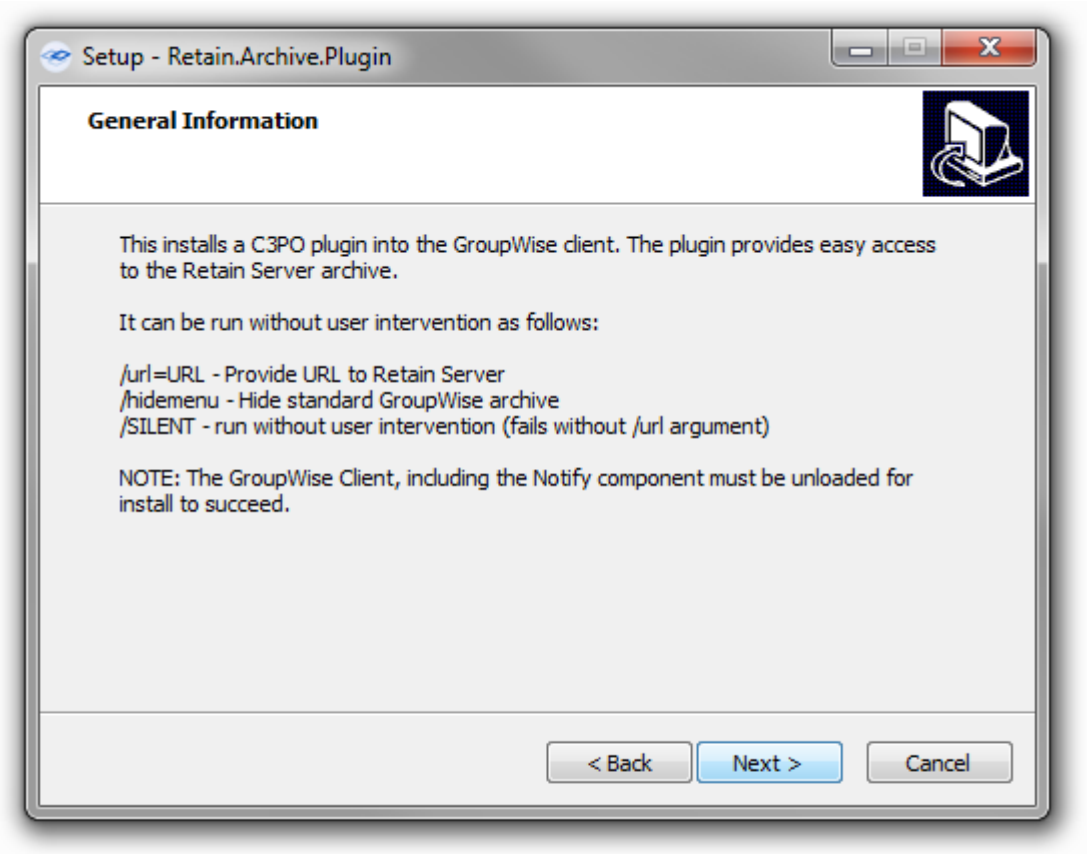
NOTE: Because the plugin provides transparent single sign-on features, installing the client may introduce some security concerns on a public workstation. It is not recommended for installation on any public workstation.

8.17.2 Install

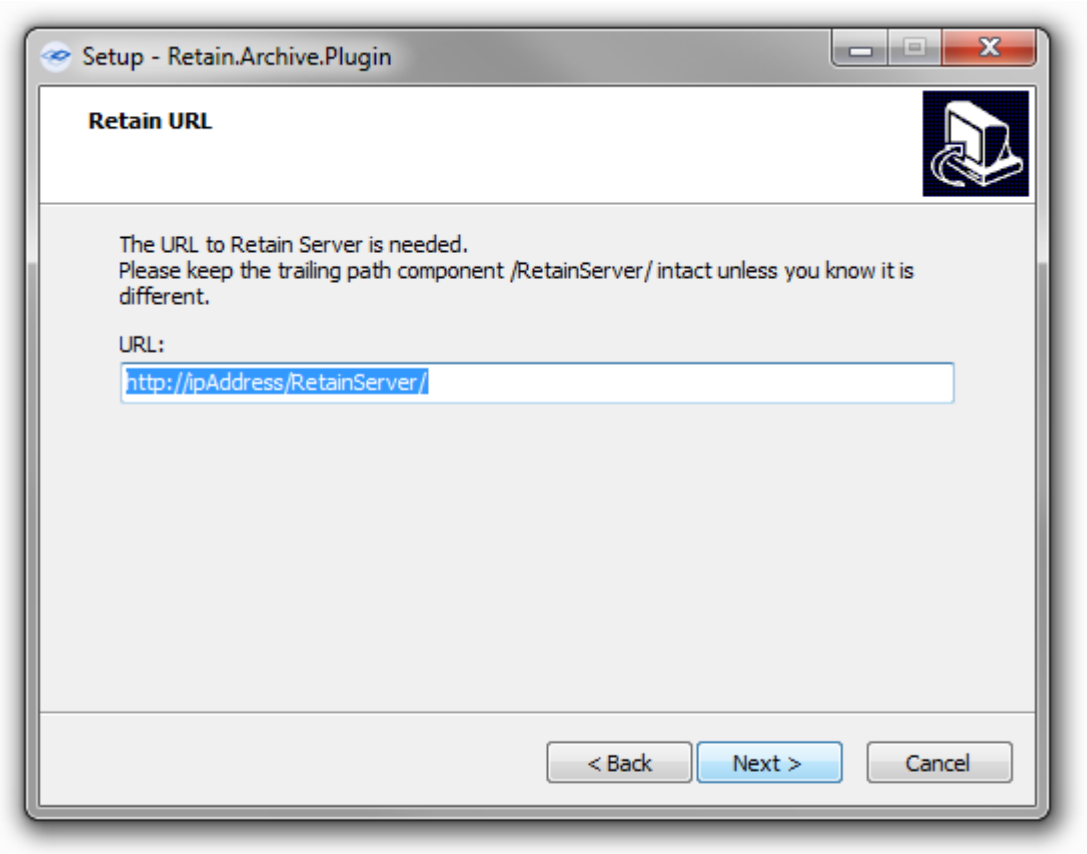
To install the plugin, download the plugin to the desired workstation and run the executable file.



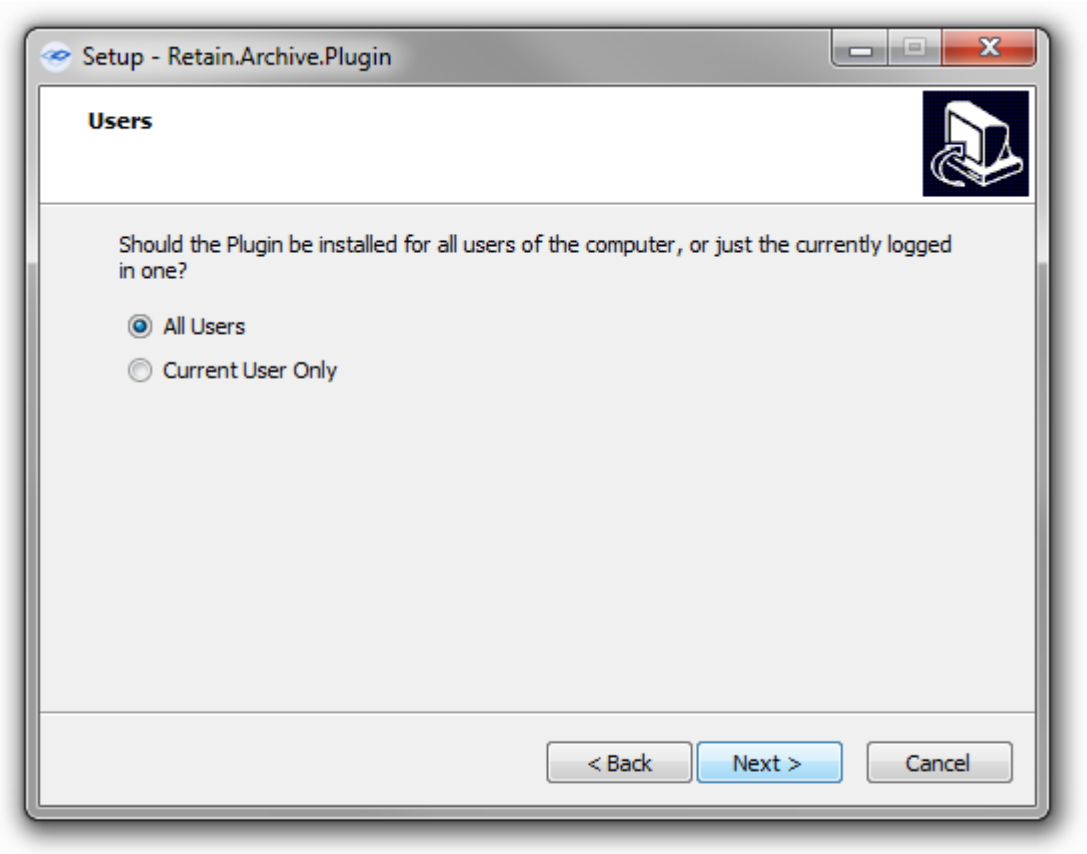
The installation gives an opening summary. Click 'Next' to continue.



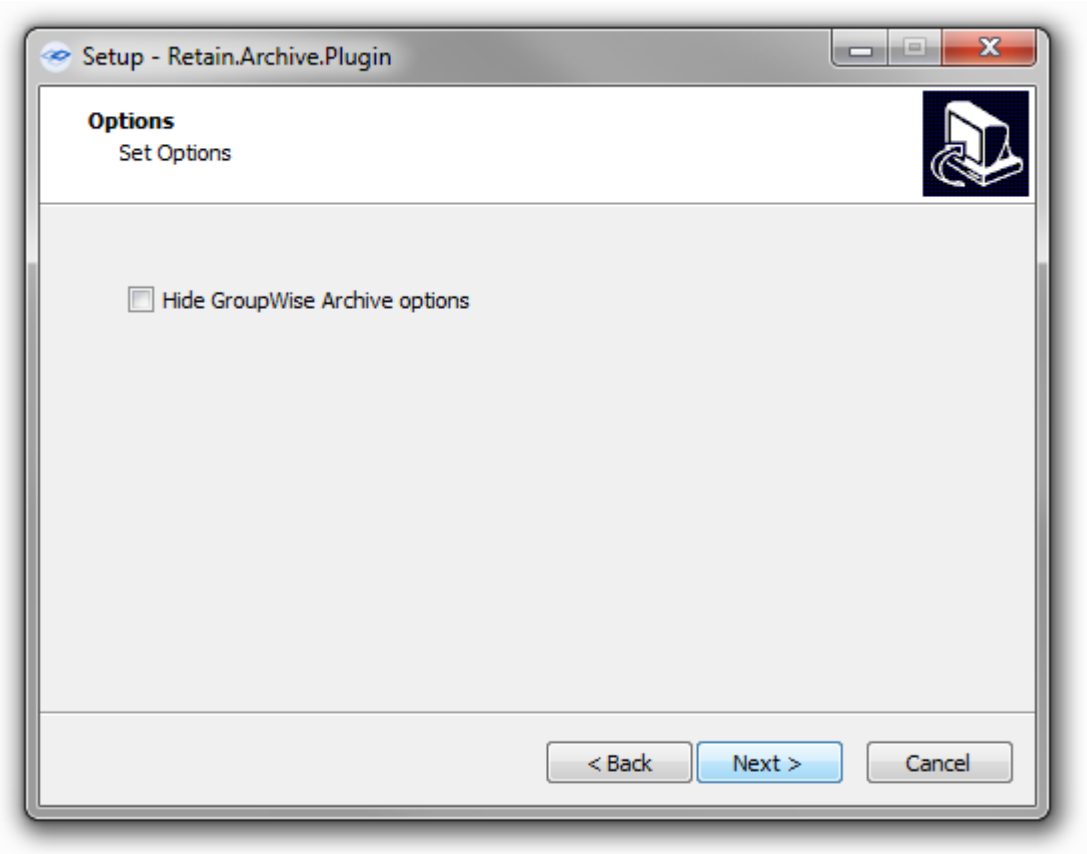
An installation summary is displayed. If a silent install is desired, the install may be performed with the following commands.



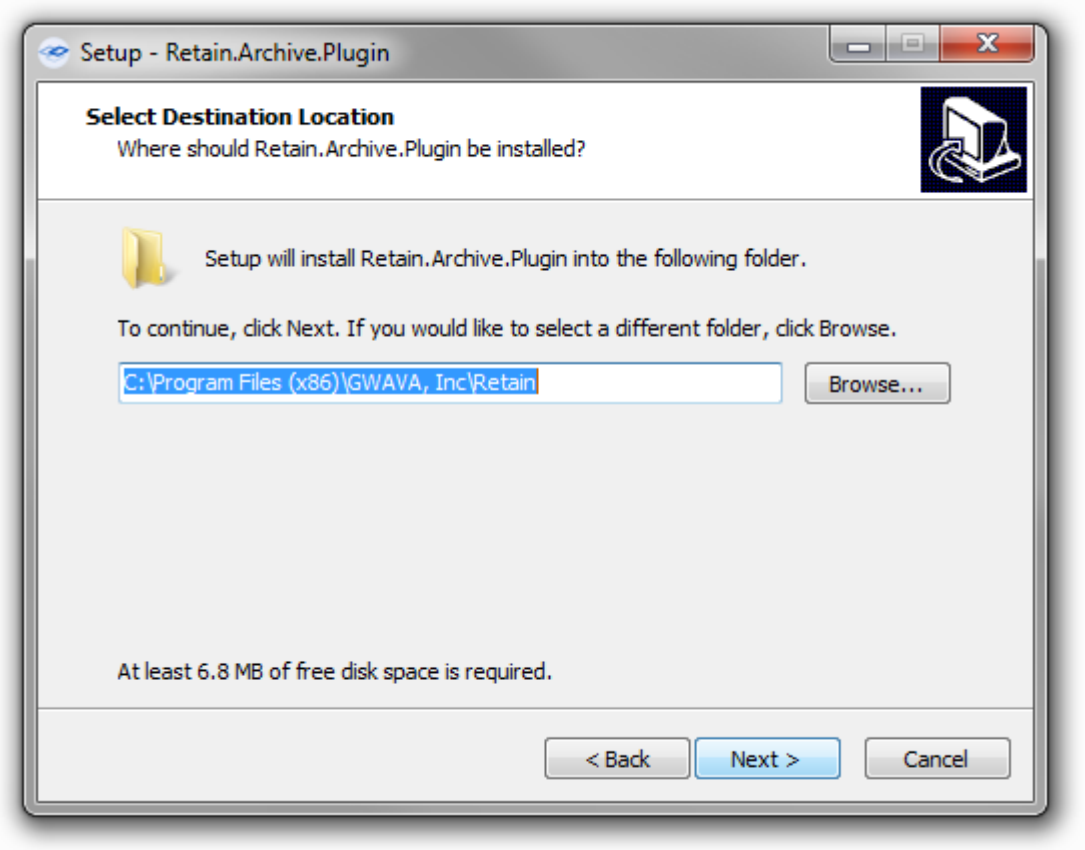
The Retain URL is the connection URL for the Retain Server. Check the URL to make sure the URL works before continuing. If the URL is incorrect the plugin will not function correctly.



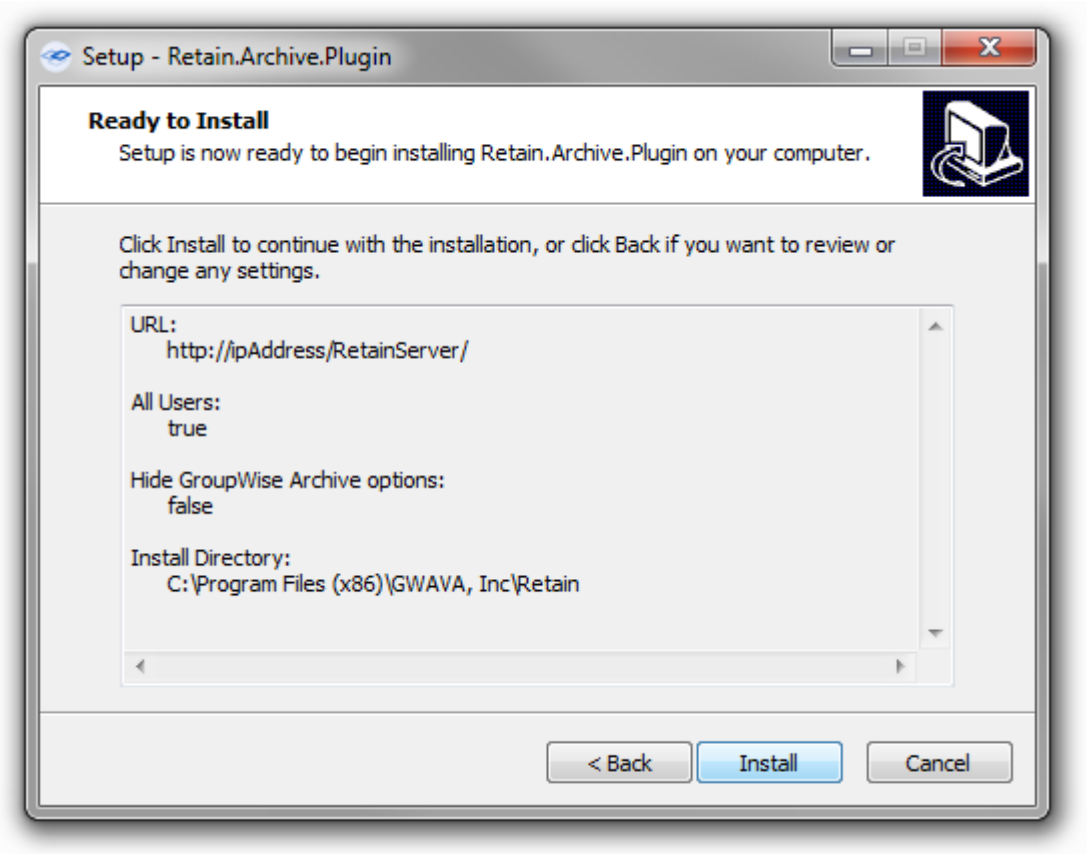
Select whether to install for only current, or all users on the system.



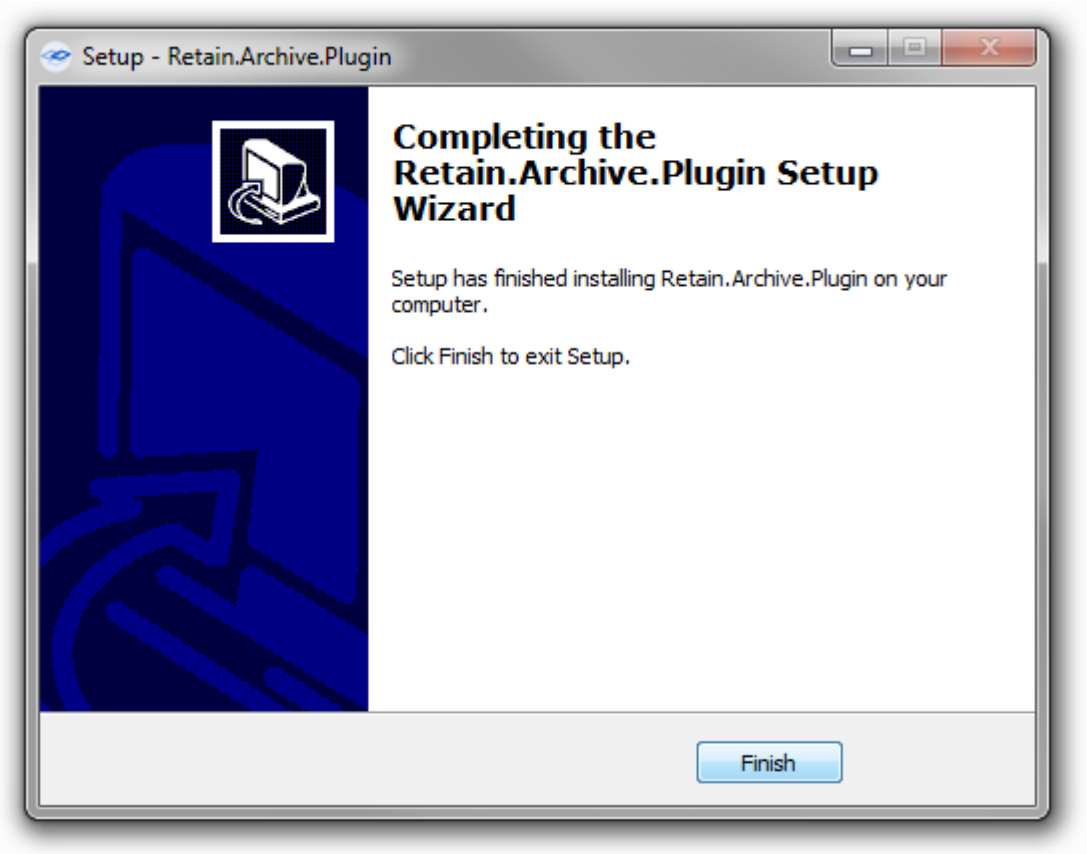
The Retain plugin may be used instead of the GroupWIse Archive. To hide the GroupWise Archive in favor of the Retain Archive, select the box and click 'Next' to continue.



The install location for the plugin is displayed. The location may be changed. Default is shown.

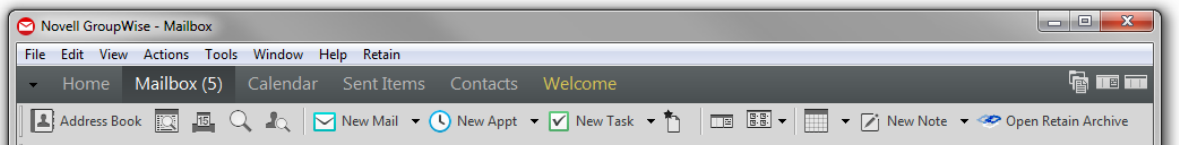


A final summary of the selected settings is displayed. To change any setting, utilize the 'Back' button. Otherwise, select 'Install' to continue.



Once the install has been completed, the GroupWise Client needs to be started, or restarted, to load the plugin.

Once started/restarted, login to the Client as usual.



The Retain plugin will have a Retain menu item as well as a Retain Archive button on the top ribbon. To access the Retain Archive, a user simply selects the 'Open Retain Archive' button from the ribbon, or the same item from the Retain menu from the top bar. Once selected the user's Retain Archive will open in a new browser window.

8.18 GroupWise WebAccess

The GroupWise WebAccess plugin adds a Retain link inside WebAccess which links directly to the Retain archive.

8.18.1 Requirements

- Java 1.7+
- File and directory access to the WebAccess server (for install only).

8.18.2 Install

Windows

Install on Windows is simple. If installing from a workstation, ensure that the Novell client is logged in with an account which provides access to the WebAccess directory on the server. File access is required for installation.

- Download the tool from the tools page in the Retain Server
- Unzip and locate the install.bat file
- Run the install.bat file either from the server, or from a workstation with file access to the server

Linux

- Download the tool from the tools page in the Retain Server (It is recommended to run the Linux install on the WebAccess server; file access is required.)
- Unzip and browse to the location with a terminal session
- Make the install scripts executable by running:

  chmod +x *.sh

There are two install files to choose from: installgui.sh and installtext.sh. Both installers result in a full installation. Which install is used depends on the desired installation interface; graphical or text based.

- If the graphical installation is desired, run the 'installgui.sh' executable in a terminal session initiated from the graphical interface. Run:

  ./installgui.sh
- If a text based installation is desired, the installation may be executed from any terminal session. Locate and execute the following file:

  ./installtext.sh

Install

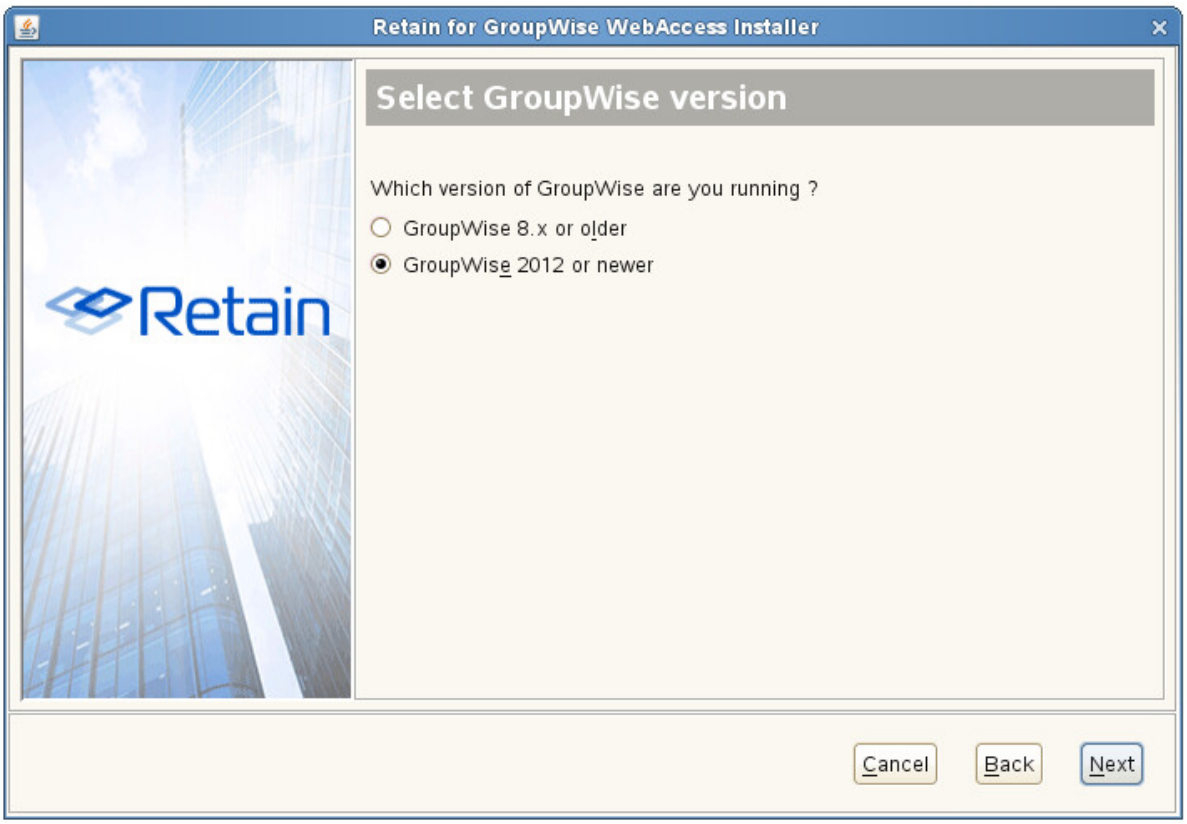
The installation is a java script which is identical between both platforms. Text-based installation prompts mirror the graphical installation's information and requirements. Graphical installation prompts are shown below.



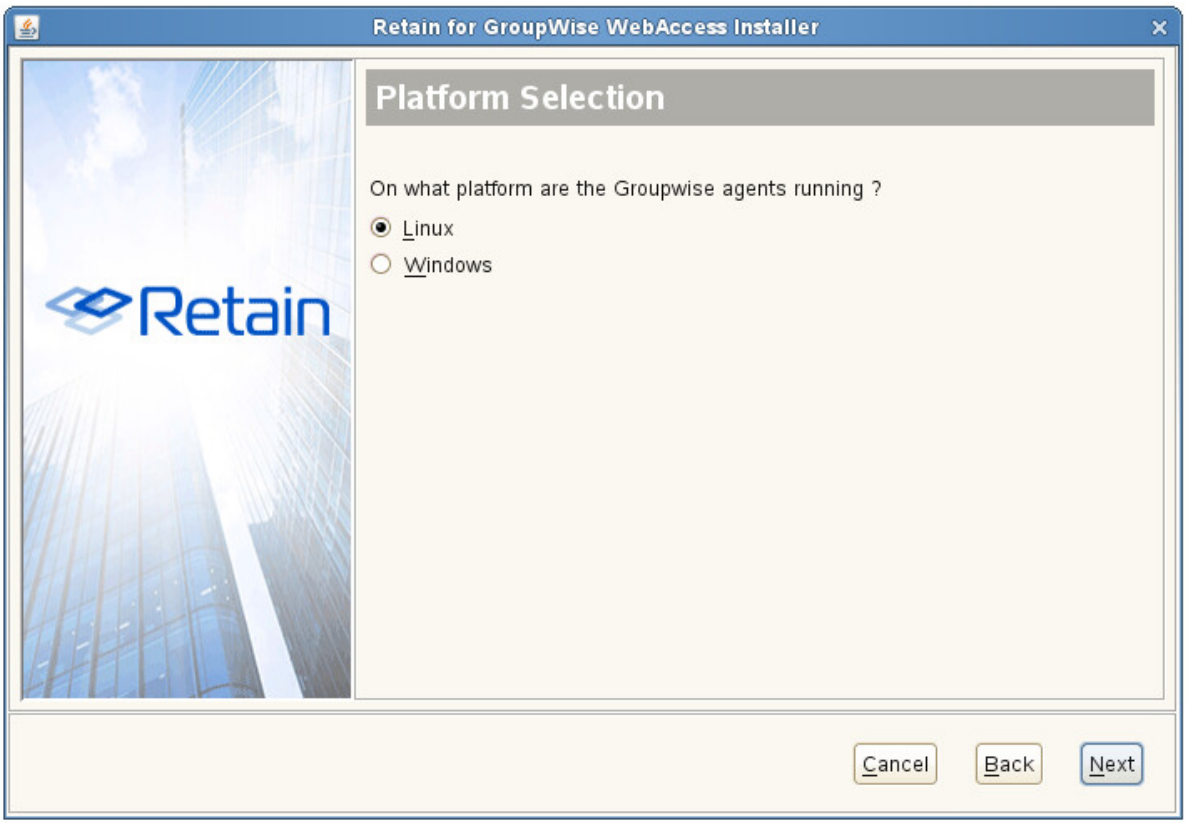
When the installation is initiated, the splash screen is displayed.



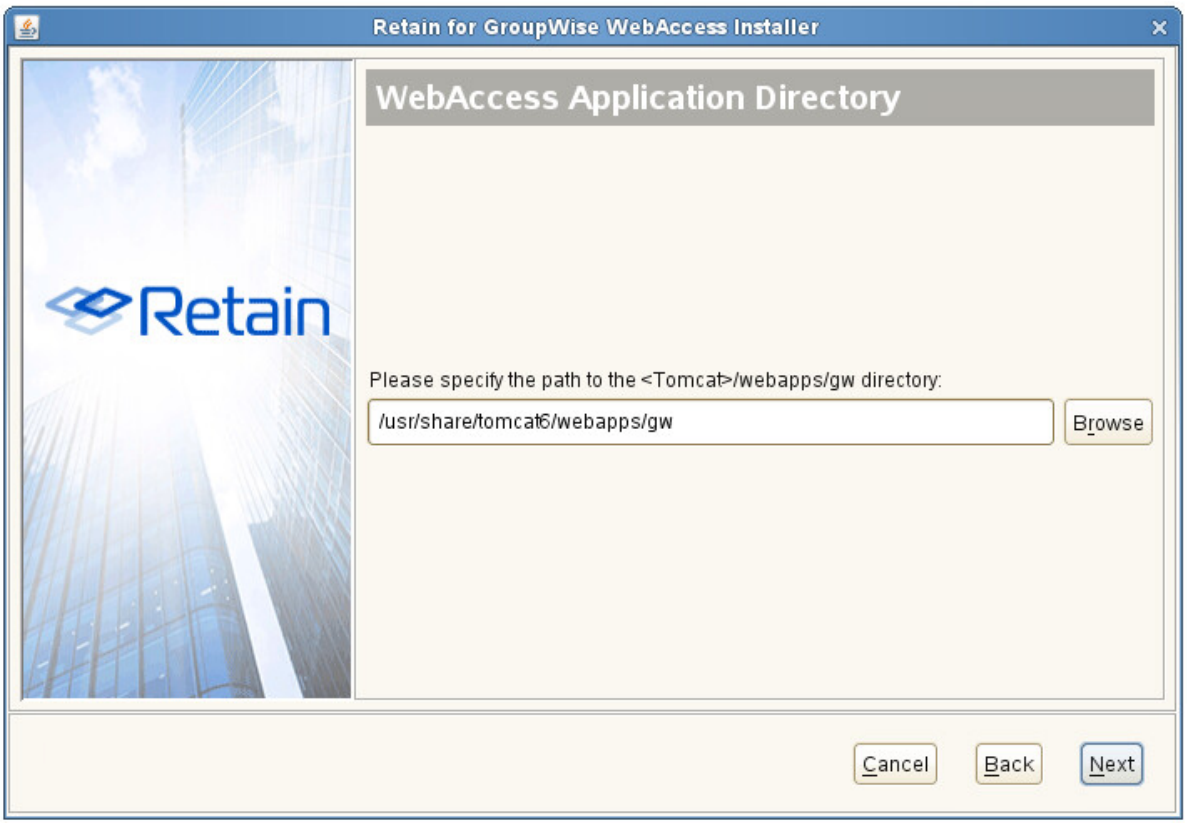
Agree to the license agreement to continue.



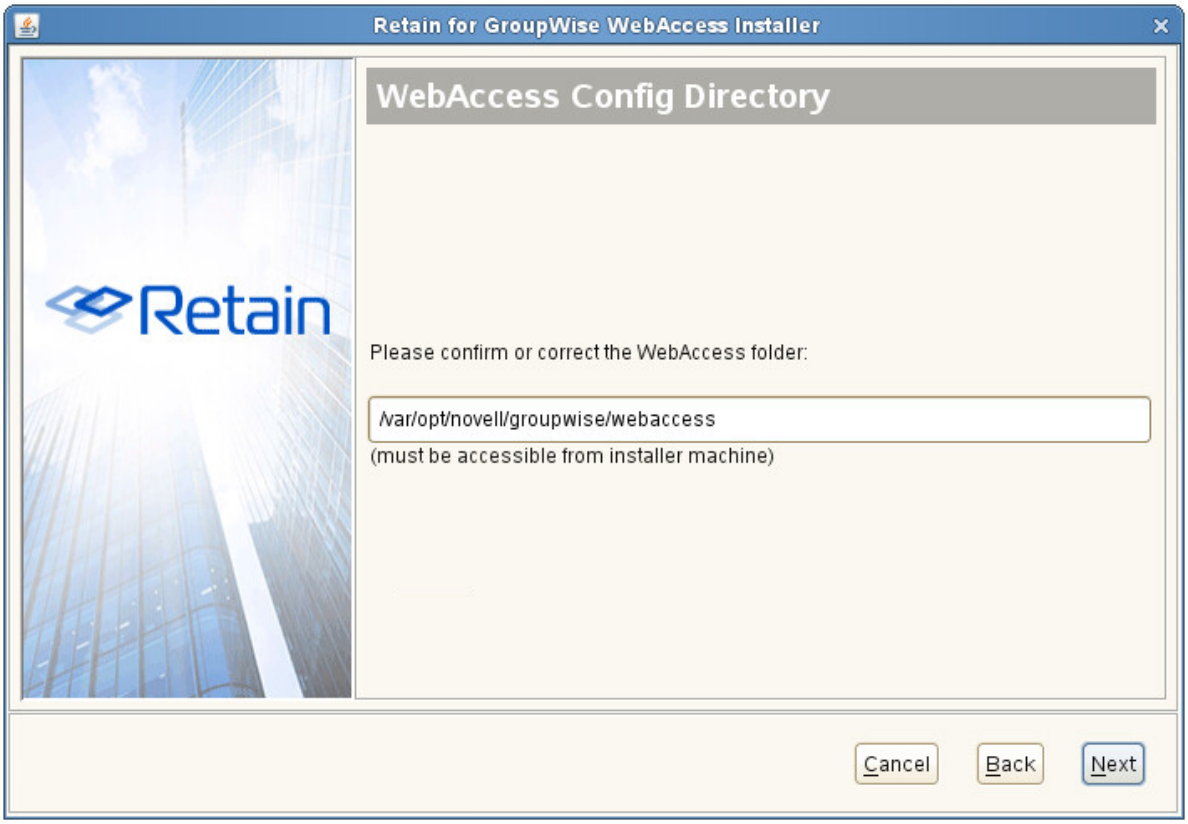
Specify the version of GroupWise WebAccess.



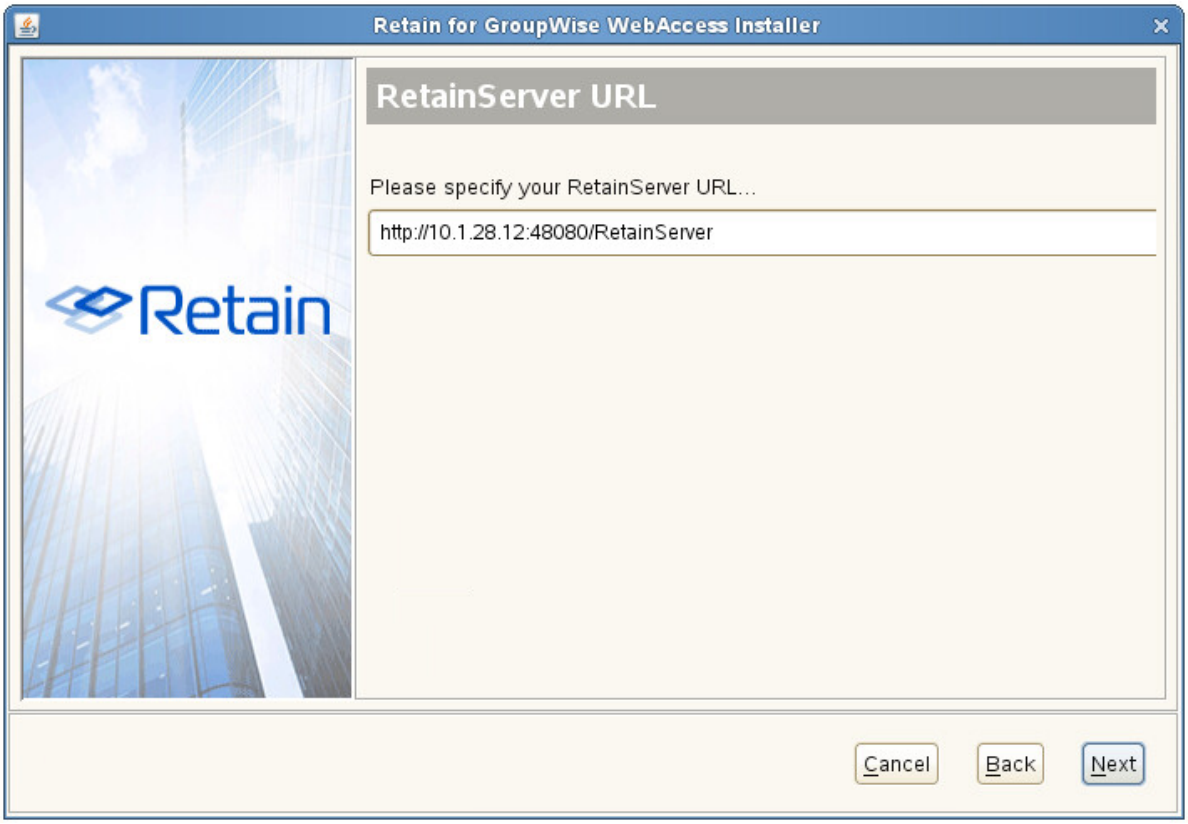
Specify the platform of the GroupWise WebAccess sysetm.



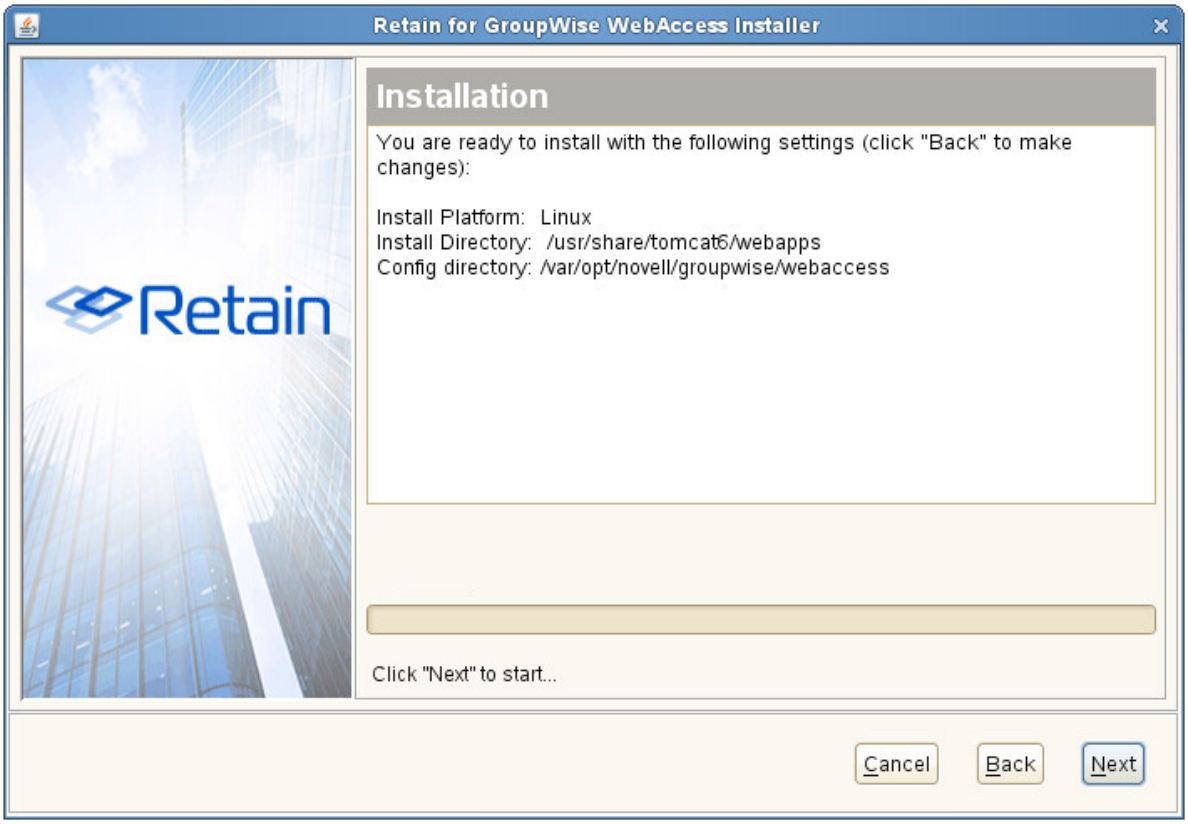
Specify the path to the Tomcat GroupWise directory. (<Tomcat>/webapps/gw)



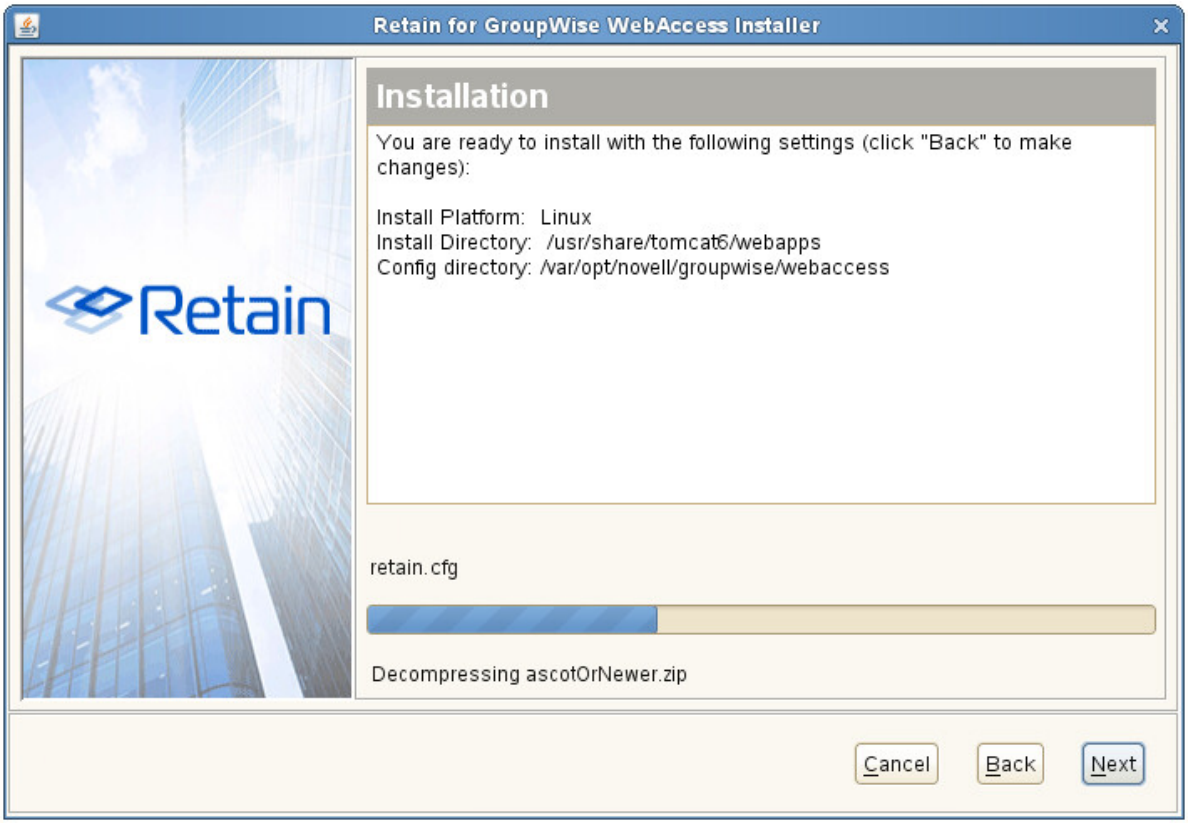
Specify the location of the WebAccess folder on the server.



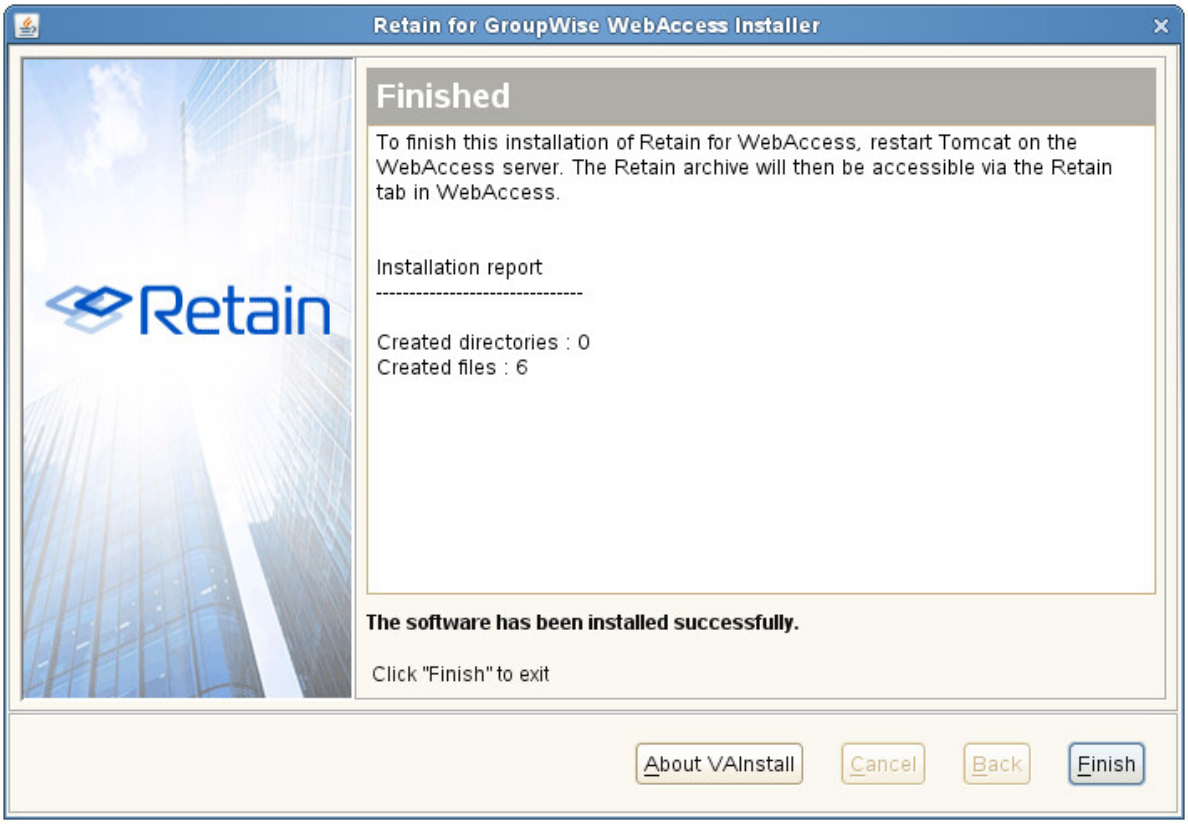
Specify the Retain Server URL. Use a DNS name if possible. (NOTE: If SSL has been configured on the apache or IIS server housing Retain, omit the port number in this field. ie: http;//10.1.28.10/RetainServer)



An installation summary is displayed. If any changes are desired, use the 'Back' button to make changes. Otherwise, click 'Next' to initialize the installation.

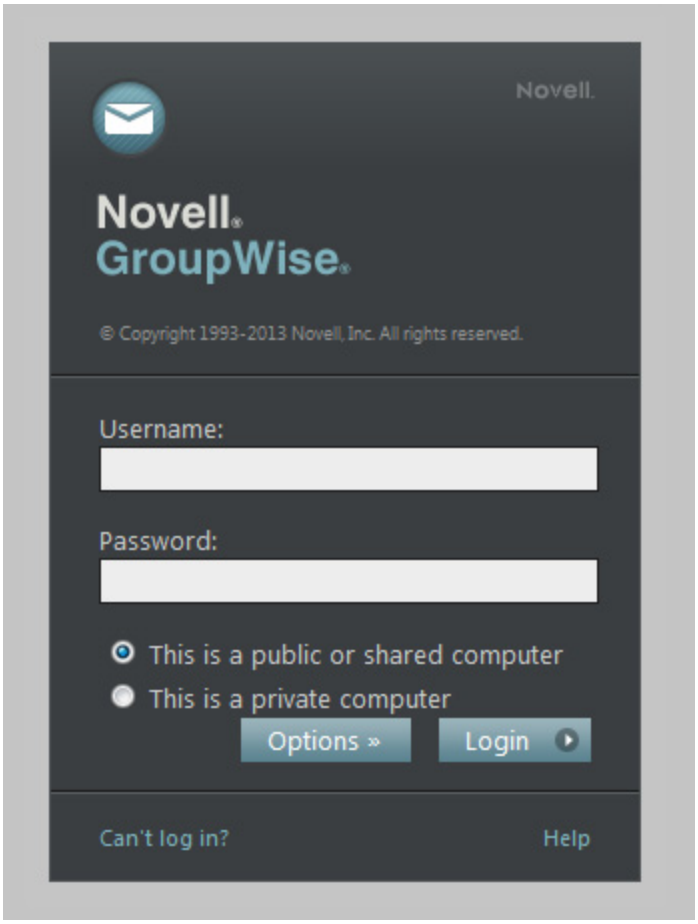


Installation proceeds as stated.

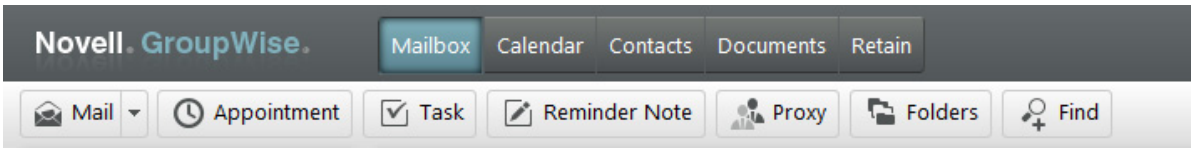


To complete the installation, restart WebAccess tomcat to load the plugin applet.

Once restarted, browse to the WebAccess login page.



Log into WebAccess as normal.



Clicking the Retain button in the top ribbon opens the user's Retain archive inside the existing GroupWise window.

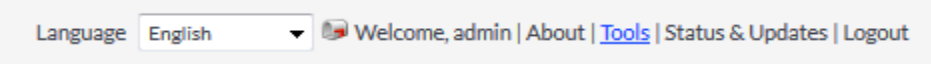
NOTE: If GroupWise WebAccess uses a secure, HTTPS connection, the plugin will require the Retain Archive to utilize a secure connection as well. Additionally, when prompted by the install for the RetainServer URL, omit the port, (:48080), and this will direct the link for Retain in WebAccess to point to the SSL-enabled webserver, (apache or IIS) on the Retain Server rather than directly pointing at Tomcat, which is not SSL-enabled.

8.19 GroupWise Archive Migration Tool

8.19.1 GroupWise Personal Archive Migration Tool

This migration tool is designed to import native GroupWise archives directly into Retain.

The GroupWise Archive Migration Tool is found on the tools page. To access the tools page, select the 'tools' link from the top right of the Retain Server administration page.



There are two modes:

1. Run at all workstations mode.

   From a command line, you'll run this and it will log into the users mailbox, and then open the local archive. It will then feed messages directly to Retain.

2. Centralized run.

   This only works if you have all archives stored at a central point.

   In this case, one can conveniently run all of the migrations from a single workstation, and not involve the users at all.

8.19.2 System Requirements

- A Windows workstation running GroupWise client 7.0.1 or up. GW 2014 and up use a slightly different file structure so there may be errors.
- Access to the archives you want to migrate.
- An HTTP connection to Retain.

8.19.3 Overview; Modes of Operation

The migration tool uses the GroupWise client to access the archives you want to migrate.

- In single-user mode, the migration tool will access the mailbox of the logged-in GroupWise user and it will read the archives available to the GroupWise client and send them to the Retain Server you specify.
- In multi-user mode, it will use the trusted application key to log in to the mailboxes of the users specified in a chosen GroupWise distribution list, attempt to access the archives defined in their mailbox and will send them to the Retain Server you specify.

8.19.4 How it Works

The migration tool does not access the archive files directly. Due to the way GroupWise stores its archives, you need the GroupWise client to read the archives and to hand the data to the migration tool. The migration tool then sends the archives to a Retain Server, much like a worker would.

In normal operation, a Retain Worker reads data from the live GroupWise system and sends it to a Retain Server. With the migration tool, the tool itself is a worker that reads the archived data from the Group-Wise client and sends it to a Retain Server.

The migration tool reads the path to the archives as defined in the GroupWise mailbox that you're accessing. That path must be accessible by the workstation you're on. It will then ask GroupWise to open the archives and it will read the items in the archive and send them to the Retain Server you have speci-fied.

If you run in single-user mode, the migration tool will access the mailbox currently logged into by the GroupWise client. If none is logged in, GroupWise will be opened and you will be asked to log in.

If you run in multi-user mode, the migration tool will open GroupWise, you can be logged in as any valid user, and it will use the trusted application key to log in to each user specified in a GroupWise distribution list, one at a time.

8.19.5 Prerequisites

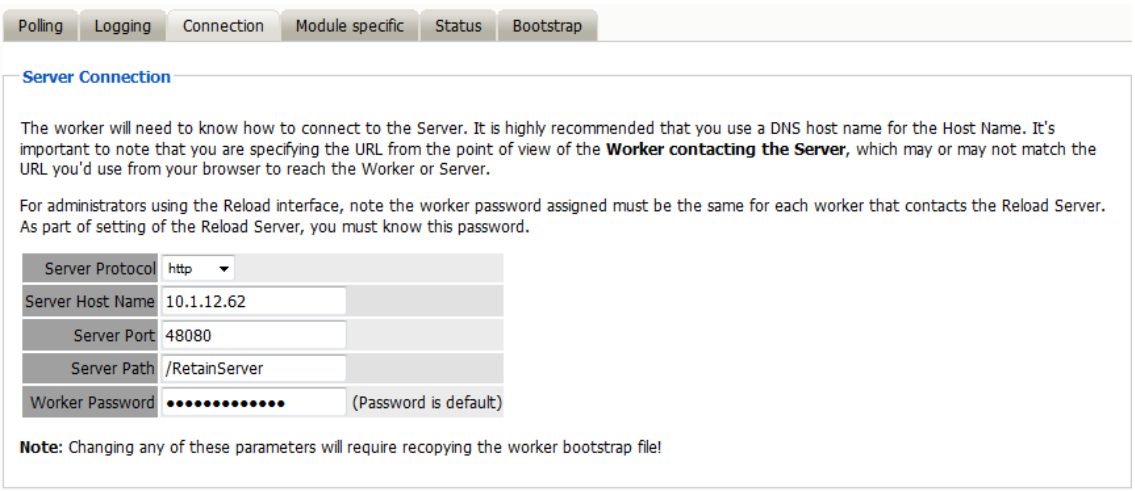
The GroupWise client on the workstation you're using must be able to open the archives you wish to migrate. The process depends on GroupWise being able to access these archives. If you're migrating the archives of another user, as you would be doing in multi-user mode, you must be able to access that user's archives from the workstation you're on.

It is not enough that the archives are on a network-accessible volume. GroupWise must be aware of the path to these archives and that path must be available from the workstation you're going to use for migration. Refer to the error handling section for important configuration.

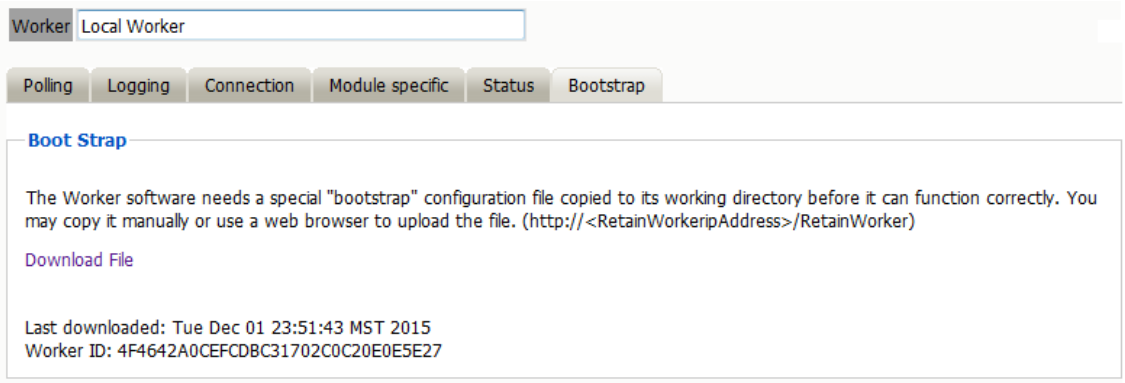
8.19.6 Setting It All Up

In multi-user mode, you can use one workstation to migrate the archives of multiple users. In so doing, you can migrate the data without affecting the user(s) in question. In single-user mode, you will run the migration tool on a workstation that has access to the archives you want to migrate (typically the work-station of the user whose archives you're migrating) and you'll migrate just one mailbox.

1. Copy the migration tool files to any directory accessible to your workstation
2. Log in to the Retain Server you want to migrate the archives to Create a new worker. You're not actually installing a new Retain Worker anywhere; you're creating a new bootstrap file for the use of the migration tool so it knows which Retain Server to send the archives to. Be sure the set-tings listed are accessible and usable by the workstation where you will be running the migration tool. The migration tool will use these settings to log into the Retain Server



3. Once you've saved changes, you will be able to save the bootstrap file. Save it in the directory of the migration tool.



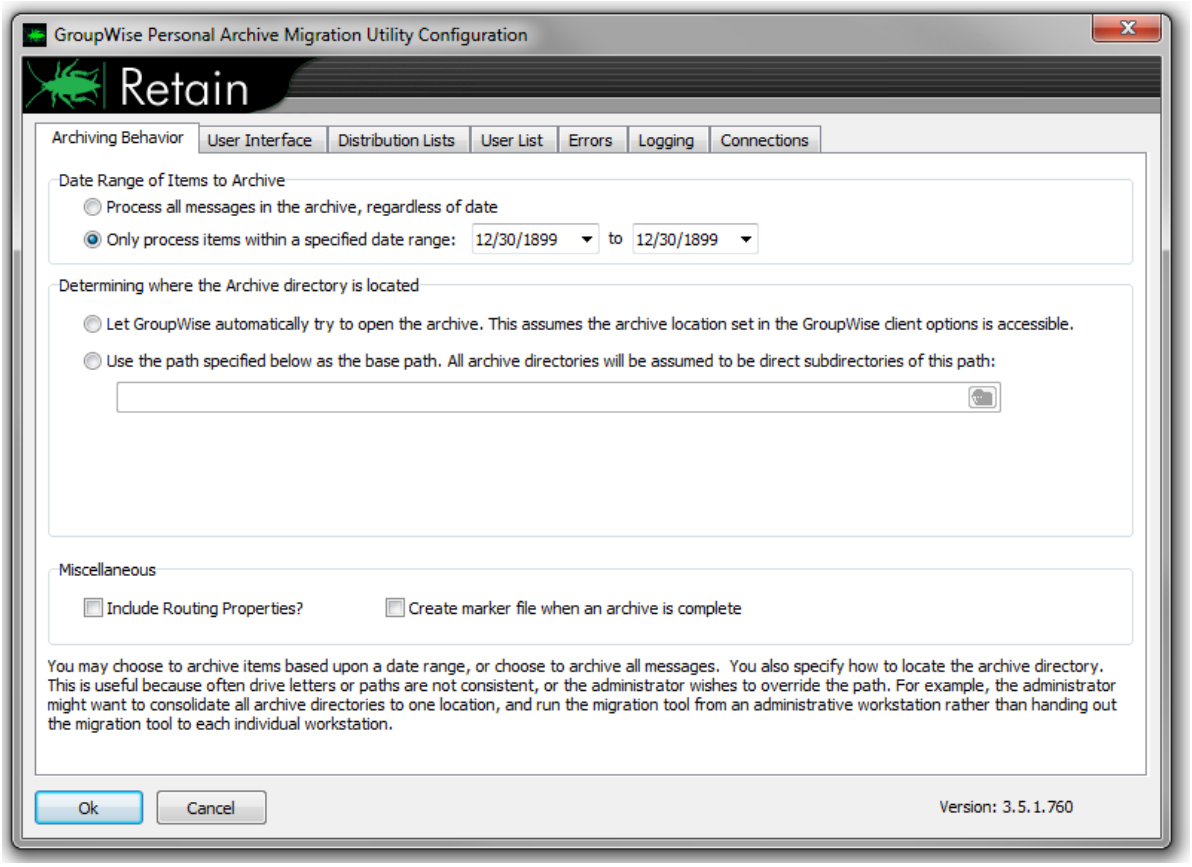
4. Now start up the migration tool configuration utility. A window will open on top of the configuration utility requiring the location to find the bootstrap file you have just saved.

```
vm-sles10-retain:/retain # l /retain
total 1
drwxr--r--  3 tomcat root  72 Mar  7 16:06 ./
drwxr-xr-x 25 root   root 608 Mar  7 16:01 ../
drwxr--r--  2 tomcat root  48 Mar  7 16:06 archives/
vm-sles10-retain:/retain #
```

(This window will also accept saved configurations from previous sessions.)

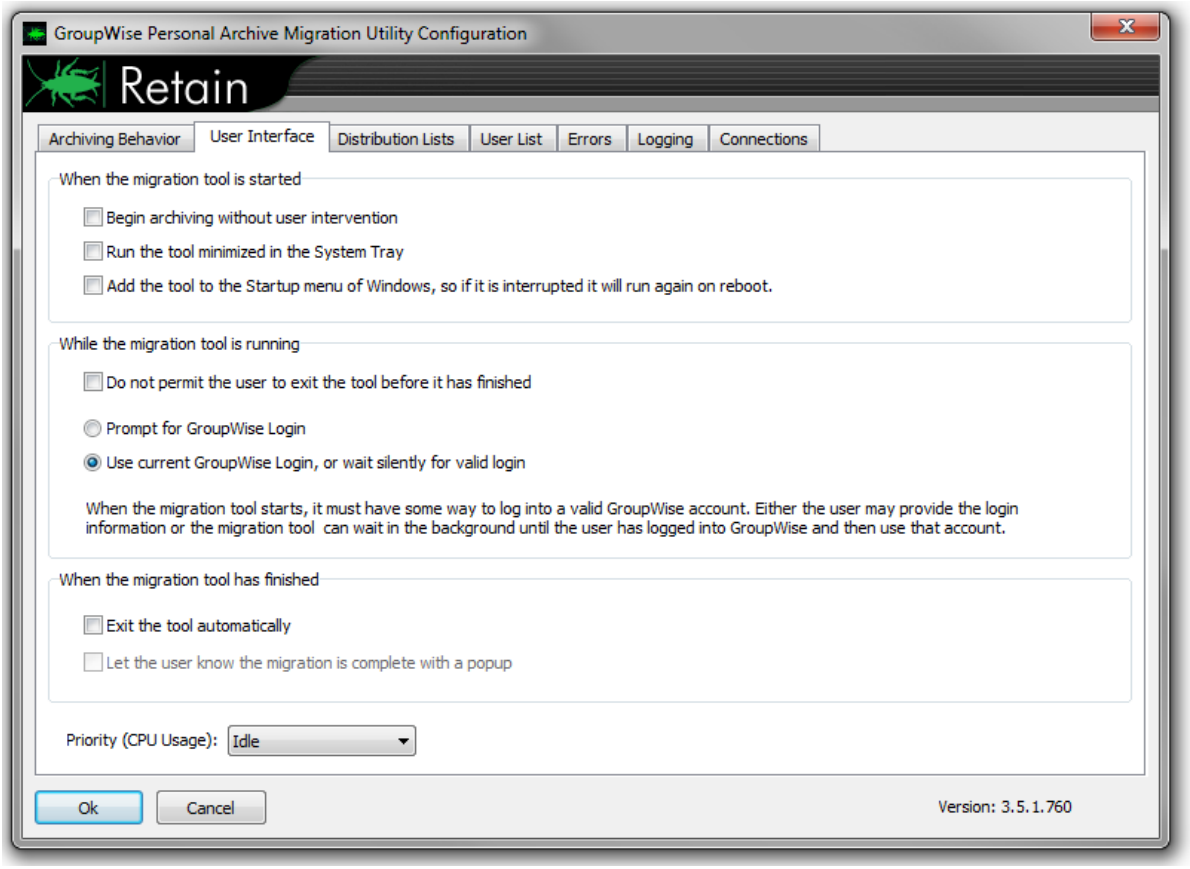
8.19.7 Archiving Behavior

- Date Range – limits or specifies specific mail to process based on a date range.
- Archive Directory – The location of the archive or archives to be migrated. The default path is assumed unless selected to be specified.
- Miscellaneous – contains the options to preserve the routing properties in the archive or to add a marker file when the migration is complete.



8.19.8 User Interface

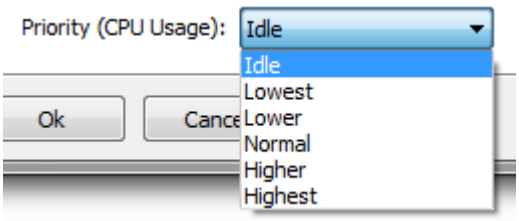
When the Migration Tool is Started - Begin without user intervention the archive process immediately begins without any prompts. Run minimized sets the migration utility to run the migration in the background while the workstation completes other work. The tool may also be set to run automatically on system boot, in case something interrupts the archive job.



When Migration Tool is Running - Contains options to prohibit the user exiting the program before the archive migration is complete and GroupWise login settings. The tool must be able to login and gain access to the archives.

When the Migration Tool has Finished – Contains the behavior of the migration utility on completion; automatic exit, notification, or neither.

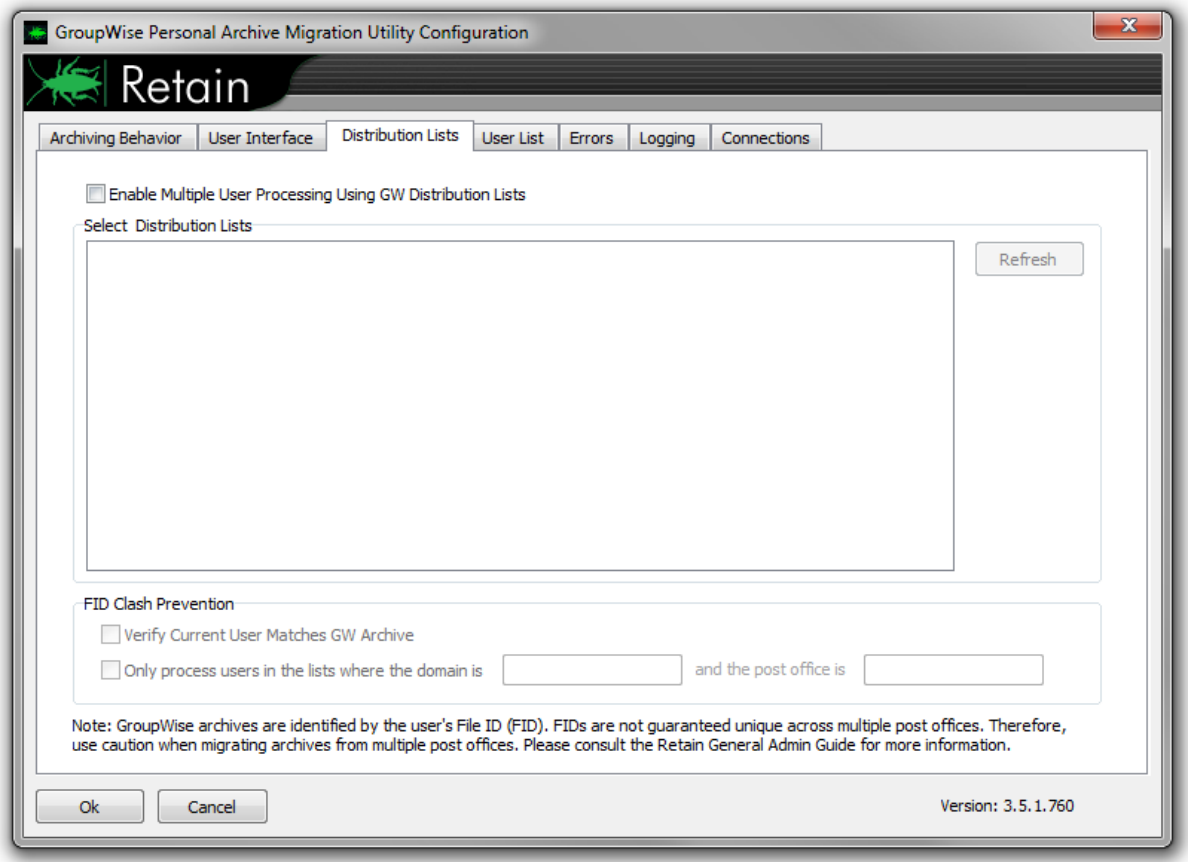
CPU priority usage is also set here.



8.19.9 Distribution Lists

In order to process more than one user archive the option must be enabled at the top of this page and the users must be verified against the GroupWise system.

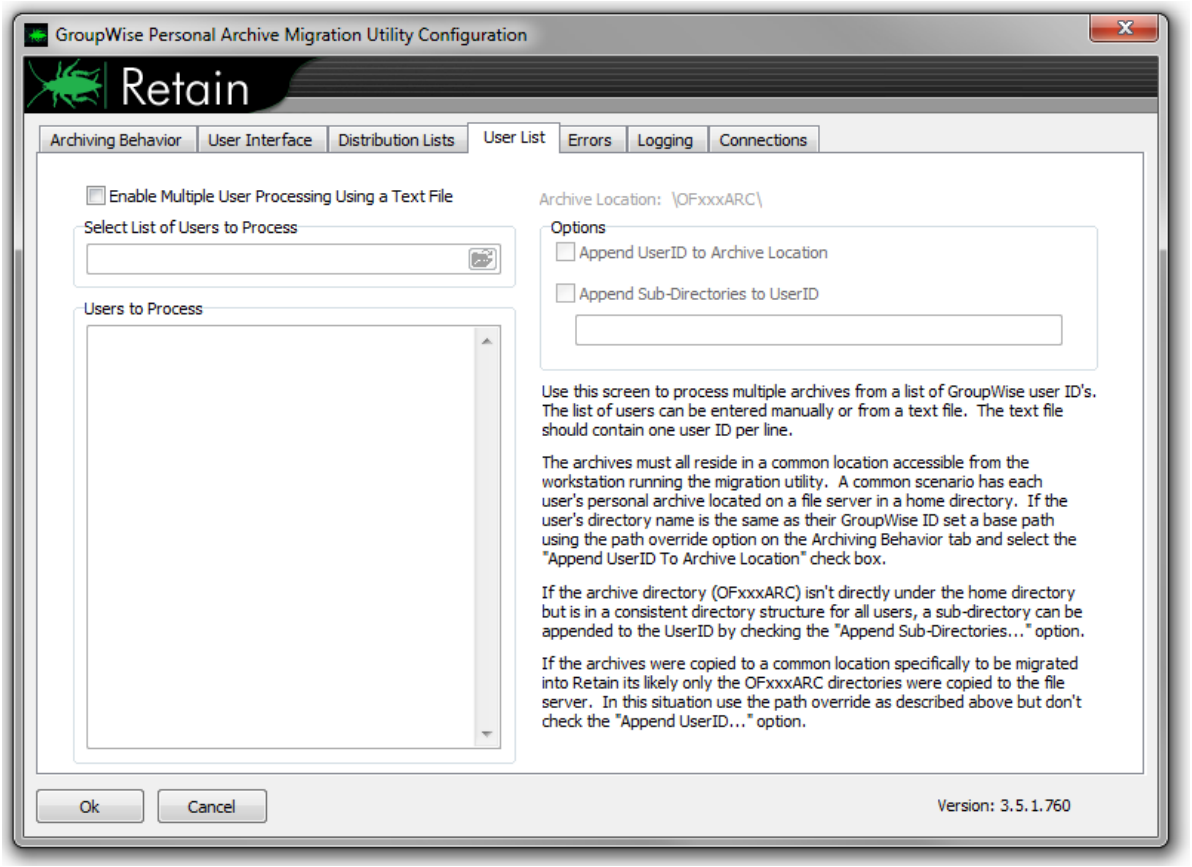
It is recommended to only process multiple archives from the same post office.



The Distribution Lists will only be populated if it is enabled and if the ‘Refresh’ button is selected. It dis-plays the available distribution lists. To have the connection work correctly, the GroupWise Client Param-eters on the ‘Connections’ tab must be filled-out with the GroupWise Server IP and port. Displayed lists can be selected. The archives belonging to the users in the selected distribution list(s) will be migrated. If users from multiple Post Offices are to be migrated, employ some kind of FID clash protection by verify-ing archives or only processing users which can be verified by a specified Domain and Post Office, (speci-fied in DNS:port or IP address:port).

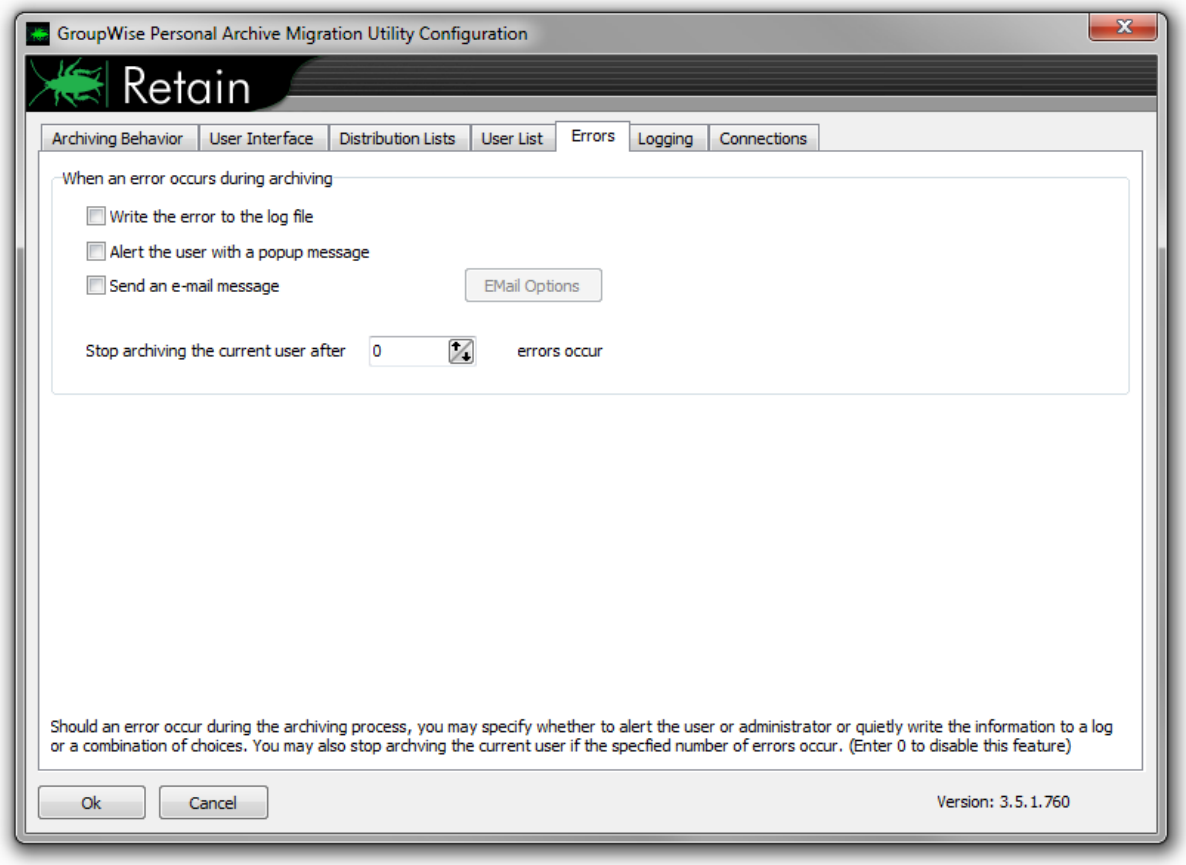
8.19.10 User List

The User List screen allows the migration of multiple users specified by a user list. The user list is a plain text file with one user ID per line. To enable the process, select the ‘Enable Multiple User Processing Using a Text File checkbox, then browse to and select the desired text file.

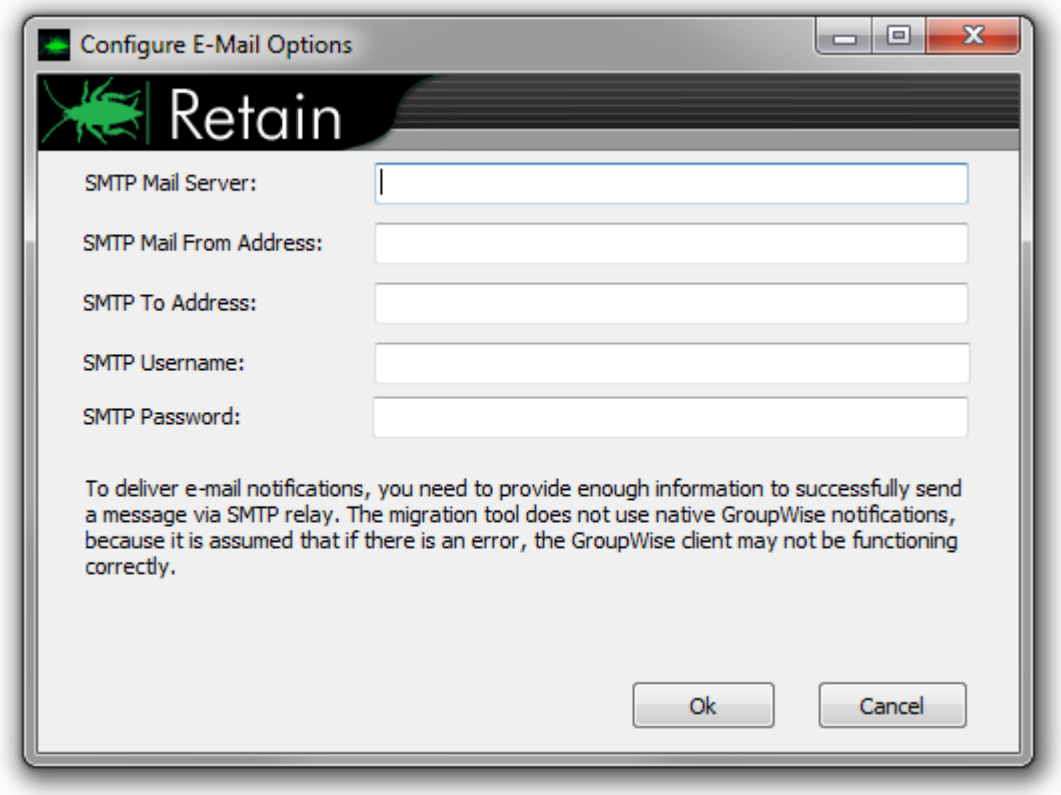


8.19.11 Errors

The error limit and actions taken when an error is encountered are specified on this tab.



Errors may be written to the log file, (specified in the next tab), sent in a pop-up message to the user, or sent as an email to the specified address.



The limit of how many errors may be encountered before the migration utility quits the current user is specified. A setting of ‘0’ sets the limit to infinite.

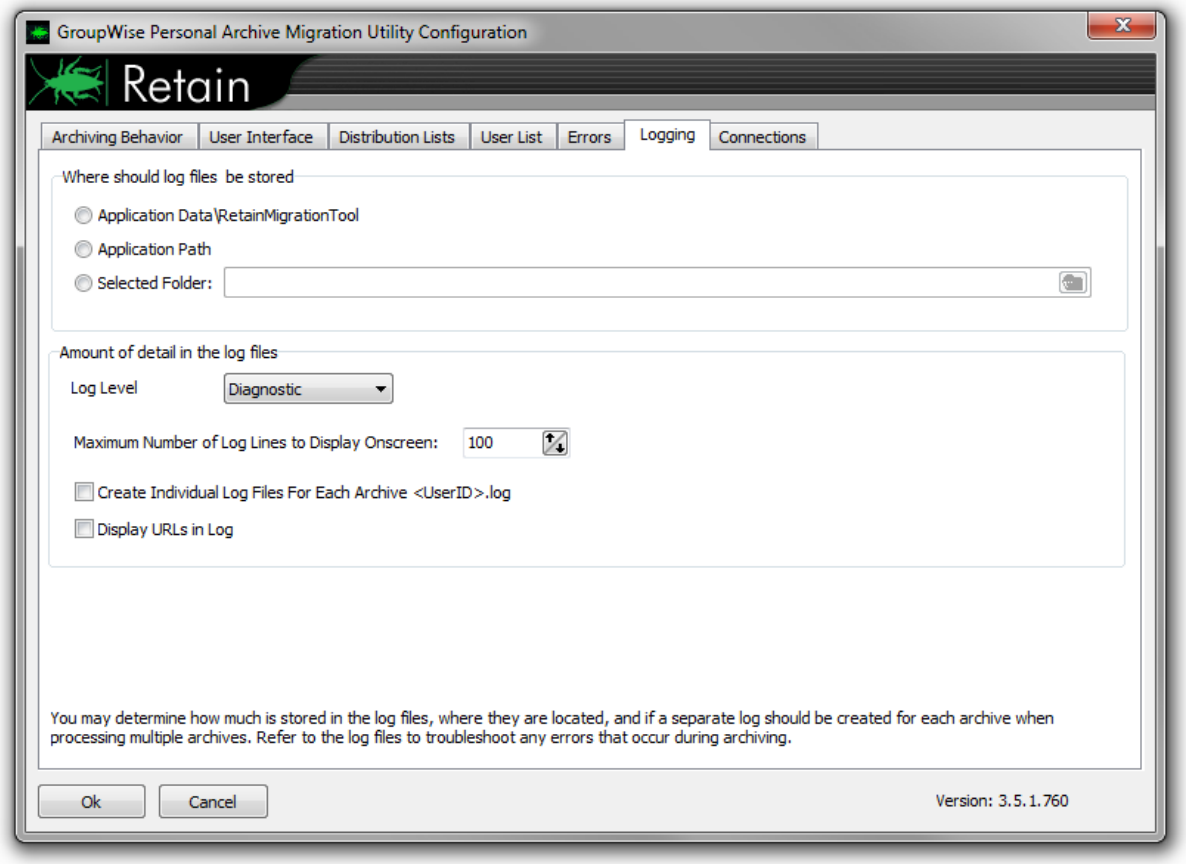
Important: By default, the current user will be skipped after 20 errors by default.

If an email is desired for each error encountered, email settings must be configured. Select ‘Email Options ‘and input the connection information, source address, destination address, and user login information for the SMTP Mail Server.

8.19.12 Logging

The location, logging levels, and buffer of the logging screen are all configured on this tab.

Where Should Log Files be Stored - A specified path, the path to the current location of the application, or a specified path must be selected.

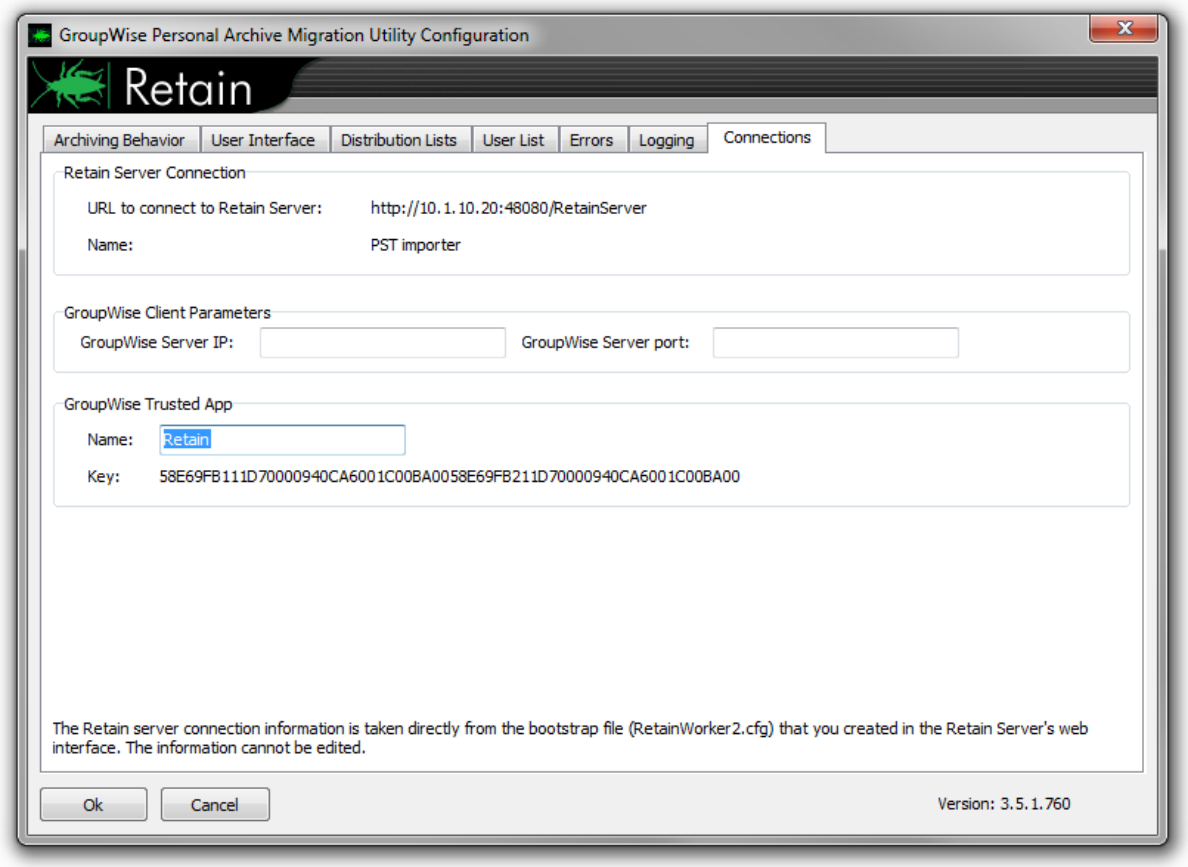


Amount of Detail in the Log Files – The logging level determines how much information is provided in the log files. Unless troubleshooting, log levels of ‘normal’ and ‘errors only’ are sufficient. The buffer size for the logging screen determines how much history the running log screen of the migration tool contains.

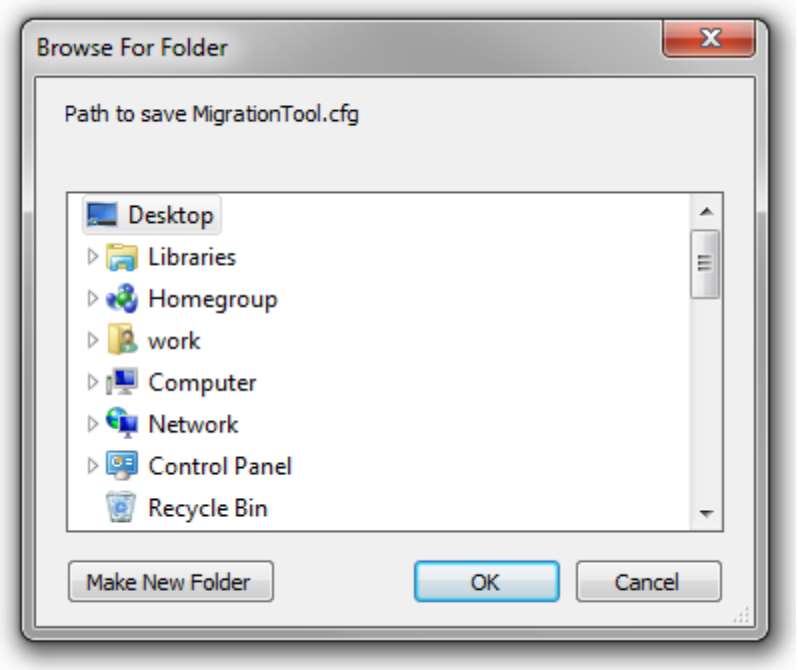
8.19.13 Connections

This is an informational page showing the connection settings from the RetainWorker.cfg created in the Retain Server. If this information is incorrect, the Retain Server Connection settings must be modified under the worker configuration in the Retain Server, and a new bootstrap file downloaded for use.

The GroupWise Client Parameters contains the connection information used by the Distribution Lists con-figuration to allow the Migrator to connect to the GroupWise system.



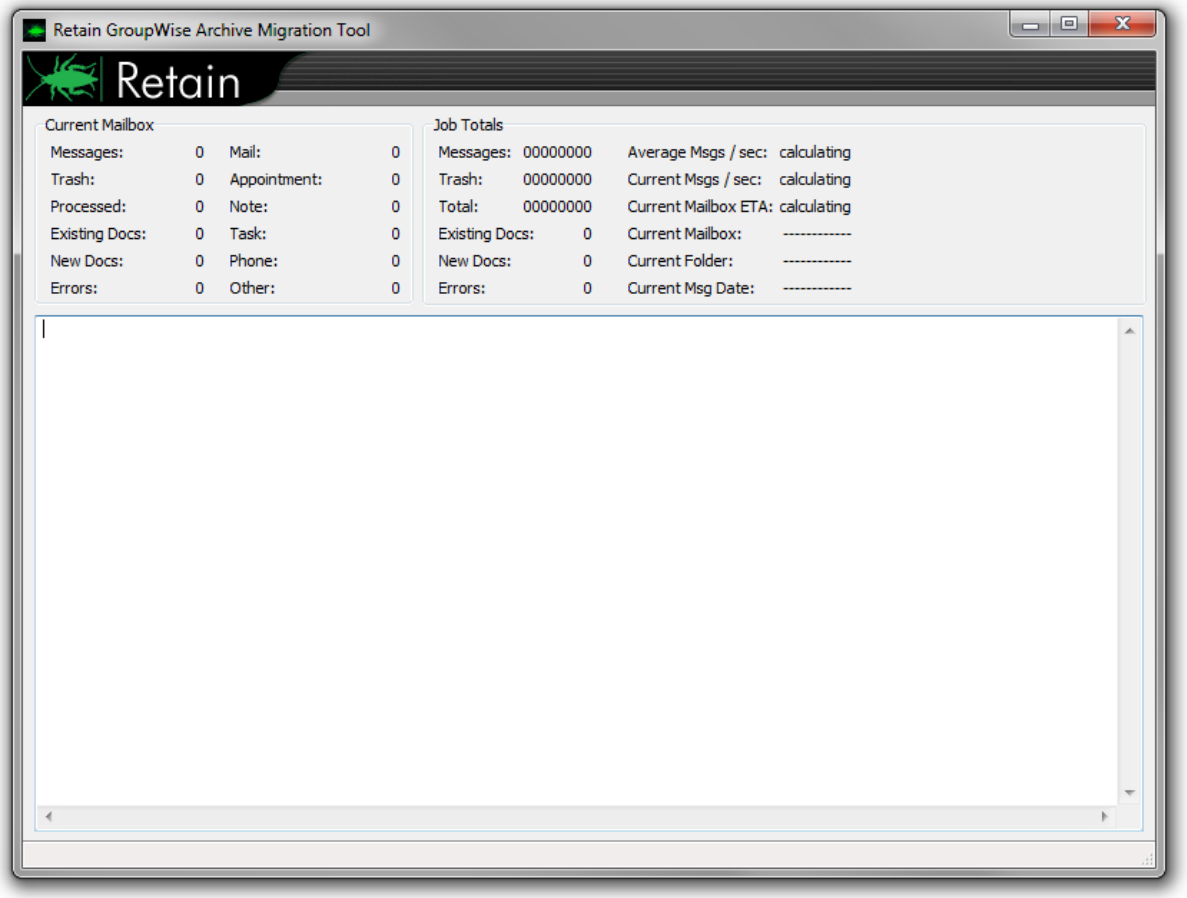
Once the configuration has been setup as desired, click ‘Ok’ and select the location to save the changes into a ‘MigrationTool.cfg’ file.



8.19.14 Running the Migration

The migration tool will use the configuration file to run the job as you have specified in the configuration tool. You may run the configuration tool as often as you need to set up the migration job to your liking.

When run, the tool will prompt for the configuration file, or automatically load the configuration file if the configuration is found it the same directory, and immediately begin the migration. When the tool runs, you will see a screen like this:



You will see any errors in this screen and it will display a notification when the job is complete.

For each user in a multi-user job, it will read the location of the archives from GroupWise and then it will attempt to open the archives. It is critically important that the workstation used for the migration can access the archive files.

8.19.15 Limit of 5 Concurrent Migrations

There is a maximum of 5 migrations permitted at any given time. For instance, if you're running the migration tool from the workstations of your users, only 5 at a time can be migrated, the others will wait their turn. This is to prevent all the migration requests from overwhelming the server.

8.19.16 Other Notes

The migration tool may be deployed using management tools such as Novell ZENworks. This is one method to collect archives which have been stored on users' workstations. Users may run the tool manu-ally or it can be run automatically.

UUID errors

[ERROR] Exception: Error getting UUID.

There are a various reasons for a UUID error:

- The User does not exist in the GroupWise system.

  If the user does not exist in GroupWise anymore but the personal archive needs to be moved into Retain, Retain will create the user mailbox and assign it a Retain UUID (unique user ID). Thus, this is also why the address book needs to be refreshed prior to using the PAM tool. If the user exists in the GroupWise system, you would get two different user mailboxes in Retain if you archived with the PAM tool first and then refreshed the address book in the Retain system.
- The Users's visibility is set too low in GroupWise.

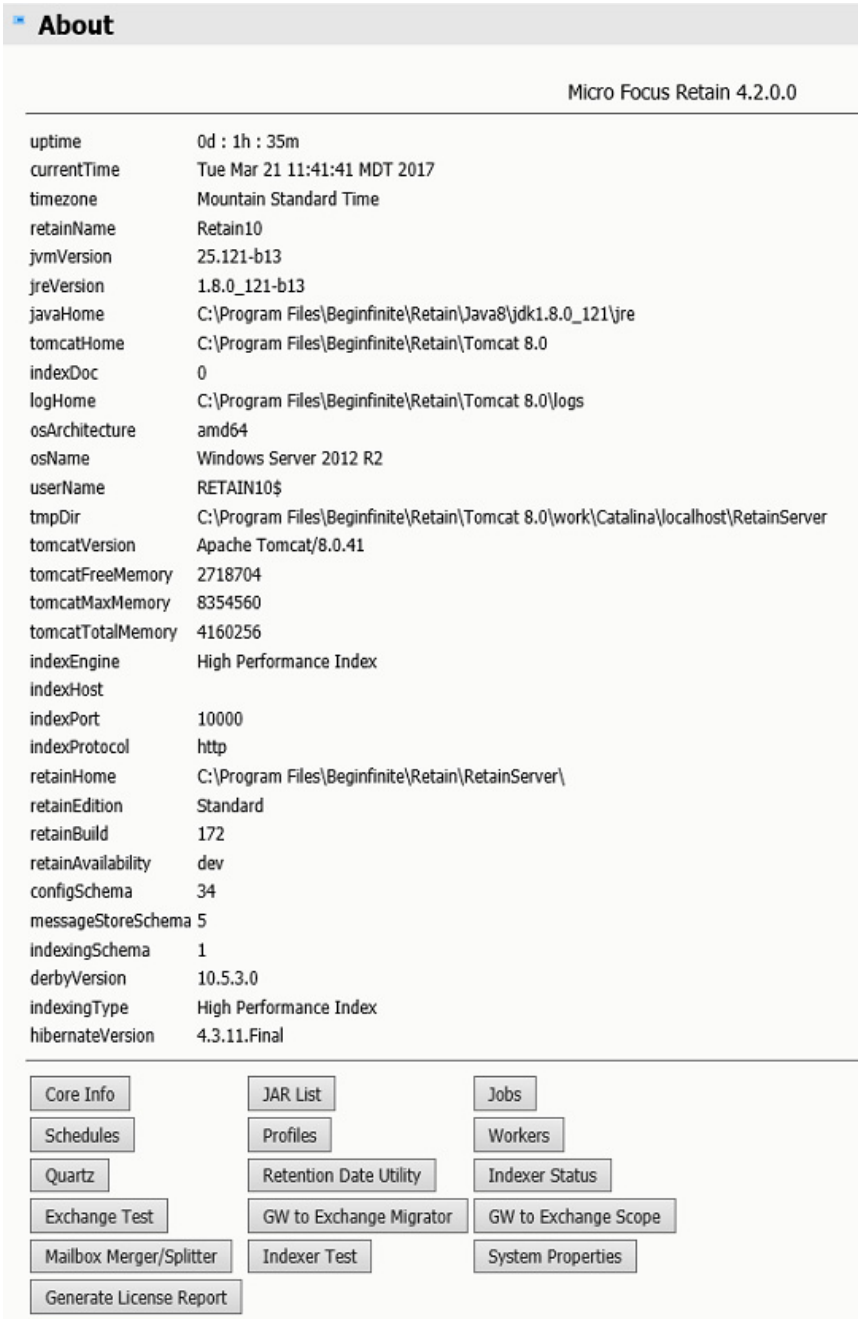
  In GroupWise, edit the account of the user and set the visibility to System.
- Apache has been configured to redirect traffic to a secure HTTPS connection.

  Edit the bootstrap file to connect to http://[Retain Server IP address or hostname]:48080 to connect directly to tomcat and bypass the apache redirect.

9 About Retain

9.1 About and Diagnostics Page

To get this page from anywhere in the administrative console of Retain, click on the Retain logo in the top left corner, or the ‘about’ link in the top right.



This page displays statistics and general program information as well as linking to valuable diagnostics. When tuning the memory usage of the Retain compo-nents, it is useful to refer to these diagnostics and statistics.

Operational statistics are listed such as uptime, component ver-sion numbers, directory paths and memory statistics.

Below that, the buttons will give you more detailed diagnostic information. Generally, you will only use these at the request of technical support.

Please pay particular attention to the Tomcat memory statistics. This information is needed to fine tune Tomcat.

9.1.1 Diagnostic Buttons

In usual operation, none of these buttons will be used without instruction of support. The diag-nostic buttons found on the diag-nostics page are used in conjunction with support and can be dangerous for the system if used without understanding or instruction. The utilities con-tained here display or modify core settings and functions of the Retain Server. Mismanagement of these functions may render the Retain Server unmanageable and in a nonfunctioning state.

9.1.2 Core info, JAR List, Jobs, Schedules, Profiles, Work-ers

The Core info, JAR List, Jobs, Schedules, Profiles, and Workers buttons initiate a download of the configu-ration .xml for each of those categories. All configuration tied to those items which have been configured in the Retain Server will be contained in the downloaded .xml and can be used to examine recorded con-figuration.

9.1.3 Quartz

The Quartz button displays a page containing all information and configuration for the Quartz scheduler. These settings determine when jobs start, when maintenance is, and time in the system.

9.1.4 Retention Date Utility

The Retention Date Utility tests the retention and purge flags in the mail systems. It is set in UNIX time stamp format, and should not be used without express guidance from Support.

9.1.5 Indexer Status

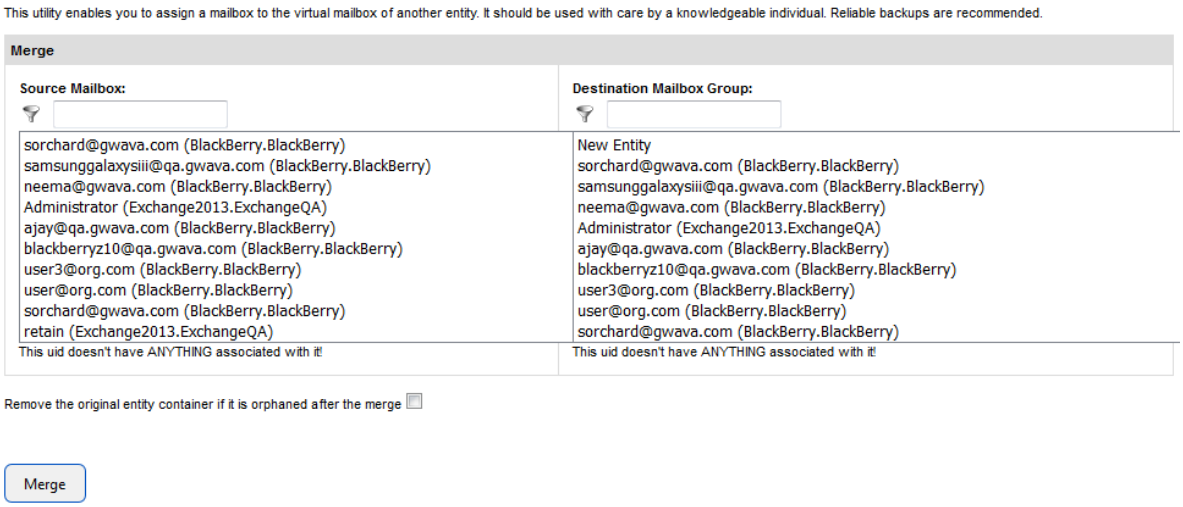
The Indexer Status displays the current health of the indexer engine along with the total unindexed items and associated memory load. There is also an option to shut off or to turn on the indexer. As the indexer should be alive at ‘on’ at all times during normal operation, only use this utility under instruction from Support.

9.1.6 Exchange Test

The *Exchange Test* utility attempts to connect to and login to the specified mailbox in the Exchange mail system, utilizing current configuration from the Exchange module, if configured. This is used to ensure that sufficient rights and access is open to Retain in the Exchange system. A target mailbox different from the Retain ‘service’ account must be specified.

9.1.7 Mailbox Merger / Splitter

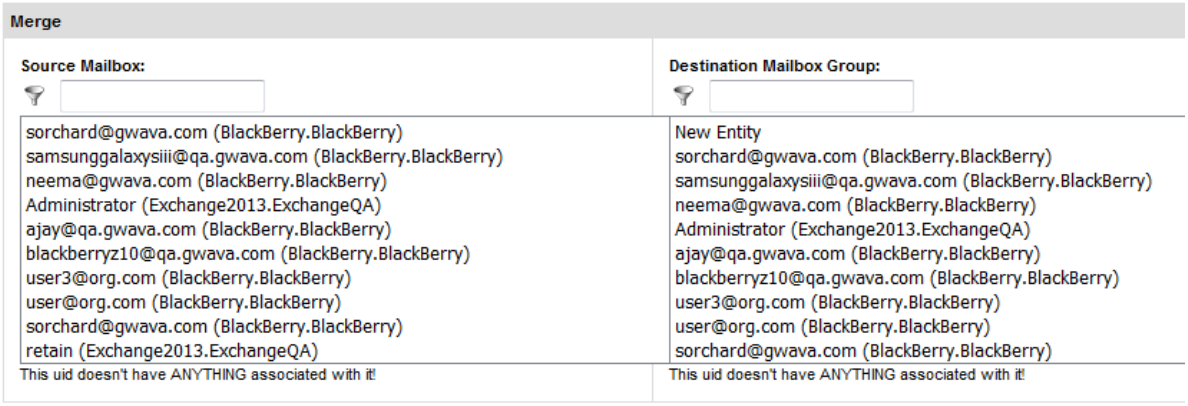
Mailboxes may be automatically merged by the Retain system if configured correctly under the Module Configuration |Mailbox Mapping options screen, but it must be done in correct order or the system will not recognize that a change has occurred and the merge will fail. Automatic merge can only be per-formed once.





Manually merging and splitting mailboxes is repeatable and is also the only way to undo a merge, auto-matic or otherwise. A batch load may be created and processed, but requires specific formatting, see the ‘mergesample.xml’ for format information.

Merging accounts creates a connection between two existing entities in the Retain database by merging their UID’s in Retain, and allows a user with two differing email accounts, or one account repeated on two different systems, to manage their mail from within the same interface without having to search for and select a different account, or separately logging-in to two different systems.



To merge two addresses, first select the source mailbox and then select the destination mailbox. After the desired mailboxes have been selected, click the green ‘merge’ button.

To remove currently merged accounts, select the account you wish to sever as the source mailbox and merge it to a ‘new entity’.

Because the merge utility relies on tying entities to accounts in the database, the ‘new entity’ simply unties the account from any other connections. Tying a user to a new entity does not create any new, floating entities or make any changes in the mail systems; it only severs the selected account from previ-ously merged accounts.

9.1.8 Indexer Test

The indexer test utility provides a simply and quick query to the index engine and database. This is used as a low-level diagnostic tool for Support and is best used in conjunction with Support. Normal operation does not require any use of this tool.

9.1.9 System Properties

The System Properties utility displays basic system information and settings about the Retain Server, host environment, and functioning parts.

9.2 Retain Component Licensing

Copyright 2017 © Micro Focus International plc. All rights reserved.

Retain uses many components covered under the following licenses. The licenses follow the list-ing below.

License	Components
Apache 2.0 License	Apache, Open Symphony Quartz, Apache Log4j & Log4j Additions, Apache Derby, Apache XML-RPC, Apache TagLib, Apache Commons, CGLib, Apache Tika, Jet 53 Toolkit, SPring Framework
LGPL	EDS, Hibernate, C3PO Connection Pooling, Jericho HTML Parser, Textmining, 7-zip

License	Components
BSD License	XStream, JDOM, MetaStuff DOM4J, Object Web ASM, ANTLR
BOuncy Castle APIS	Special license included below
PDFBox	Special license included below
Open Software License v.3.0	openXML Indexer
Mozilla Public License	iText
GNU GPL V. 2	VAInstall
Bitstream Vera Fonts Copyright	DejaVuSans Fonts
ICU Licence	Java Utilities

9.2.1 Apache 2.0 License

Apache License - Version 2.0, January 2004

http://www.apache.org/licenses/

TERMS AND CONDITIONS FOR USE, REPRODUCTION, AND DISTRIBUTION

1.Definitions.

"License" shall mean the terms and conditions for use, reproduction, and distribution as defined by Sec-tions 1 through 9 of this document.

"Licensor" shall mean the copyright owner or entity authorized by the copyright owner that is granting the License.

"Legal Entity" shall mean the union of the acting entity and all other entities that control, are controlled by, or are under common control with that entity. For the purposes of this definition, "control" means (i) the power, direct or indirect, to cause the direction or management of such entity, whether by contract or otherwise, or (ii) ownership of fifty percent (50%) or more of the outstanding shares, or (iii) beneficial ownership of such entity.

"You" (or "Your") shall mean an individual or Legal Entity exercising permissions granted by this License.

"Source" form shall mean the preferred form for making modifications, including but not limited to soft-ware source code, documentation source, and configuration files.

"Object" form shall mean any form resulting from mechanical transformation or translation of a Source form, including but not limited to compiled object code, generated documentation, and conversions to other media types.

"Work" shall mean the work of authorship, whether in Source or Object form, made available under the License, as indicated by a copyright notice that is included in or attached to the work (an example is pro-vided in the Appendix below).

"Derivative Works" shall mean any work, whether in Source or Object form, that is based on (or derived from) the Work and for which the editorial revisions, annotations, elaborations, or other modifications represent, as a whole, an original work of authorship. For the purposes of this License, Derivative Works shall not include works that remain separable from, or merely link (or bind by name) to the interfaces of, the Work and Derivative Works thereof.

“Contribution” shall mean any work of authorship, including the original version of the Work and any modifications or additions to that Work or Derivative Works thereof, that is intentionally submitted to Licensor for inclusion in the Work by the copyright owner or by an individual or Legal Entity authorized to submit on behalf of the copyright owner. For the purposes of this definition, “submitted” means any form of electronic, verbal, or written communication sent to the Licensor or its representatives, including but not limited to communication on electronic mailing lists, source code control systems, and issue tracking systems that are managed by, or on behalf of, the Licensor for the purpose of discussing and improving the Work, but excluding communication that is conspicuously marked or otherwise designated in writing by the copyright owner as "Not a Contribution."

"Contributor" shall mean Licensor and any individual or Legal Entity on behalf of whom a Contribution has been received by Licensor and subsequently incorporated within the Work.

2.Grant of Copyright License. Subject to the terms and conditions of this License, each Contribu-tor hereby grants to You a perpetual, worldwide, non-exclusive, no-charge, royalty-free, irrevoca-ble copyright license to reproduce, prepare Derivative Works of, publicly display, publicly perform, sublicense, and distribute the Work and such Derivative Works in Source or Object form.

3.Grant of Patent License. Subject to the terms and conditions of this License, each Contributor hereby grants to You a perpetual, worldwide, non-exclusive, no-charge, royalty-free, irrevocable (except as stated in this section) patent license to make, have made, use, offer to sell, sell, import, and otherwise transfer the Work, where such license applies only to those patent claims licensable by such Contributor that are necessarily infringed by their Contribution(s) alone or by combination of their Contribution(s) with the Work to which such Contribution(s) was submitted. If You institute patent litigation against any entity (including a cross-claim or counterclaim in a lawsuit) alleging that the Work or a Contribution incorporated within the Work constitutes direct or contributory patent infringement, then any patent licenses granted to You under this License for that Work shall terminate as of the date such litigation is filed.

4.Redistribution. You may reproduce and distribute copies of the Work or Derivative Works thereof in any medium, with or without modifications, and in Source or Object form, provided that You meet the following conditions:

a)You must give any other recipients of the Work or Derivative Works a copy of this License; and

b)You must cause any modified files to carry prominent notices stating that You changed the files; and

c)You must retain, in the Source form of any Derivative Works that You distribute, all copy-right, patent, trademark, and attribution notices from the Source form of the Work, exclud-ing those notices that do not pertain to any part of the Derivative Works; and

d)If the Work includes a "NOTICE" text file as part of its distribution, then any Derivative Works that You distribute must include a readable copy of the attribution notices contained within such NOTICE file, excluding those notices that do not pertain to any part of the Deriva-tive Works, in at least one of the following places: within a NOTICE text file distributed as part of the Derivative Works; within the Source form or documentation, if provided along with the Derivative Works; or, within a display generated by the Derivative Works, if and wherever such third-party notices normally appear. The contents of the NOTICE file are for informa-tional purposes only and do not modify the License. You may add Your own attribution notices within Derivative Works that You distribute, alongside or as an addendum to the NOTICE text from the Work, provided that such additional attribution notices cannot be con-strued as modifying the License.

You may add Your own copyright statement to Your modifications and may provide additional or differ-ent license terms and conditions for use, reproduction, or distribution of Your modifications, or for any such Derivative Works as a whole, provided Your use, reproduction, and distribution of the Work other-wise complies with the conditions stated in this License.

5.Submission of Contributions. Unless You explicitly state otherwise, any Contribution intention-ally submitted for inclusion in the Work by You to the Licensor shall be under the terms and con-ditions of this License, without any additional terms or conditions. Notwithstanding the above, nothing herein shall supersede or modify the terms of any separate license agreement you may have executed with Licensor regarding such Contributions.

6.Trademarks. This License does not grant permission to use the trade names, trademarks, ser-vice marks, or product names of the Licensor, except as required for reasonable and customary use in describing the origin of the Work and reproducing the content of the NOTICE file.

7.Disclaimer of Warranty. Unless required by applicable law or agreed to in writing, Licensor provides the Work (and each Contributor provides its Contributions) on an "AS IS" BASIS, WITH-OUT WARRANTIES OR CONDITIONS OF ANY KIND, either express or implied, including, without limitation, any warranties or conditions of TITLE, NON-INFRINGEMENT, MERCHANTABILITY, or FITNESS FOR A PARTICULAR PURPOSE. You are solely responsible for determining the appropri-ateness of using or redistributing the Work and assume any risks associated with Your exercise of permissions under this License.

8.Limitation of Liability. In no event and under no legal theory, whether in tort (including negli-gence), contract, or otherwise, unless required by applicable law (such as deliberate and grossly negligent acts) or agreed to in writing, shall any Contributor be liable to You for damages, includ-ing any direct, indirect, special, incidental, or consequential damages of any character arising as a result of this License or out of the use or inability to use the Work (including but not limited to damages for loss of goodwill, work stoppage, computer failure or malfunction, or any and all other commercial damages or losses), even if such Contributor has been advised of the possibility of such damages.

9.Accepting Warranty or Additional Liability. While redistributing the Work or Derivative Works thereof, You may choose to offer, and charge a fee for, acceptance of support, warranty, indem-nity, or other liability obligations and/or rights consistent with this License. However, in accept-ing such obligations, You may act only on Your own behalf and on Your sole responsibility, not on behalf of any other Contributor, and only if You agree to indemnify, defend, and hold each Con-tributor harmless for any liability incurred by, or claims asserted against, such Contributor by rea-son of your accepting any such warranty or additional liability.

END OF TERMS AND CONDITIONS

APPENDIX: How to apply the Apache License to your work.

To apply the Apache License to your work, attach the following boilerplate notice, with the fields enclosed by brackets "[]" replaced with your own identifying information. (Don't include the brackets!) The text should be enclosed in the appropriate comment syntax for the file format. We also recommend that a file or class name and description of purpose be included on the same "printed page" as the copy-right notice for easier identification within third-party archives.

Copyright [yyyy] [name of copyright owner]

Licensed under the Apache License, Version 2.0 (the "License"); you may not use this file except in com-pliance with the License.

You may obtain a copy of the License at

http://www.apache.org/licenses/LICENSE-2.0

Unless required by applicable law or agreed to in writing, software distributed under the License is dis-tributed on an "AS IS" BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY KIND, either express or implied. See the License for the specific language governing permissions and limitations under the License.

9.3 LGPL – GNU Lesser General Public License

GNU LESSER GENERAL PUBLIC LICENSE

Version 3, 29 June 2007

Copyright (C) 2007 Free Software Foundation, Inc. http://fsf.org/

Everyone is permitted to copy and distribute verbatim copies of this license document, but changing it is not allowed.

This version of the GNU Lesser General Public License incorporates the terms and conditions of version 3 of the GNU General Public License, supplemented by the additional permissions listed below.

10.Additional Definitions.

As used herein, “this License” refers to version 3 of the GNU Lesser General Public License, and the “GNU GPL” refers to version 3 of the GNU General Public License.

“The Library” refers to a covered work governed by this License, other than an Application or a Combined Work as defined below.

An “Application” is any work that makes use of an interface provided by the Library, but which is not oth-erwise based on the Library. Defining a subclass of a class defined by the Library is deemed a mode of using an interface provided by the Library.

A “Combined Work” is a work produced by combining or linking an Application with the Library. The par-ticular version of the Library with which the Combined Work was made is also called the “Linked Ver-sion”.

The “Minimal Corresponding Source” for a Combined Work means the Corresponding Source for the Combined Work, excluding any source code for portions of the Combined Work that, considered in isola-tion, are based on the Application, and not on the Linked Version.

The “Corresponding Application Code” for a Combined Work means the object code and/or source code for the Application, including any data and utility programs needed for reproducing the Combined Work from the Application, but excluding the System Libraries of the Combined Work.

11.Exception to Section 3 of the GNU GPL.

You may convey a covered work under sections 3 and 4 of this License without being bound by section 3 of the GNU GPL.

12.Conveying Modified Versions.

If you modify a copy of the Library, and, in your modifications, a facility refers to a function or data to be supplied by an Application that uses the facility (other than as an argument passed when the facility is invoked), then you may convey a copy of the modified version:

a)under this License, provided that you make a good faith effort to ensure that, in the event an application does not supply the function or data, the facility still operates, and performs whatever part of its purpose remains meaningful, or

b)under the GNU GPL, with none of the additional permissions of this License applicable to that copy.

13.Object Code Incorporating Material from Library Header Files.

The object code form of an Application may incorporate material from a header file that is part of the Library. You may convey such object code under terms of your choice, provided that, if the incorporated material is not limited to numerical parameters, data structure layouts and accessors, or small macros, inline functions and templates (ten or fewer lines in length), you do both of the following:

a)Give prominent notice with each copy of the object code that the Library is used in it and that the Library and its use are covered by this License.

b)Accompany the object code with a copy of the GNU GPL and this license document.

14.Combined Works.

You may convey a Combined Work under terms of your choice that, taken together, effectively do not restrict modification of the portions of the Library contained in the Combined Work and reverse engineering for debugging such modifications, if you also do each of the following:

a)Give prominent notice with each copy of the Combined Work that the Library is used in it and that the Library and its use are covered by this License.

b)Accompany the Combined Work with a copy of the GNU GPL and this license document.

c)For a Combined Work that displays copyright notices during execution, include the copyright notice for the Library among these notices, as well as a reference directing the user to the copies of the GNU GPL and this license document.

d)Do one of the following:

0)Convey the Minimal Corresponding Source under the terms of this License, and the Corresponding Application Code in a form suitable for, and under terms that permit, the user to recombine or relink the Application with a modified version of the Linked Version to produce a modified Combined Work, in the manner specified by section 6 of the GNU GPL for conveying Corresponding Source.

1)Use a suitable shared library mechanism for linking with the Library. A suitable mecha-nism is one that (a) uses at run time a copy of the Library already present on the user's computer system, and (b) will operate properly with a modified version of the Library that is interface-compatible with the Linked Version.

e)Provide Installation Information, but only if you would otherwise be required to provide such information under section 6 of the GNU GPL, and only to the extent that such informa-tion is necessary to install and execute a modified version of the Combined Work produced by recombining or relinking the Application with a modified version of the Linked Version. (If you use option 14d0, the Installation Information must accompany the Minimal Correspond-ing Source and Corresponding Application Code. If you use option 14d1, you must provide the Installation Information in the manner specified by section 6 of the GNU GPL for conveying Corresponding Source.)

15.Combined Libraries.

You may place library facilities that are a work based on the Library side by side in a single library together with other library facilities that are not Applications and are not covered by this License, and convey such a combined library under terms of your choice, if you do both of the following:

a)Accompany the combined library with a copy of the same work based on the Library, uncombined with any other library facilities, conveyed under the terms of this License.

b)Give prominent notice with the combined library that part of it is a work based on the Library, and explaining where to find the accompanying uncombined form of the same work.

16.Revised Versions of the GNU Lesser General Public License.

The Free Software Foundation may publish revised and/or new versions of the GNU Lesser General Public License from time to time. Such new versions will be similar in spirit to the present version, but may differ in detail to address new problems or concerns.

Each version is given a distinguishing version number. If the Library as you received it specifies that a cer-tain numbered version of the GNU Lesser General Public License “or any later version” applies to it, you have the option of following the terms and conditions either of that published version or of any later ver-sion published by the Free Software Foundation. If the Library as you received it does not specify a ver-sion number of the GNU Lesser General Public License, you may choose any version of the GNU Lesser General Public License ever published by the Free Software Foundation.

If the Library as you received it specifies that a proxy can decide whether future versions of the GNU Lesser General Public License shall apply, that proxy's public statement of acceptance of any version is permanent authorization for you to choose that version for the Library.

9.4 BSD License – Individual Component Licenses

9.4.1 ANTLR 3 License

[The BSD License]

Copyright (c) 2003-2007, Terence Parr

All rights reserved.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

1.Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.

2.Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the dis-tribution.

3.Neither the name of the author nor the names of its contributors may be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDERS AND CONTRIBUTORS "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE COPYRIGHT OWNER OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOW-EVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

9.4.2 ASM

Copyright (c) Object Web, ASM project.

All rights reserved.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

1.Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.

2.Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the dis-tribution.

3.Neither the name of the ASM Project nor the names of its contributors may be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDERS AND CONTRIBUTORS "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE COPYRIGHT OWNER OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOW-EVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

9.4.3 DOM4J

Redistribution and use of this software and associated documentation ("Software"), with or without modification, are permitted provided that the following conditions are met:

1.Redistributions of source code must retain copyright statements and notices. Redistributions must also contain a copy of this document.

2.Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the dis-tribution.

3.The name "DOM4J" must not be used to endorse or promote products derived from this Soft-ware without prior written permission of MetaStuff, Ltd. For written permission, please contact dom4j-info@metastuff.com.

4.Products derived from this Software may not be called "DOM4J" nor may "DOM4J" appear in their names without prior written permission of MetaStuff, Ltd. DOM4J is a registered trademark of MetaStuff, Ltd.

Due credit should be given to the DOM4J Project - http://www.dom4j.org

THIS SOFTWARE IS PROVIDED BY METASTUFF, LTD. AND CONTRIBUTORS ``AS IS'' AND ANY EXPRESSED OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABI-LITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL METASTUFF, LTD. OR ITS CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER

CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLI-GENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

Copyright 2001-2005 (C) MetaStuff, Ltd. All Rights Reserved.

9.4.4 JDOM

Copyright (C) 2000-2004 Jason Hunter & Brett McLaughlin.

All rights reserved.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

1.Redistributions of source code must retain the above copyright notice, this list of conditions, and the following disclaimer.

2.Redistributions in binary form must reproduce the above copyright notice, this list of conditions, and the disclaimer that follows these conditions in the documentation and/or other materials provided with the distribution.

3.The name "JDOM" must not be used to endorse or promote products derived from this software without prior written permission. For written permission, please contact request@jdom.org.

Products derived from this software may not be called "JDOM", nor may "JDOM" appear in their name, without prior written permission from the JDOM Project Management request@jdom.org.

In addition, we request (but do not require) that you include in the end-user documentation provided with the redistribution and/or in the software itself an acknowledgement equivalent to the following:

"This product includes software developed by the JDOM Project (http://www.jdom.org/)."

Alternatively, the acknowledgment may be graphical using the logos available at http://www.jdom.org/images/logos.

THIS SOFTWARE IS PROVIDED ``AS IS'' AND ANY EXPRESSED OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE JDOM AUTHORS OR THE PROJECT CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAM-AGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIA-BILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

This software consists of voluntary contributions made by many individuals on behalf of the JDOM Proj-ect and was originally created by Jason Hunter <jhunter@jdom.org> and Brett McLaughlin <brett@jdom.org>. For more information on the JDOM Project, please see <http://www.jdom.org/>.

9.4.5 XStream

Copyright (c) 2003-2006, Joe Walnes

Copyright (c) 2006-2007, XStream Committers

All rights reserved.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

Redistributions of source code must retain the above copyright notice, this list of conditions and the fol-lowing disclaimer. Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.

Neither the name of XStream nor the names of its contributors may be used to endorse or promote prod-ucts derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDERS AND CONTRIBUTORS "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE COPYRIGHT OWNER OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOW-EVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

9.4.6 PAL Artistic License

The Artistic License

Preamble

The intent of this document is to state the conditions under which a Package may be copied, such that the Copyright Holder maintains some semblance of artistic control over the development of the package, while giving the users of the package the right to use and distribute the Package in a more-or-less cus-tomary fashion, plus the right to make reasonable modifications.

"Package" refers to the collection of files distributed by the Copyright Holder, and derivatives of that col-lection of files created through textual modification.

"Standard Version" refers to such a Package if it has not been modified, or has been modified in accor-dance with the wishes of the Copyright Holder.

"Copyright Holder" is whoever is named in the copyright or copyrights for the package.

"You" is you, if you're thinking about copying or distributing this Package.

"Reasonable copying fee" is whatever you can justify on the basis of media cost, duplication charges, time of people involved, and so on. (You will not be required to justify it to the Copyright Holder, but only to the computing community at large as a market that must bear the fee.)

"Freely Available" means that no fee is charged for the item itself, though there may be fees involved in handling the item. It also means that recipients of the item may redistribute it under the same conditions they received it.

1.You may make and give away verbatim copies of the source form of the Standard Version of this Package without restriction, provided that you duplicate all of the original copyright notices and associated disclaimers.

2.You may apply bug fixes, portability fixes and other modifications derived from the Public Domain or from the Copyright Holder. A Package modified in such a way shall still be considered the Standard Version.

3.You may otherwise modify your copy of this Package in any way, provided that you insert a prominent notice in each changed file stating how and when you changed that file, and provided that you do at least ONE of the following:

a)place your modifications in the Public Domain or otherwise make them Freely Avail-able, such as by posting said modifications to Usenet or an equivalent medium, or placing the modifications on a major archive site such as ftp.uu.net, or by allowing the Copyright Holder to include your modifications in the Standard Version of the Package.

b)use the modified Package only within your corporation or organization.

c)rename any non-standard executables so the names do not conflict with standard executables, which must also be provided, and provide a separate manual page for each non-standard executable that clearly documents how it differs from the Stan-dard Version.

d)make other distribution arrangements with the Copyright Holder.

4.You may distribute the programs of this Package in object code or executable form, provided that you do at least ONE of the following:

a)distribute a Standard Version of the executables and library files, together with instructions (in the manual page or equivalent) on where to get the Standard Version.

b)accompany the distribution with the machine-readable source of the Package with your modifications.

c)accompany any non-standard executables with their corresponding Standard Ver-sion executables, giving the non-standard executables non-standard names, and clearly documenting the differences in manual pages (or equivalent), together with instructions on where to get the Standard Version.

d)make other distribution arrangements with the Copyright Holder.

5.You may charge a reasonable copying fee for any distribution of this Package. You may charge any fee you choose for support of this Package. You may not charge a fee for this Package itself. However, you may distribute this Package in aggregate with other (possi-bly commercial) programs as part of a larger (possibly commercial) software distribution provided that you do not advertise this Package as a product of your own.

6.The scripts and library files supplied as input to or produced as output from the pro-grams of this package do not automatically fall under the copyright of this Package, but belong to whomever generated them, and may be sold commercially, and may be aggre-gated with this Package.

7.C or perl subroutines supplied by you and linked into this Package shall not be consid-ered part of this Package.

8.The name of the Copyright Holder may not be used to endorse or promote products derived from this software without specific prior written permission.

9.THIS PACKAGE IS PROVIDED "AS IS" AND WITHOUT ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, WITHOUT LIMITATION, THE IMPLIED WARRANTIES OF MERCHANTIBILITY AND FITNESS FOR A PARTICULAR PURPOSE.

The End

9.4.7 Bouncy Castle APIs

Bouncy Castle

Copyright (c) 2000 - 2006 The Legion Of The Bouncy Castle (http://www.bouncycastle.org)

Permission is hereby granted, free of charge, to any person obtaining a copy of this software and associ-ated documentation files (the "Software"), to deal in the Software without restriction, including without limitation the rights to use, copy, modify, merge, publish, distribute, sublicense, and/or sell copies of the Software, and to permit persons to whom the Software is furnished to do so, subject to the following conditions:

The above copyright notice and this permission notice shall be included in all copies or substantial por-tions of the Software.

THE SOFTWARE IS PROVIDED "AS IS", WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, INCLUD-ING BUT NOT LIMITED TO THE WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PUR-POSE AND NONINFRINGEMENT. IN NO EVENT SHALL THE AUTHORS OR COPYRIGHT HOLDERS BE LIABLE FOR ANY CLAIM, DAMAGES OR OTHER LIABILITY, WHETHER IN AN ACTION OF CONTRACT, TORT OR OTH-ERWISE, ARISING FROM, OUT OF OR IN CONNECTION WITH THE SOFTWARE OR THE USE OR OTHER DEAL-INGS IN THE SOFTWARE.

9.4.8 PDF Box

PDFBOX

Copyright (c) 2003-2005, www.pdfbox.org

All rights reserved.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

1.Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.

2.Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.

3.Neither the name of pdfbox; nor the names of its contributors may be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDERS AND CONTRIBUTORS "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE REGENTS OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEM-PLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTI-TUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER

CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

XFire

Copyright (c) 2005 Envoi Solutions LLC

Permission is hereby granted, free of charge, to any person obtaining a copy of this software and associ-ated documentation files (the "Software"), to deal in the Software without restriction, including without limitation the rights to use, copy, modify, merge, publish, distribute, sublicense, and/or sell copies of the Software, and to permit persons to whom the Software is furnished to do so, subject to the following conditions:

The above copyright notice and this permission notice shall be included in all copies or substantial por-tions of the Software.

THE SOFTWARE IS PROVIDED "AS IS", WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, INCLUD-ING BUT NOT LIMITED TO THE WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PUR-POSE AND NONINFRINGEMENT. IN NO EVENT SHALL THE AUTHORS OR COPYRIGHT HOLDERS BE LIABLE FOR ANY CLAIM, DAMAGES OR OTHER LIABILITY, WHETHER IN AN ACTION OF CONTRACT, TORT OR OTH-ERWISE, ARISING FROM, OUT OF OR IN CONNECTION WITH THE SOFTWARE OR THE USE OR OTHER DEAL-INGS IN THE SOFTWARE.

Jaxen

$Id: LICENSE.txt,v 1.5 2006/02/05 21:49:04 elharo Exp $

Copyright 2003-2006 The Werken Company. All Rights Reserved.Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

* Redistributions of source code must retain the above copyright notice, this list of conditions and the fol-lowing disclaimer.

* Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.

* Neither the name of the Jaxen Project nor the names of its contributors may be used to endorse or pro-mote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDERS AND CONTRIBUTORS "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE COPYRIGHT OWNER OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOW-EVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

9.4.9 WSDL4J

Common Public License v1.0

03 Mar 2005

THE ACCOMPANYING PROGRAM IS PROVIDED UNDER THE TERMS OF THIS COMMON PUBLIC LICENSE ("AGREEMENT"). ANY USE, REPRODUCTION OR DISTRIBUTION OF THE PROGRAM CONSTITUTES RECIPI-ENT'S ACCEPTANCE OF THIS AGREEMENT.

DEFINITIONS

"Contribution" means:

in the case of the initial Contributor, the initial code and documentation distributed under this Agree-ment, and

in the case of each subsequent Contributor:

changes to the Program, and

additions to the Program;

where such changes and/or additions to the Program originate from and are distributed by that particu-lar Contributor. A Contribution 'originates' from a Contributor if it was added to the Program by such Contributor itself or anyone acting on such Contributor's behalf. Contributions do not include additions to the Program which: (i) are separate modules of software distributed in conjunction with the Program under their own license agreement, and (ii) are not derivative works of the Program.

"Contributor" means any person or entity that distributes the Program.

"Licensed Patents " mean patent claims licensable by a Contributor which are necessarily infringed by the use or sale of its Contribution alone or when combined with the Program.

"Program" means the Contributions distributed in accordance with this Agreement.

"Recipient" means anyone who receives the Program under this Agreement, including all Contributors.

GRANT OF RIGHTS

Subject to the terms of this Agreement, each Contributor hereby grants Recipient a non-exclusive, world-wide, royalty-free copyright license to reproduce, prepare derivative works of, publicly display, publicly perform, distribute and sublicense the Contribution of such Contributor, if any, and such derivative works, in source code and object code form.

Subject to the terms of this Agreement, each Contributor hereby grants Recipient a non-exclusive, world-wide, royalty-free patent license under Licensed Patents to make, use, sell, offer to sell, import and oth-erwise transfer the Contribution of such Contributor, if any, in source code and object code form. This patent license shall apply to the combination of the Contribution and the Program if, at the time the Con-tribution is added by the Contributor, such addition of the Contribution causes such combination to be covered by the Licensed Patents. The patent license shall not apply to any other combinations which include the Contribution. No hardware per se is licensed hereunder.

Recipient understands that although each Contributor grants the licenses to its Contributions set forth herein, no assurances are provided by any Contributor that the Program does not infringe the patent or other intellectual property rights of any other entity. Each Contributor disclaims any liability to Recipient for claims brought by any other entity based on infringement of intellectual property rights or otherwise. As a condition to exercising the rights and licenses granted hereunder, each Recipient hereby assumes sole responsibility to secure any other intellectual property rights needed, if any. For example, if a third party patent license is required to allow Recipient to distribute the Program, it is Recipient's responsibil-ity to acquire that license before distributing the Program.

Each Contributor represents that to its knowledge it has sufficient copyright rights in its Contribution, if any, to grant the copyright license set forth in this Agreement.

REQUIREMENTS

A Contributor may choose to distribute the Program in object code form under its own license agreement, provided that:

it complies with the terms and conditions of this Agreement; and

its license agreement:

effectively disclaims on behalf of all Contributors all warranties and conditions, express and implied, including warranties or conditions of title and non-infringement, and implied warranties or conditions of merchantability and fitness for a particular purpose;

effectively excludes on behalf of all Contributors all liability for damages, including direct, indirect, spe-cial, incidental and consequential damages, such as lost profits;

states that any provisions which differ from this Agreement are offered by that Contributor alone and not by any other party; and

states that source code for the Program is available from such Contributor, and informs licensees how to obtain it in a reasonable manner on or through a medium customarily used for software exchange.

When the Program is made available in source code form:

it must be made available under this Agreement; and

a copy of this Agreement must be included with each copy of the Program.

Contributors may not remove or alter any copyright notices contained within the Program.

Each Contributor must identify itself as the originator of its Contribution, if any, in a manner that reason-ably allows subsequent Recipients to identify the originator of the Contribution.

COMMERCIAL DISTRIBUTION

Commercial distributors of software may accept certain responsibilities with respect to end users, busi-ness partners and the like. While this license is intended to facilitate the commercial use of the Program, the Contributor who includes the Program in a commercial product offering should do so in a manner which does not create potential liability for other Contributors. Therefore, if a Contributor includes the Program in a commercial product offering, such Contributor ("Commercial Contributor") hereby agrees to defend and indemnify every other Contributor ("Indemnified Contributor") against any losses, dam-ages and costs (collectively "Losses") arising from claims, lawsuits and other legal actions brought by a third party against the Indemnified Contributor to the extent caused by the acts or omissions of such Commercial Contributor in connection with its distribution of the Program in a commercial product offer-ing. The obligations in this section do not apply to any claims or Losses relating to any actual or alleged intellectual property infringement. In order to qualify, an Indemnified Contributor must: a) promptly notify the Commercial Contributor in writing of such claim, and b) allow the Commercial Contributor to control, and cooperate with the Commercial Contributor in, the defense and any related settlement negotiations. The Indemnified Contributor may participate in any such claim at its own expense.

For example, a Contributor might include the Program in a commercial product offering, Product X. That Contributor is then a Commercial Contributor. If that Commercial Contributor then makes performance claims, or offers warranties related to Product X, those performance claims and warranties are such Com-mercial Contributor's responsibility alone. Under this section, the Commercial Contributor would have to defend claims against the other Contributors related to those performance claims and warranties, and if a court requires any other Contributor to pay any damages as a result, the Commercial Contributor must pay those damages.

NO WARRANTY

EXCEPT AS EXPRESSLY SET FORTH IN THIS AGREEMENT, THE PROGRAM IS PROVIDED ON AN "AS IS" BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY KIND, EITHER EXPRESS OR IMPLIED INCLUDING, WITHOUT LIMITATION, ANY WARRANTIES OR CONDITIONS OF TITLE, NON-INFRINGEMENT, MERCHANT-ABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Each Recipient is solely responsible for determining the appropriateness of using and distributing the Program and assumes all risks associated with its exer-cise of rights under this Agreement, including but not limited to the risks and costs of program errors, compliance with applicable laws, damage to or loss of data, programs or equipment, and unavailability or interruption of operations.

DISCLAIMER OF LIABILITY

EXCEPT AS EXPRESSLY SET FORTH IN THIS AGREEMENT, NEITHER RECIPIENT NOR ANY CONTRIBUTORS SHALL HAVE ANY LIABILITY FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSE-QUENTIAL DAMAGES (INCLUDING WITHOUT LIMITATION LOST PROFITS), HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OR DISTRIBUTION OF THE PROGRAM OR THE EXERCISE OF ANY RIGHTS GRANTED HEREUNDER, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAM-AGES.

GENERAL

If any provision of this Agreement is invalid or unenforceable under applicable law, it shall not affect the validity or enforceability of the remainder of the terms of this Agreement, and without further action by the parties hereto, such provision shall be reformed to the minimum extent necessary to make such pro-vision valid and enforceable.

If Recipient institutes patent litigation against a Contributor with respect to a patent applicable to soft-ware (including a cross-claim or counterclaim in a lawsuit), then any patent licenses granted by that Con-tributor to such Recipient under this Agreement shall terminate as of the date such litigation is filed. In addition, if Recipient institutes patent litigation against any entity (including a cross-claim or counter-claim in a lawsuit) alleging that the Program itself (excluding combinations of the Program with other software or hardware) infringes such Recipient's patent(s), then such Recipient's rights granted under Section 2(b) shall terminate as of the date such litigation is filed.

All Recipient's rights under this Agreement shall terminate if it fails to comply with any of the material terms or conditions of this Agreement and does not cure such failure in a reasonable period of time after becoming aware of such noncompliance. If all Recipient's rights under this Agreement terminate, Recipi-ent agrees to cease use and distribution of the Program as soon as reasonably practicable. However, Recipient's obligations under this Agreement and any licenses granted by Recipient relating to the Pro-gram shall continue and survive.

Everyone is permitted to copy and distribute copies of this Agreement, but in order to avoid inconsis-tency the Agreement is copyrighted and may only be modified in the following manner. The Agreement Steward reserves the right to publish new versions (including revisions) of this Agreement from time to time. No one other than the Agreement Steward has the right to modify this Agreement. IBM is the initial Agreement Steward. IBM may assign the responsibility to serve as the Agreement Steward to a suitable separate entity. Each new version of the Agreement will be given a distinguishing version number. The Program (including Contributions) may always be distributed subject to the version of the Agreement under which it was received. In addition, after a new version of the Agreement is published, Contributor may elect to distribute the Program (including its Contributions) under the new version. Except as expressly stated in Sections 2(a) and 2(b) above, Recipient receives no rights or licenses to the intellectual property of any Contributor under this Agreement, whether expressly, by implication, estoppel or other-wise. All rights in the Program not expressly granted under this Agreement are reserved.

This Agreement is governed by the laws of the State of New York and the intellectual property laws of the United States of America. No party to this Agreement will bring a legal action under this Agreement more than one year after the cause of action arose. Each party waives its rights to a jury trial in any resulting lit-igation.

9.4.10 BSD License

The PostgreSQL JDBC driver is distributed under the BSD license, same as the server. The simplest expla-nation of the licensing terms is that you can do whatever you want with the product and source code as long as you don't claim you wrote it or sue us. You should give it a read though, it's only half a page.

Copyright (c) 1997-2008, PostgreSQL Global Development Group

All rights reserved.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

1. Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.

2. Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.

3. Neither the name of the PostgreSQL Global Development Group nor the names of its contributors may be used to endorse or promote products derived from this software without specific prior written per-mission.

THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDERS AND CONTRIBUTORS "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE COPYRIGHT OWNER OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOW-EVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

9.4.11 Open Software License v. 3.0

This Open Software License (the "License") applies to any original work of authorship (the "Original Work") whose owner (the "Licensor") has placed the following licensing notice adjacent to the copyright notice for the Original Work:

Licensed under the Open Software License version 3.0

1) *Grant of Copyright License.* Licensor grants You a worldwide, royalty-free, non-exclusive, sublicens-able license, for the duration of the copyright, to do the following:

a) to reproduce the Original Work in copies, either alone or as part of a collective work;

b) to translate, adapt, alter, transform, modify, or arrange the Original Work, thereby creating derivative works ("Derivative Works") based upon the Original Work;

c) to distribute or communicate copies of the Original Work and Derivative Works to the public, with the proviso that copies of Original Work or Derivative Works that You distribute or communicate shall be licensed under this Open Software License;

d) to perform the Original Work publicly; and

e) to display the Original Work publicly.

2) *Grant of Patent License.* Licensor grants You a worldwide, royalty-free, non-exclusive, sublicensable license, under patent claims owned or controlled by the Licensor that are embodied in the Original Work as furnished by the Licensor, for the duration of the patents, to make, use, sell, offer for sale, have made, and import the Original Work and Derivative Works.

3) *Grant of Source Code License.* The term "Source Code" means the preferred form of the Original Work for making modifications to it and all available documentation describing how to modify the Original Work. Licensor agrees to provide a machine-readable copy of the Source Code of the Original Work along with each copy of the Original Work that Licensor distributes. Licensor reserves the right to satisfy this obligation by placing a machine-readable copy of the Source Code in an information repository reason-ably calculated to permit inexpensive and convenient access by You for as long as Licensor continues to distribute the Original Work.

4) *Exclusions From License Grant.* Neither the names of Licensor, nor the names of any contributors to the Original Work, nor any of their trademarks or service marks, may be used to endorse or promote products derived from this Original Work without express prior permission of the Licensor. Except as expressly stated herein, nothing in this License grants any license to Licensorï¿½s trademarks, copyrights, patents, trade secrets or any other intellectual property. No patent license is granted to make, use, sell, offer for sale, have made, or import embodiments of any patent claims other than the licensed claims defined in Section 2. No license is granted to the trademarks of Licensor even if such marks are included in the Original Work. Nothing in this License shall be interpreted to prohibit Licensor from licensing under terms different from this License any Original Work that Licensor otherwise would have a right to license.

5) *External Deployment.* The term "External Deployment" means the use, distribution, or communication of the Original Work or Derivative Works in any way such that the Original Work or Derivative Works may be used by anyone other than You, whether those works are distributed or communicated to those persons or made available as an application intended for use over a network. As an express condition for the grants of license hereunder, You must treat any External Deployment by You of the Original Work or a Derivative Work as a distribution under section 1(c).

6) *Attribution Rights.* You must retain, in the Source Code of any Derivative Works that You create, all copyright, patent, or trademark notices from the Source Code of the Original Work, as well as any notices of licensing and any descriptive text identified therein as an "Attribution Notice." You must cause the Source Code for any Derivative Works that You create to carry a prominent Attribution Notice reasonably calculated to inform recipients that You have modified the Original Work.

7) *Warranty of Provenance and Disclaimer of Warranty.* Licensor warrants that the copyright in and to the Original Work and the patent rights granted herein by Licensor are owned by the Licensor or are sub-licensed to You under the terms of this License with the permission of the contributor(s) of those copy-rights and patent rights. Except as expressly stated in the immediately preceding sentence, the Original Work is provided under this License on an "AS IS" BASIS and WITHOUT WARRANTY, either express or implied, including, without limitation, the warranties of non-infringement, merchantability or fitness for a particular purpose. THE ENTIRE RISK AS TO THE QUALITY OF THE ORIGINAL WORK IS WITH YOU. This DISCLAIMER OF WARRANTY constitutes an essential part of this License. No license to the Original Work is granted by this License except under this disclaimer.

8) *Limitation of Liability.* Under no circumstances and under no legal theory, whether in tort (including negligence), contract, or otherwise, shall the Licensor be liable to anyone for any indirect, special, inci-dental, or consequential damages of any character arising as a result of this License or the use of

the Orig-inal Work including, without limitation, damages for loss of goodwill, work stoppage, computer failure or malfunction, or any and all other commercial damages or losses. This limitation of liability shall not apply to the extent applicable law prohibits such limitation.

9) *Acceptance and Termination.* If, at any time, You expressly assented to this License, that assent indi-cates your clear and irrevocable acceptance of this License and all of its terms and conditions. If You dis-tribute or communicate copies of the Original Work or a Derivative Work, You must make a reasonable effort under the circumstances to obtain the express assent of recipients to the terms of this License. This License conditions your rights to undertake the activities listed in Section 1, including your right to create Derivative Works based upon the Original Work, and doing so without honoring these terms and condi-tions is prohibited by copyright law and international treaty. Nothing in this License is intended to affect copyright exceptions and limitations (including ï¿½fair useï¿½ or ï¿½fair dealingï¿½). This License shall terminate immediately and You may no longer exercise any of the rights granted to You by this License upon your failure to honor the conditions in Section 1(c).

10) *Termination for Patent Action.* This License shall terminate automatically and You may no longer exercise any of the rights granted to You by this License as of the date You commence an action, including a cross-claim or counterclaim, against Licensor or any licensee alleging that the Original Work infringes a patent. This termination provision shall not apply for an action alleging patent infringement by combina-tions of the Original Work with other software or hardware.

11) *Jurisdiction, Venue and Governing Law.* Any action or suit relating to this License may be brought only in the courts of a jurisdiction wherein the Licensor resides or in which Licensor conducts its primary business, and under the laws of that jurisdiction excluding its conflict-of-law provisions. The application of the United Nations Convention on Contracts for the International Sale of Goods is expressly excluded. Any use of the Original Work outside the scope of this License or after its termination shall be subject to the requirements and penalties of copyright or patent law in the appropriate jurisdiction. This section shall survive the termination of this License.

12) *Attorneysï¿½ Fees.* In any action to enforce the terms of this License or seeking damages relating thereto, the prevailing party shall be entitled to recover its costs and expenses, including, without limita-tion, reasonable attorneys' fees and costs incurred in connection with such action, including any appeal of such action. This section shall survive the termination of this License.

13) *Miscellaneous.* If any provision of this License is held to be unenforceable, such provision shall be reformed only to the extent necessary to make it enforceable.

14) *Definition of "You" in This License.* "You" throughout this License, whether in upper or lower case, means an individual or a legal entity exercising rights under, and complying with all of the terms of, this License. For legal entities, "You" includes any entity that controls, is controlled by, or is under common control with you. For purposes of this definition, "control" means (i) the power, direct or indirect, to cause the direction or management of such entity, whether by contract or otherwise, or (ii) ownership of fifty percent (50%) or more of the outstanding shares, or (iii) beneficial ownership of such entity.

15) *Right to Use.* You may use the Original Work in all ways not otherwise restricted or conditioned by this License or by law, and Licensor promises not to interfere with or be responsible for such uses by You.

16) *Modification of This License.* This License is Copyright ï¿½ 2005 Lawrence Rosen. Permission is granted to copy, distribute, or communicate this License without modification. Nothing in this License permits You to modify this License as applied to the Original Work or to Derivative Works. However, You may modify the text of this License and copy, distribute or communicate your modified version (the "Modified License") and apply it to other original works of authorship subject to the following conditions: (i) You may not indicate in any way that your Modified License is the "Open Software License" or "OSL" and you may not use those names in the name of your Modified License; (ii) You must replace the notice specified in the first paragraph above with the notice "Licensed under <insert your license name here>" or with a notice of your own that is not confusingly similar to the notice in

this License; and (iii) You may not claim that your original works are open source software unless your Modified License has been approved by Open Source Initiative (OSI) and You comply with its license review and certification process.

9.4.12 Mozilla Public License Version 1.1

1. Definitions.

1.0.1. "Commercial Use"

means distribution or otherwise making the Covered Code available to a third party.

1.1. "Contributor"

means each entity that creates or contributes to the creation of Modifications.

1.2. "Contributor Version"

means the combination of the Original Code, prior Modifications used by a Contributor, and the Modifi-cations made by that particular Contributor.

1.3. "Covered Code"

means the Original Code or Modifications or the combination of the Original Code and Modifications, in each case including portions thereof.

1.4. "Electronic Distribution Mechanism"

means a mechanism generally accepted in the software development community for the electronic transfer of data.

1.5. "Executable"

means Covered Code in any form other than Source Code.

1.6. "Initial Developer"

means the individual or entity identified as the Initial Developer in the Source Code notice required by Exhibit A.

1.7. "Larger Work"

means a work which combines Covered Code or portions thereof with code not governed by the terms of this License.

1.8. "License"

means this document.

1.8.1. "Licensable"

means having the right to grant, to the maximum extent possible, whether at the time of the initial grant or subsequently acquired, any and all of the rights conveyed herein.

1.9. "Modifications"

means any addition to or deletion from the substance or structure of either the Original Code or any pre-vious Modifications. When Covered Code is released as a series of files, a Modification is:

Any addition to or deletion from the contents of a file containing Original Code or previous Modifications.

Any new file that contains any part of the Original Code or previous Modifications.

1.10. "Original Code"

means Source Code of computer software code which is described in the Source Code notice required by Exhibit A as Original Code, and which, at the time of its release under this License is not already Covered Code governed by this License.

1.10.1. "Patent Claims"

means any patent claim(s), now owned or hereafter acquired, including without limitation, method, pro-cess, and apparatus claims, in any patent Licensable by grantor.

1.11. "Source Code"

means the preferred form of the Covered Code for making modifications to it, including all modules it contains, plus any associated interface definition files, scripts used to control compilation and installation of an Executable, or source code differential comparisons against either the Original Code or another well known, available Covered Code of the Contributor's choice. The Source Code can be in a compressed or archival form, provided the appropriate decompression or de-archiving software is widely available for no charge.

1.12. "You" (or "Your")

means an individual or a legal entity exercising rights under, and complying with all of the terms of, this License or a future version of this License issued under Section 6.1. For legal entities, "You" includes any entity which controls, is controlled by, or is under common control with You. For purposes of this defini-tion, "control" means (a) the power, direct or indirect, to cause the direction or management of such entity, whether by contract or otherwise, or (b) ownership of more than fifty percent (50%) of the out-standing shares or beneficial ownership of such entity.

2. Source Code License.

2.1. The Initial Developer Grant.

The Initial Developer hereby grants You a world-wide, royalty-free, non-exclusive license, subject to third party intellectual property claims:

under intellectual property rights (other than patent or trademark) Licensable by Initial Developer to use, reproduce, modify, display, perform, sublicense and distribute the Original Code (or portions thereof) with or without Modifications, and/or as part of a Larger Work; and under Patents Claims infringed by the making, using or selling of Original Code, to make, have made, use, practice, sell, and offer for sale, and/or otherwise dispose of the Original Code (or portions thereof).

the licenses granted in this Section 2.1 (a) and (b) are effective on the date Initial Developer first distributes Original Code under the terms of this License.

Notwithstanding Section 2.1 (b) above, no patent license is granted: 1) for code that You delete from the Original Code; 2) separate from the Original Code; or 3) for infringements caused by: i) the modification of the Original Code or ii) the combination of the Original Code with other software or devices.

2.2. Contributor Grant.

Subject to third party intellectual property claims, each Contributor hereby grants You a world-wide, roy-alty-free, non-exclusive license

under intellectual property rights (other than patent or trademark) Licensable by Contributor, to use, reproduce, modify, display, perform, sublicense and distribute the Modifications created by such Con-tributor (or portions thereof) either on an unmodified basis, with other Modifications, as Covered Code and/or as part of a Larger Work; and

under Patent Claims infringed by the making, using, or selling of Modifications made by that Contributor either alone and/or in combination with its Contributor Version (or portions of such combination), to make, use, sell, offer for sale, have made, and/or otherwise dispose of: 1) Modifications made by that Contributor (or portions thereof); and 2) the combination of Modifications made by that Contributor with its Contributor Version (or portions of such combination).

the licenses granted in Sections 2.2 (a) and 2.2 (b) are effective on the date Contributor first makes Commercial Use of the Covered Code.

Notwithstanding Section 2.2 (b) above, no patent license is granted: 1) for any code that Contributor has deleted from the Contributor Version; 2) separate from the Contributor Version; 3) for infringements caused by: i) third party modifica-tions of Contributor Version or ii) the combination of Modifications made by that Contributor with other software (except as part of the Contributor Version) or other devices; or 4) under Patent Claims infringed by Covered Code in the absence of Modifications made by that Contributor.

3. Distribution Obligations.

3.1. Application of License.

The Modifications which You create or to which You contribute are governed by the terms of this License, including without limitation Section 2.2. The Source Code version of Covered Code may be distributed only under the terms of this License or a future version of this License released under Section 6.1, and You must include a copy of this License with every copy of the Source Code You distribute. You may not offer or impose any terms on any Source Code version that alters or restricts the applicable version of this License or the recipients' rights hereunder. However, You may include an additional document offering the additional rights described in Section 3.5.

3.2. Availability of Source Code.

Any Modification which You create or to which You contribute must be made available in Source Code form under the terms of this License either on the same media as an Executable version or via an accepted Electronic Distribution Mechanism to anyone to whom you made an Executable version avail-able; and if made available via Electronic Distribution Mechanism, must remain available for at least twelve (12) months after the date it initially became available, or at least six (6) months after a subse-quent version of that particular Modification has been made available to such recipients. You are respon-sible for ensuring that the Source Code version remains available even if the Electronic Distribution Mechanism is maintained by a third party.

3.3. Description of Modifications.

You must cause all Covered Code to which You contribute to contain a file documenting the changes You made to create that Covered Code and the date of any change. You must include a prominent statement that the Modification is derived, directly or indirectly, from Original Code provided by the Initial Devel-oper and including the name of the Initial Developer in (a) the Source Code, and (b) in any notice in an Executable version or related documentation in which You describe the origin or ownership of the Cov-ered Code.

3.4. Intellectual Property Matters

(a) Third Party Claims

If Contributor has knowledge that a license under a third party's intellectual property rights is required to exercise the rights granted by such Contributor under Sections 2.1 or 2.2, Contributor must include a text file with the Source Code distribution titled "LEGAL" which describes the claim and the party making the claim in sufficient detail that a recipient will know whom to contact. If Contributor obtains such knowl-edge after the Modification is made available as described in Section 3.2, Contributor

shall promptly modify the LEGAL file in all copies Contributor makes available thereafter and shall take other steps (such as notifying appropriate mailing lists or newsgroups) reasonably calculated to inform those who received the Covered Code that new knowledge has been obtained.

(b) Contributor APIs

If Contributor's Modifications include an application programming interface and Contributor has knowl-edge of patent licenses which are reasonably necessary to implement that API, Contributor must also include this information in the *legal* file.

(c) Representations.

Contributor represents that, except as disclosed pursuant to Section 3.4 (a) above, Contributor believes that Contributor's Modifications are Contributor's original creation(s) and/or Contributor has sufficient rights to grant the rights conveyed by this License.

3.5. Required Notices.

You must duplicate the notice in Exhibit A in each file of the Source Code. If it is not possible to put such notice in a particular Source Code file due to its structure, then You must include such notice in a location (such as a relevant directory) where a user would be likely to look for such a notice. If You created one or more Modification(s) You may add your name as a Contributor to the notice described in Exhibit A. You must also duplicate this License in any documentation for the Source Code where You describe recipi-ents' rights or ownership rights relating to Covered Code. You may choose to offer, and to charge a fee for, warranty, support, indemnity or liability obligations to one or more recipients of Covered Code. How-ever, You may do so only on Your own behalf, and not on behalf of the Initial Developer or any Contribu-tor. You must make it absolutely clear than any such warranty, support, indemnity or liability obligation is offered by You alone, and You hereby agree to indemnify the Initial Developer and every Contributor for any liability incurred by the Initial Developer or such Contributor as a result of warranty, support, indem-nity or liability terms You offer.

3.6. Distribution of Executable Versions.

You may distribute Covered Code in Executable form only if the requirements of Sections 3.1, 3.2, 3.3, 3.4 and 3.5 have been met for that Covered Code, and if You include a notice stating that the Source Code version of the Covered Code is available under the terms of this License, including a description of how and where You have fulfilled the obligations of Section 3.2. The notice must be conspicuously included in any notice in an Executable version, related documentation or collateral in which You describe recipients' rights relating to the Covered Code. You may distribute the Executable version of Covered Code or own-ership rights under a license of Your choice, which may contain terms different from this License, pro-vided that You are in compliance with the terms of this License and that the license for the Executable version does not attempt to limit or alter the recipient's rights in the Source Code version from the rights set forth in this License. If You distribute the Executable version under a different license You must make it absolutely clear that any terms which differ from this License are offered by You alone, not by the Initial Developer or any Contributor. You hereby agree to indemnify the Initial Developer and every Contributor for any liability incurred by the Initial Developer or such Contributor as a result of any such terms You offer.

3.7. Larger Works.

You may create a Larger Work by combining Covered Code with other code not governed by the terms of this License and distribute the Larger Work as a single product. In such a case, You must make sure the requirements of this License are fulfilled for the Covered Code.

4. Inability to Comply Due to Statute or Regulation.

If it is impossible for You to comply with any of the terms of this License with respect to some or all of the Covered Code due to statute, judicial order, or regulation then You must: (a) comply with the terms of this License to the maximum extent possible; and (b) describe the limitations and the code

they affect. Such description must be included in the *legal* file described in Section 3.4 and must be included with all distributions of the Source Code. Except to the extent prohibited by statute or regulation, such descrip-tion must be sufficiently detailed for a recipient of ordinary skill to be able to understand it.

5. Application of this License.

This License applies to code to which the Initial Developer has attached the notice in Exhibit A and to related Covered Code.

6. Versions of the License.

6.1. New Versions

Netscape Communications Corporation ("Netscape") may publish revised and/or new versions of the License from time to time. Each version will be given a distinguishing version number.

6.2. Effect of New Versions

Once Covered Code has been published under a particular version of the License, You may always con-tinue to use it under the terms of that version. You may also choose to use such Covered Code under the terms of any subsequent version of the License published by Netscape. No one other than Netscape has the right to modify the terms applicable to Covered Code created under this License.

6.3. Derivative Works

If You create or use a modified version of this License (which you may only do in order to apply it to code which is not already Covered Code governed by this License), You must (a) rename Your license so that the phrases "Mozilla", "MOZILLAPL", "MOZPL", "Netscape", "MPL", "NPL" or any confusingly similar phrase do not appear in your license (except to note that your license differs from this License) and (b) otherwise make it clear that Your version of the license contains terms which differ from the Mozilla Pub-lic License and Netscape Public License. (Filling in the name of the Initial Developer, Original Code or Con-tributor in the notice described in Exhibit A shall not of themselves be deemed to be modifications of this License.)

7. Disclaimer of warranty

Covered code is provided under this license on an "as is" basis, without warranty of any kind, either expressed or implied, including, without limitation, warranties that the covered code is free of defects, merchantable, fit for a particular purpose or non-infringing. The entire risk as to the quality and perfor-mance of the covered code is with you. Should any covered code prove defective in any respect, you (not the initial developer or any other contributor) assume the cost of any necessary servicing, repair or cor-rection. This disclaimer of warranty constitutes an essential part of this license. No use of any covered code is authorized hereunder except under this disclaimer.

8. Termination

8.1. This License and the rights granted hereunder will terminate automatically if You fail to comply with terms herein and fail to cure such breach within 30 days of becoming aware of the breach. All sublicenses to the Covered Code which are properly granted shall survive any termination of this License. Provisions which, by their nature, must remain in effect beyond the termination of this License shall survive.

8.2. If You initiate litigation by asserting a patent infringement claim (excluding declatory judgment actions) against Initial Developer or a Contributor (the Initial Developer or Contributor against whom You file such action is referred to as "Participant") alleging that:

such Participant's Contributor Version directly or indirectly infringes any patent, then any and all rights granted by such Participant to You under Sections 2.1 and/or 2.2 of this License shall, upon 60 days notice from Participant terminate prospectively, unless if within 60 days after receipt of notice You

either: (i) agree in writing to pay Participant a mutually agreeable reasonable royalty for Your past and future use of Modifications made by such Participant, or (ii) withdraw Your litigation claim with respect to the Con-tributor Version against such Participant. If within 60 days of notice, a reasonable royalty and payment arrangement are not mutually agreed upon in writing by the parties or the litigation claim is not with-drawn, the rights granted by Participant to You under Sections 2.1 and/or 2.2 automatically terminate at the expiration of the 60 day notice period specified above.

any software, hardware, or device, other than such Participant's Contributor Version, directly or indirectly infringes any patent, then any rights granted to You by such Participant under Sections 2.1(b) and 2.2(b) are revoked effective as of the date You first made, used, sold, distributed, or had made, Modifications made by that Participant.

8.3. If You assert a patent infringement claim against Participant alleging that such Participant's Contribu-tor Version directly or indirectly infringes any patent where such claim is resolved (such as by license or settlement) prior to the initiation of patent infringement litigation, then the reasonable value of the licenses granted by such Participant under Sections 2.1 or 2.2 shall be taken into account in determining the amount or value of any payment or license.

8.4. In the event of termination under Sections 8.1 or 8.2 above, all end user license agreements (exclud-ing distributors and resellers) which have been validly granted by You or any distributor hereunder prior to termination shall survive termination.

9. Limitation of liability

Under no circumstances and under no legal theory, whether tort (including negligence), contract, or oth-erwise, shall you, the initial developer, any other contributor, or any distributor of covered code, or any supplier of any of such parties, be liable to any person for any indirect, special, incidental, or consequen-tial damages of any character including, without limitation, damages for loss of goodwill, work stoppage, computer failure or malfunction, or any and all other commercial damages or losses, even if such party shall have been informed of the possibility of such damages. This limitation of liability shall not apply to liability for death or personal injury resulting from such party's negligence to the extent applicable law prohibits such limitation. Some jurisdictions do not allow the exclusion or limitation of incidental or con-sequential damages, so this exclusion and limitation may not apply to you.

10. U.S. government end users

The Covered Code is a "commercial item," as that term is defined in 48 C.F.R. 2.101 (Oct. 1995), consist-ing of "commercial computer software" and "commercial computer software documentation," as such terms are used in 48 C.F.R. 12.212 (Sept. 1995). Consistent with 48 C.F.R. 12.212 and 48 C.F.R. 227.7202-1 through 227.7202-4 (June 1995), all U.S. Government End Users acquire Covered Code with only those rights set forth herein.

11. Miscellaneous

This License represents the complete agreement concerning subject matter hereof. If any provision of this License is held to be unenforceable, such provision shall be reformed only to the extent necessary to make it enforceable. This License shall be governed by California law provisions (except to the extent applicable law, if any, provides otherwise), excluding its conflict-of-law provisions. With respect to dis-putes in which at least one party is a citizen of, or an entity chartered or registered to do business in the United States of America, any litigation relating to this License shall be subject to the jurisdiction of the Federal Courts of the Northern District of California, with venue lying in Santa Clara County, California, with the losing party responsible for costs, including without limitation, court costs and reasonable attor-neys' fees and expenses. The application of the United Nations Convention on Contracts for the Interna-tional Sale of Goods is expressly excluded. Any law or regulation which provides that the language of a contract shall be construed against the drafter shall not apply to this License.

12. Responsibility for claims

As between Initial Developer and the Contributors, each party is responsible for claims and damages aris-ing, directly or indirectly, out of its utilization of rights under this License and You agree to work with Ini-tial Developer and Contributors to distribute such responsibility on an equitable basis. Nothing herein is intended or shall be deemed to constitute any admission of liability.

13. Multiple-licensed code

Initial Developer may designate portions of the Covered Code as "Multiple-Licensed". "Multiple-Licensed" means that the Initial Developer permits you to utilize portions of the Covered Code under Your choice of the MPL or the alternative licenses, if any, specified by the Initial Developer in the file described in Exhibit A.

Exhibit A - Mozilla Public License.

"The contents of this file are subject to the Mozilla Public License Version 1.1 (the "License"); you may not use this file except in compliance with the License. You may obtain a copy of the License at http://www.mozilla.org/MPL/

Software distributed under the License is distributed on an "AS IS" basis, WITHOUT WARRANTY OF ANY KIND, either express or implied. See the License for the specific language governing rights and limitations

under the License.

The Original Code is _____.

The Initial Developer of the Original Code is _____.

Portions created by _____ are Copyright (C) _____ _____. All Rights Reserved.

Contributor(s): _____.

Alternatively, the contents of this file may be used under the terms of the _____ license (the "[___] License"), in which case the provisions of [_____] License are applicable instead of those above. If you wish to allow use of your version of this file only under the terms of the [____] License and not to allow others to use your version of this file under the MPL, indicate your decision by deleting the provisions above and replace them with the notice and other provisions required by the [___] License. If you do not delete the provisions above, a recipient may use your version of this file under either the MPL or the [___] License."

NOTE: The text of this Exhibit A may differ slightly from the text of the notices in the Source Code files of the Original Code. You should use the text of this Exhibit A rather than the text found in the Original Code Source Code for Your Modifications.

9.4.13 GNU GENERAL PUBLIC LICENSE

Version 2, June 1991

Copyright (C) 1989, 1991 Free Software Foundation, Inc.

675 Mass Ave, Cambridge, MA 02139, USA

Everyone is permitted to copy and distribute verbatim copies of this license document, but changing it is not allowed.

Preamble

The licenses for most software are designed to take away your freedom to share and change it. By con-trast, the GNU General Public License is intended to guarantee your freedom to share and change free software--to make sure the software is free for all its users. This General Public License applies to most of the Free Software Foundation's software and to any other program whose authors commit to

using it. (Some other Free Software Foundation software is covered by the GNU Library General Public License instead.) You can apply it to your programs, too.

When we speak of free software, we are referring to freedom, not price. Our General Public Licenses are designed to make sure that you have the freedom to distribute copies of free software (and charge for this service if you wish), that you receive source code or can get it if you want it, that you can change the software or use pieces of it in new free programs; and that you know you can do these things.

To protect your rights, we need to make restrictions that forbid anyone to deny you these rights or to ask you to surrender the rights. These restrictions translate to certain responsibilities for you if you distribute copies of the software, or if you modify it.

For example, if you distribute copies of such a program, whether gratis or for a fee, you must give the recipients all the rights that you have. You must make sure that they, too, receive or can get the source code. And you must show them these terms so they know their rights.

We protect your rights with two steps: (1) copyright the software, and (2) offer you this license which gives you legal permission to copy, distribute and/or modify the software.

Also, for each author's protection and ours, we want to make certain that everyone understands that there is no warranty for this free software. If the software is modified by someone else and passed on, we want its recipients to know that what they have is not the original, so that any problems introduced by others will not reflect on the original authors' reputations.

Finally, any free program is threatened constantly by software patents. We wish to avoid the danger that redistributors of a free program will individually obtain patent licenses, in effect making the program pro-prietary. To prevent this, we have made it clear that any patent must be licensed for everyone's free use or not licensed at all.

The precise terms and conditions for copying, distribution and modification follow.

GNU GENERAL PUBLIC LICENSE

TERMS AND CONDITIONS FOR COPYING, DISTRIBUTION AND MODIFICATION

0. This License applies to any program or other work which contains a notice placed by the copyright holder saying it may be distributed under the terms of this General Public License. The "Program", below, refers to any such program or work, and a "work based on the Program" means either the Pro-gram or any derivative work under copyright law: that is to say, a work containing the Program or a por-tion of it, either verbatim or with modifications and/or translated into another language. (Hereinafter, translation is included without limitation in the term "modification".) Each licensee is addressed as "you".

Activities other than copying, distribution and modification are not covered by this License; they are out-side its scope. The act of running the Program is not restricted, and the output from the Program

is covered only if its contents constitute a work based on the Program (independent of having been made by running the Program). Whether that is true depends on what the Program does.

1. You may copy and distribute verbatim copies of the Program's source code as you receive it, in any medium, provided that you conspicuously and appropriately publish on each copy an appropriate copy-right notice and disclaimer of warranty; keep intact all the notices that refer to this License and to the absence of any warranty; and give any other recipients of the Program a copy of this License along with the Program.

You may charge a fee for the physical act of transferring a copy, and you may at your option offer warranty protection in exchange for a fee.

2. You may modify your copy or copies of the Program or any portion of it, thus forming a work based on the Program, and copy and distribute such modifications or work under the terms of Section 1 above, provided that you also meet all of these conditions:

a) You must cause the modified files to carry prominent notices stating that you changed the files and the date of any change.

b) You must cause any work that you distribute or publish, that in whole or in part contains or is derived from the Program or any part thereof, to be licensed as a whole at no charge to all third parties under the terms of this License.

c) If the modified program normally reads commands interactively when run, you must cause it, when started running for such interactive use in the most ordinary way, to print or display an announcement including an appropriate copyright notice and a notice that there is no warranty (or else, saying that you provide a warranty) and that users may redistribute the program under these conditions, and telling the user how to view a copy of this License. (Exception: if the Program itself is interactive but does not nor-mally print such an announcement, your work based on the Program is not required to print an announcement.)

These requirements apply to the modified work as a whole. If identifiable sections of that work are not derived from the Program, and can be reasonably considered independent and separate works in them-selves, then this License, and its terms, do not apply to those sections when you distribute them as sepa-rate works. But when you distribute the same sections as part of a whole which is a work based on the Program, the distribution of the whole must be on the terms of this License, whose permissions for other licensees extend to the entire whole, and thus to each and every part regardless of who wrote it.

Thus, it is not the intent of this section to claim rights or contest your rights to work written entirely by you; rather, the intent is to exercise the right to control the distribution of derivative or collective works based on the Program.

In addition, mere aggregation of another work not based on the Program with the Program (or with a work based on the Program) on a volume of a storage or distribution medium does not bring the other work under the scope of this License.

3. You may copy and distribute the Program (or a work based on it, under Section 2) in object code or executable form under the terms of Sections 1 and 2 above provided that you also do one of the follow-ing:

a) Accompany it with the complete corresponding machine-readable source code, which must be distributed under the terms of Sections 1 and 2 above on a medium customarily used for software interchange; or,

b) Accompany it with a written offer, valid for at least three years, to give any third party, for a charge no more than your cost of physically performing source distribution, a complete machine-readable copy of the corresponding source code, to be distributed under the terms of Sections 1 and 2 above on a medium customarily used for software interchange; or,

c) Accompany it with the information you received as to the offer to distribute corresponding source code. (This alternative is allowed only for noncommercial distribution and only if you received the pro-gram in object code or executable form with such an offer, in accord with Subsection b above.)

The source code for a work means the preferred form of the work for making modifications to it. For an executable work, complete source code means all the source code for all modules it contains, plus any associated interface definition files, plus the scripts used to control compilation and installation of the executable. However, as a special exception, the source code distributed need not include anything that is normally distributed (in either source or binary form) with the major components (compiler, kernel, and so on) of the operating system on which the executable runs, unless that component itself accompa-nies the executable.

If distribution of executable or object code is made by offering access to copy from a designated place, then offering equivalent access to copy the source code from the same place counts as distribution of the source code, even though third parties are not compelled to copy the source along with the object code.

4. You may not copy, modify, sublicense, or distribute the Program except as expressly provided under this License. Any attempt otherwise to copy, modify, sublicense or distribute the Program is void, and will automatically terminate your rights under this License. However, parties who have received copies, or rights, from you under this License will not have their licenses terminated so long as such parties remain in full compliance.

5. You are not required to accept this License, since you have not signed it. However, nothing else grants you permission to modify or distribute the Program or its derivative works. These actions are pro-hibited by law if you do not accept this License. Therefore, by modifying or distributing the Program (or any work based on the Program), you indicate your acceptance of this License to do so, and all its terms and conditions for copying, distributing or modifying the Program or works based on it.

6. Each time you redistribute the Program (or any work based on the Program), the recipient automati-cally receives a license from the original licensor to copy, distribute or modify the Program subject to these terms and conditions. You may not impose any further restrictions on the recipients' exercise of the rights granted herein. You are not responsible for enforcing compliance by third parties to

this License.

7. If, as a consequence of a court judgment or allegation of patent infringement or for any other reason (not limited to patent issues), conditions are imposed on you (whether by court order, agreement or oth-erwise) that contradict the conditions of this License, they do not excuse you from the conditions of this License. If you cannot distribute so as to satisfy simultaneously your obligations under this License and any other pertinent obligations, then as a consequence you may not distribute the Program at all. For example, if a patent license would not permit royalty-free redistribution of the Program by all those who receive copies directly or indirectly through you, then the only way you could satisfy both it and this License would be to refrain entirely from distribution of the Program.

If any portion of this section is held invalid or unenforceable under any particular circumstance, the bal-ance of the section is intended to apply and the section as a whole is intended to apply in other circum-stances.

It is not the purpose of this section to induce you to infringe any patents or other property right claims or to contest validity of any such claims; this section has the sole purpose of protecting the integrity of the free software distribution system, which is implemented by public license practices. Many people have made generous contributions to the wide range of software distributed through that system in reliance on consistent application of that system; it is up to the author/donor to decide if he or she is willing to distribute software through any other system and a licensee cannot impose that choice.

This section is intended to make thoroughly clear what is believed to be a consequence of the rest of this License.

8. If the distribution and/or use of the Program is restricted in certain countries either by patents or by copyrighted interfaces, the original copyright holder who places the Program under this License may add an explicit geographical distribution limitation excluding those countries, so that distribution is permitted only in or among countries not thus excluded. In such case, this License incorporates the limitation as if written in the body of this License.

9. The Free Software Foundation may publish revised and/or new versions of the General Public License from time to time. Such new versions will be similar in spirit to the present version, but may differ in detail to address new problems or concerns.

Each version is given a distinguishing version number. If the Program specifies a version number of this License which applies to it and "any later version", you have the option of following the terms and condi-tions either of that version or of any later version published by the Free Software Foundation. If the Pro-gram does not specify a version number of this License, you may choose any version ever published by the Free Software Foundation.

10. If you wish to incorporate parts of the Program into other free programs whose distribution conditions are different, write to the author to ask for permission. For software which is copyrighted by the Free Software Foundation, write to the Free Software Foundation; we sometimes make exceptions for this. Our decision will be guided by the two goals of preserving the free status of all derivatives of our free software and of promoting the sharing and reuse of software generally.

NO WARRANTY

11. BECAUSE THE PROGRAM IS LICENSED FREE OF CHARGE, THERE IS NO WARRANTY FOR THE PRO-GRAM, TO THE EXTENT PERMITTED BY APPLICABLE LAW. EXCEPT WHEN OTHERWISE STATED IN WRIT-ING THE COPYRIGHT HOLDERS AND/OR OTHER PARTIES PROVIDE THE PROGRAM "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE. THE ENTIRE RISK AS TO THE QUALITY AND PERFORMANCE OF THE PROGRAM IS WITH YOU. SHOULD THE PROGRAM PROVE DEFECTIVE, YOU ASSUME THE COST OF ALL NECESSARY SERVICING, REPAIR OR CORRECTION.

12. IN NO EVENT UNLESS REQUIRED BY APPLICABLE LAW OR AGREED TO IN WRITING WILL ANY COPY-RIGHT HOLDER, OR ANY OTHER PARTY WHO MAY MODIFY AND/OR REDISTRIBUTE THE PROGRAM AS PERMITTED ABOVE, BE LIABLE TO YOU FOR DAMAGES, INCLUDING ANY GENERAL, SPECIAL, INCIDENTAL OR CONSEQUENTIAL DAMAGES ARISING OUT OF THE USE OR INABILITY TO USE THE PROGRAM (INCLUD-ING BUT NOT LIMITED TO LOSS OF DATA OR DATA BEING RENDERED INACCURATE OR LOSSES SUSTAINED BY YOU OR THIRD PARTIES OR A FAILURE OF THE PROGRAM TO OPERATE WITH ANY OTHER PROGRAMS), EVEN IF SUCH HOLDER OR OTHER PARTY HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

END OF TERMS AND CONDITIONS

Appendix: How to Apply These Terms to Your New Programs

If you develop a new program, and you want it to be of the greatest possible use to the public, the best way to achieve this is to make it free software which everyone can redistribute and change under these terms.

To do so, attach the following notices to the program. It is safest to attach them to the start of each source file to most effectively convey the exclusion of warranty; and each file should have at least the "copyright" line and a pointer to where the full notice is found.

<one line to give the program's name and a brief idea of what it does.>

Copyright (C) 19yy <name of author>

This program is free software; you can redistribute it and/or modify it under the terms of the GNU Gen-eral Public License as published by the Free Software Foundation; either version 2 of the License, or

(at your option) any later version.

This program is distributed in the hope that it will be useful, but WITHOUT ANY WARRANTY; without even the implied warranty of MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE. See the GNU General Public License for more details.

You should have received a copy of the GNU General Public License along with this program; if not, write to the Free Software Foundation, Inc., 675 Mass Ave, Cambridge, MA 02139, USA.

Also add information on how to contact you by electronic and paper mail.

If the program is interactive, make it output a short notice like this when it starts in an interactive mode:

Gnomovision version 69, Copyright (C) 19yy name of author Gnomovision comes with ABSOLUTELY NO WARRANTY; for details type `show w'. This is free software, and you are welcome to redistribute it under certain conditions; type `show c' for details.

The hypothetical commands `show w' and `show c' should show the appropriate parts of the General Public License. Of course, the commands you use may be called something other than `show w' and `show c'; they could even be mouse-clicks or menu items--whatever suits your program. You should also get your employer (if you work as a programmer) or your school, if any, to sign a "copyright disclaimer" for the program, if necessary. Here is a sample; alter the names:

Yoyodyne, Inc., hereby disclaims all copyright interest in the program

`Gnomovision' (which makes passes at compilers) written by James Hacker.

<signature of Ty Coon>, 1 April 1989

Ty Coon, President of Vice

This General Public License does not permit incorporating your program into proprietary programs. If your program is a subroutine library, you may consider it more useful to permit linking proprietary appli-cations with the library. If this is what you want to do, use the GNU Library General Public License instead of this License.

Fonts are (c) Bitstream (see below). DejaVu changes are in public domain. Glyphs imported from Arev fonts are (c) Tavmjong Bah (see below)

9.4.14 Bitstream Vera Fonts Copyright

Copyright (c) 2003 by Bitstream, Inc. All Rights Reserved. Bitstream Vera is a trademark of Bitstream, Inc.

Permission is hereby granted, free of charge, to any person obtaining a copy of the fonts accompanying this license ("Fonts") and associated documentation files (the "Font Software"), to reproduce and distrib-ute the Font Software, including without limitation the rights to use, copy, merge, publish, distribute, and/or sell copies of the Font Software, and to permit persons to whom the Font Software is furnished to do so, subject to the following conditions:

The above copyright and trademark notices and this permission notice shall be included in all copies of one or more of the Font Software typefaces.

The Font Software may be modified, altered, or added to, and in particular the designs of glyphs or char-acters in the Fonts may be modified and additional glyphs or characters may be added to the Fonts, only if the fonts are renamed to names not containing either the words "Bitstream" or the word "Vera".

This License becomes null and void to the extent applicable to Fonts or Font Software that has been mod-ified and is distributed under the "Bitstream Vera" names.

The Font Software may be sold as part of a larger software package but no copy of one or more of the Font Software typefaces may be sold by itself.

THE FONT SOFTWARE IS PROVIDED "AS IS", WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT OF COPYRIGHT, PATENT, TRADEMARK, OR OTHER RIGHT. IN NO EVENT SHALL BITSTREAM OR THE GNOME FOUNDATION BE LIABLE FOR ANY CLAIM, DAMAGES OR OTHER LIABILITY, INCLUDING ANY GENERAL, SPECIAL, INDIRECT, INCIDENTAL, OR CONSEQUENTIAL DAM-AGES, WHETHER IN AN ACTION OF CONTRACT, TORT OR OTHERWISE, ARISING FROM, OUT OF THE USE OR INABILITY TO USE THE FONT SOFTWARE OR FROM OTHER DEALINGS IN THE FONT SOFTWARE.

Except as contained in this notice, the names of Gnome, the Gnome Foundation, and Bitstream Inc., shall not be used in advertising or otherwise to promote the sale, use or other dealings in this Font Software without prior written authorization from the Gnome Foundation or Bitstream Inc., respectively. For fur-ther information, contact: fonts at gnome dot org.

9.4.15 Arev Fonts Copyright

Copyright (c) 2006 by Tavmjong Bah. All Rights Reserved.

Permission is hereby granted, free of charge, to any person obtaining a copy of the fonts accompanying this license ("Fonts") and associated documentation files (the "Font Software"), to reproduce and distrib-ute the modifications to the Bitstream Vera Font Software, including without limitation the rights to use, copy, merge, publish, distribute, and/or sell copies of the Font Software, and to permit persons to whom the Font Software is furnished to do so, subject to the following conditions:

The above copyright and trademark notices and this permission notice shall be included in all copies of one or more of the Font Software typefaces.

The Font Software may be modified, altered, or added to, and in particular the designs of glyphs or char-acters in the Fonts may be modified and additional glyphs or characters may be added to the Fonts, only if the fonts are renamed to names not containing either the words "Tavmjong Bah" or the word "Arev".

This License becomes null and void to the extent applicable to Fonts or Font Software that has been mod-ified and is distributed under the "Tavmjong Bah Arev" names.

The Font Software may be sold as part of a larger software package but no copy of one or more of the Font Software typefaces may be sold by itself.

THE FONT SOFTWARE IS PROVIDED "AS IS", WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT OF COPYRIGHT, PATENT, TRADEMARK, OR OTHER RIGHT. IN NO EVENT SHALL TAVMJONG BAH BE LIABLE FOR ANY CLAIM, DAMAGES OR OTHER LIABILITY, INCLUDING ANY GENERAL, SPECIAL, INDIRECT, INCIDENTAL, OR CONSEQUENTIAL DAMAGES, WHETHER IN AN ACTION OF CONTRACT, TORT OR OTHERWISE, ARISING FROM, OUT OF THE USE OR INABILITY TO USE THE FONT SOFTWARE OR FROM OTHER DEALINGS IN THE FONT SOFTWARE.

Except as contained in this notice, the name of Tavmjong Bah shall not be used in advertising or other-wise to promote the sale, use or other dealings in this Font Software without prior written authorization from Tavmjong Bah. For further information, contact: tavmjong @ free. fr.

$Id: LICENSE 2133 2007-11-28 02:46:28Z lechimp $

9.4.16 Java Util Concurrent Public Domain

ICU (license below)

ICU License - ICU 1.8.1 and later

COPYRIGHT AND PERMISSION NOTICE

Copyright (c) 1995-2011 International Business Machines Corporation and others

All rights reserved.

Permission is hereby granted, free of charge, to any person obtaining a copy of this software and associ-ated documentation files (the "Software"), to deal in the Software without restriction, including without limitation the rights to use, copy, modify, merge, publish, distribute, and/or sell copies of the Software, and to permit persons to whom the Software is furnished to do so, provided that the above copyright notice(s) and this permission notice appear in all copies of the Software and that both the above copyright notice(s) and this permission notice appear in supporting documentation.

THE SOFTWARE IS PROVIDED "AS IS", WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, INCLUD-ING BUT NOT LIMITED TO THE WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PUR-POSE AND NONINFRINGEMENT OF THIRD PARTY RIGHTS. IN NO EVENT SHALL THE COPYRIGHT HOLDER OR HOLDERS INCLUDED IN THIS NOTICE BE LIABLE FOR ANY CLAIM, OR ANY SPECIAL INDIRECT OR CONSE-QUENTIAL DAMAGES, OR ANY DAMAGES WHATSOEVER RESULTING FROM LOSS OF USE, DATA OR PROF-ITS, WHETHER IN AN ACTION OF CONTRACT, NEGLIGENCE OR OTHER TORTIOUS ACTION, ARISING OUT OF OR IN CONNECTION WITH THE USE OR PERFORMANCE OF THIS SOFTWARE.

Except as contained in this notice, the name of a copyright holder shall not be used in advertising or other-wise to promote the sale, use or other dealings in this Software without prior written authorization of the copyright holder.

All trademarks and registered trademarks mentioned herein are the property of their respective owners.

iText	LGPL
JavaAssist	LGPL
JSoup	Apache 2.0
JMDNS	Apache 2.0
Password Complexity Validator	(Fork of Apache 2.0 Product)
Snowtide, Oracle JDBC, NetApp, Centera	(Commercial License)
Apache POI	Apache 2.0
PostGres JDBC	BSD
Aperture	(WP Extractor) (Special License, below)

Aperture License (for releases from 1.3.0 onward)

The Aperture Framework itself is distributed under a BSD-style license. The copyright holder for each par-ticular source file is given within the file itself.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

Redistributions of source code must retain the above copyright notice, this list of conditions and the fol-lowing disclaimer.

Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution

Neither the name of the copyright holder nor the names of its contributors may be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDERS AND CONTRIBUTORS "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE COPYRIGHT OWNER OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUB-STITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

Third party libs

Licenses for third party libraries used in Aperture can be found in ./lib/

The use of BSD in the core framework allows developers to license implementations of Aperture APIs under any license they see fit, including proprietary and commercial licenses.

Rationale

The aperture framework is free to use and to extend. Through use of BSD, we experience that the frame-work is more useful and users are willing to contribute to the project.

9.4.17 SLF4J MIT License, below

Copyright (c) 2004-2008 QOS.ch All rights reserved. Permission is hereby granted, free of charge, to any person obtaining a copy of this software and associated documentation files (the "Software"), to deal in the Software without restriction, including without limitation the rights to use, copy, modify, merge, pub-lish, distribute, sublicense, and/or sell copies of the Software, and to permit persons to whom the Soft-ware is furnished to do so, subject to the following conditions: The above copyright notice and this permission notice shall be included in all copies or substantial portions of the Software. THE SOFTWARE IS PROVIDED "AS IS", WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIM-ITED TO THE WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONIN-FRINGEMENT. IN NO EVENT SHALL THE AUTHORS OR COPYRIGHT HOLDERS BE LIABLE FOR ANY CLAIM, DAMAGES OR OTHER LIABILITY, WHETHER IN AN ACTION OF CONTRACT, TORT OR OTHERWISE, ARISING FROM, OUT OF OR IN CONNECTION WITH THE SOFTWARE OR THE USE OR OTHER DEALINGS IN THE SOFT-WARE.

Spring Framework Apache 2.0

MS SQL JDBC (Commercial License)

Metro JAXWS (CDDL License, below)

COMMON DEVELOPMENT AND DISTRIBUTION LICENSE (CDDL)Version 1.1

1. Definitions.

1.1. “Contributor” means each individual or entity that creates or contributes to the creation of Modifications.

1.2. “Contributor Version” means the combination of the Original Software, prior Modifications used by a Contributor (if any), and the Modifications made by that particular Contributor.

1.3. “Covered Software” means (a) the Original Software, or (b) Modifications, or (c) the combination of files containing Original Software with files containing Modifications, in each case including portions thereof.

1.4. “Executable” means the Covered Software in any form other than Source Code.

1.5. “Initial Developer” means the individual or entity that first makes Original Software available under this License.

1.6. “Larger Work” means a work which combines Covered Software or portions thereof with code not governed by the terms of this License.

1.7. “License” means this document.

1.8. “Licensable” means having the right to grant, to the maximum extent possible, whether at the time of the initial grant or subsequently acquired, any and all of the rights conveyed herein.

1.9. “Modifications” means the Source Code and Executable form of any of the following:

A. Any file that results from an addition to, deletion from or modification of the contents of a file containing Original Software or previous Modifications;

B. Any new file that contains any part of the Original Software or previous Modification; or

C. Any new file that is contributed or otherwise made available under the terms of this License.

1.10. “Original Software” means the Source Code and Executable form of computer software code that is originally released under this License.

1.11. “Patent Claims” means any patent claim(s), now owned or hereafter acquired, including without limitation, method, process, and apparatus claims, in any patent Licensable by grantor.

1.12. “Source Code” means (a) the common form of computer software code in which modifications are made and (b) associated documentation included in or with such code.

1.13. “You” (or “Your”) means an individual or a legal entity exercising rights under, and complying with all of the terms of, this License. For legal entities, “You” includes any entity which controls, is controlled by, or is under common control with You. For purposes of this definition, “control” means (a) the power, direct or indirect, to cause the direction or management of such entity, whether by contract or otherwise, or (b) ownership of more than fifty percent (50%) of the outstanding shares or beneficial ownership of such entity.

2. License Grants.

2.1. The Initial Developer Grant.

Conditioned upon Your compliance with Section 3.1 below and subject to third party intellectual property claims, the Initial Developer hereby grants You a world-wide, royalty-free, non-exclusive license:

(a) under intellectual property rights (other than patent or trademark) Licensable by Initial Developer, to use, reproduce, modify, display, perform, sublicense and distribute the Original Software (or portions thereof), with or without Modifications, and/or as part of a Larger Work; and

(b) under Patent Claims infringed by the making, using or selling of Original Software, to make, have made, use, practice, sell, and offer for sale, and/or otherwise dispose of the Original Software (or portions thereof).

(c) The licenses granted in Sections 2.1(a) and (b) are effective on the date Initial Developer first distrib-utes or otherwise makes the Original Software available to a third party under the terms of this License.

(d) Notwithstanding Section 2.1(b) above, no patent license is granted: (1) for code that You delete from the Original Software, or (2) for infringements caused by: (i) the modification of the Original Software, or (ii) the combination of the Original Software with other software or devices.

2.2. Contributor Grant.

Conditioned upon Your compliance with Section 3.1 below and subject to third party intellectual property claims, each Contributor hereby grants You a world-wide, royalty-free, non-exclusive license:

(a) under intellectual property rights (other than patent or trademark) Licensable by Contributor to use, reproduce, modify, display, perform, sublicense and distribute the Modifications created by such Contrib-utor (or portions thereof), either on an unmodified basis, with other Modifications, as Covered Software and/or as part of a Larger Work; and

(b) under Patent Claims infringed by the making, using, or selling of Modifications made by that Contrib-utor either alone and/or in combination with its Contributor Version (or portions of such combination), to make, use, sell, offer for sale, have made, and/or otherwise dispose of: (1) Modifications made by that Contributor (or portions thereof); and (2) the combination of Modifications made by that Contributor with its Contributor Version (or portions of such combination).

(c) The licenses granted in Sections 2.2(a) and 2.2(b) are effective on the date Contributor first distrib-utes or otherwise makes the Modifications available to a third party.

(d) Notwithstanding Section 2.2(b) above, no patent license is granted: (1) for any code that Contributor has deleted from the Contributor Version; (2) for infringements caused by: (i) third party modifications of Contributor Version, or (ii) the combination of Modifications made by that Contributor with other soft-ware (except as part of the Contributor Version) or other devices; or (3) under Patent Claims infringed by Covered Software in the absence of Modifications made by that Contributor.

3. Distribution Obligations.

3.1. Availability of Source Code.

Any Covered Software that You distribute or otherwise make available in Executable form must also be made available in Source Code form and that Source Code form must be distributed only under the terms of this License. You must include a copy of this License with every copy of the Source Code form of the Cov-ered Software You distribute or otherwise make available. You must inform recipients of any such Covered Software in Executable form as to how they can obtain such Covered Software in Source Code form in a reasonable manner on or through a medium customarily used for software exchange.

3.2. Modifications.

The Modifications that You create or to which You contribute are governed by the terms of this License. You represent that You believe Your Modifications are Your original creation(s) and/or You have sufficient rights to grant the rights conveyed by this License.

3.3. Required Notices.

You must include a notice in each of Your Modifications that identifies You as the Contributor of the Modification. You may not remove or alter any copyright, patent or trademark notices contained within the Covered Software, or any notices of licensing or any descriptive text giving attribution to any Contrib-utor or the Initial Developer.

3.4. Application of Additional Terms.

You may not offer or impose any terms on any Covered Software in Source Code form that alters or restricts the applicable version of this License or the recipients' rights hereunder. You may choose to offer, and to charge a fee for, warranty, support, indemnity or liability obligations to one or more recipients of Covered Software. However, you may do so only on Your own behalf, and not on behalf of the Initial Developer or any Contributor. You must make it absolutely clear that any such warranty, support, indem-nity or liability obligation is offered by You alone, and You hereby agree to indemnify the Initial Developer and every Contributor for any liability incurred by the Initial Developer or such Contributor as a result of warranty, support, indemnity or liability terms You offer.

3.5. Distribution of Executable Versions.

You may distribute the Executable form of the Covered Software under the terms of this License or under the terms of a license of Your choice, which may contain terms different from this License, provided that You are in compliance with the terms of this License and that the license for the Executable form does not attempt to limit or alter the recipient's rights in the Source Code form from the rights set forth in this License. If You distribute the Covered Software in Executable form under a different license, You must make it absolutely clear that any terms which differ from this License are offered by You alone, not by the Initial Developer or Contributor. You hereby agree to indemnify the Initial Developer and every Contribu-tor for any liability incurred by the Initial Developer or such Contributor as a result of any such terms You offer.

3.6. Larger Works.

You may create a Larger Work by combining Covered Software with other code not governed by the terms of this License and distribute the Larger Work as a single product. In such a case, You must make sure the requirements of this License are fulfilled for the Covered Software.

4. Versions of the License.

4.1. New Versions.

Oracle is the initial license steward and may publish revised and/or new versions of this License from time to time. Each version will be given a distinguishing version number. Except as provided in Section 4.3, no one other than the license steward has the right to modify this License.

4.2. Effect of New Versions.

You may always continue to use, distribute or otherwise make the Covered Software available under the terms of the version of the License under which You originally received the Covered Software. If the Initial Developer includes a notice in the Original Software prohibiting it from being distributed or otherwise made available under any subsequent version of the License, You must distribute and make the Covered Software available under the terms of the version of the License under which You originally received the Covered Software. Otherwise, You may also choose to use, distribute or otherwise make the Covered Soft-ware available under the terms of any subsequent version of the License published by the license steward.

4.3. Modified Versions.

When You are an Initial Developer and You want to create a new license for Your Original Software, You may create and use a modified version of this License if You: (a) rename the license and remove any refer-ences to the name of the license steward (except to note that the license differs from this License); and (b) otherwise make it clear that the license contains terms which differ from this License.

5. DISCLAIMER OF WARRANTY.

COVERED SOFTWARE IS PROVIDED UNDER THIS LICENSE ON AN “AS IS” BASIS, WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED OR IMPLIED, INCLUDING, WITHOUT LIMITATION, WARRANTIES THAT THE COVERED SOFTWARE IS FREE OF DEFECTS, MERCHANTABLE, FIT FOR A PARTICULAR PURPOSE OR NON-INFRINGING. THE ENTIRE RISK AS TO THE QUALITY AND PERFORMANCE OF THE COVERED SOFTWARE IS WITH YOU. SHOULD ANY COVERED SOFTWARE PROVE DEFECTIVE IN ANY RESPECT, YOU (NOT THE INITIAL DEVELOPER OR ANY OTHER CONTRIBUTOR) ASSUME THE COST OF ANY NECESSARY SERVICING, REPAIR OR CORRECTION. THIS DISCLAIMER OF WARRANTY CONSTITUTES AN ESSENTIAL PART OF THIS LICENSE. NO USE OF ANY COVERED SOFTWARE IS AUTHORIZED HEREUNDER EXCEPT UNDER THIS DISCLAIMER.

6. TERMINATION.

6.1. This License and the rights granted hereunder will terminate automatically if You fail to comply with terms herein and fail to cure such breach within 30 days of becoming aware of the breach. Provisions which, by their nature, must remain in effect beyond the termination of this License shall survive.

6.2. If You assert a patent infringement claim (excluding declaratory judgment actions) against Initial Developer or a Contributor (the Initial Developer or Contributor against whom You assert such claim is referred to as “Participant”) alleging that the Participant Software (meaning the Contributor Version where the Participant is a Contributor or the Original Software where the Participant is the Initial Devel-oper) directly or indirectly infringes any patent, then any and all rights granted directly or indirectly to You by such Participant, the Initial Developer (if the Initial Developer is not the Participant) and all Contribu-tors under Sections 2.1 and/or 2.2 of this License shall, upon 60 days notice from Participant terminate prospectively and automatically at the expiration of such 60 day notice period, unless if within such 60 day period You withdraw Your claim with respect to the Participant Software against such Participant either unilaterally or pursuant to a written agreement with Participant.

6.3. If You assert a patent infringement claim against Participant alleging that the Participant Software directly or indirectly infringes any patent where such claim is resolved (such as by license or settlement) prior to the initiation of patent infringement litigation, then the reasonable value of the licenses granted by such Participant under Sections 2.1 or 2.2 shall be taken into account in determining the amount or value of any payment or license.

6.4. In the event of termination under Sections 6.1 or 6.2 above, all end user licenses that have been val-idly granted by You or any distributor hereunder prior to termination (excluding licenses granted to You by any distributor) shall survive termination.

7. LIMITATION OF LIABILITY.

UNDER NO CIRCUMSTANCES AND UNDER NO LEGAL THEORY, WHETHER TORT (INCLUDING NEGLI-GENCE), CONTRACT, OR OTHERWISE, SHALL YOU, THE INITIAL DEVELOPER, ANY OTHER CONTRIBUTOR, OR ANY DISTRIBUTOR OF COVERED SOFTWARE, OR ANY SUPPLIER OF ANY OF SUCH PARTIES, BE LIABLE TO ANY PERSON FOR ANY INDIRECT, SPECIAL, INCIDENTAL, OR CONSEQUENTIAL DAMAGES OF ANY CHAR-ACTER INCLUDING, WITHOUT LIMITATION, DAMAGES FOR LOSS OF GOODWILL, WORK STOPPAGE, COM-PUTER FAILURE OR MALFUNCTION, OR ANY AND ALL OTHER COMMERCIAL DAMAGES OR LOSSES, EVEN IF SUCH PARTY SHALL HAVE BEEN INFORMED OF THE POSSIBILITY OF SUCH DAMAGES. THIS LIMITATION OF LIABILITY SHALL NOT APPLY TO LIABILITY FOR DEATH OR PERSONAL INJURY RESULTING FROM SUCH PARTY'S NEGLIGENCE TO THE EXTENT APPLICABLE LAW PROHIBITS SUCH LIMITATION. SOME JURISDIC-TIONS DO NOT ALLOW THE EXCLUSION OR LIMITATION OF INCIDENTAL OR CONSEQUENTIAL DAMAGES, SO THIS EXCLUSION AND LIMITATION MAY NOT APPLY TO YOU.

8. U.S. GOVERNMENT END USERS.

The Covered Software is a “commercial item,” as that term is defined in 48 C.F.R. 2.101 (Oct. 1995), con-sisting of “commercial computer software” (as that term is defined at 48 C.F.R. § 252.227-7014(a)(1)) and “commercial computer software documentation” as such terms are used in 48 C.F.R. 12.212 (Sept. 1995). Consistent with 48 C.F.R. 12.212 and 48 C.F.R. 227.7202-1 through 227.7202-4 (June 1995), all U.S. Gov-ernment End Users acquire Covered Software with only those rights set forth herein. This U.S. Government Rights clause is in lieu of, and supersedes, any other FAR, DFAR, or other clause or provision that addresses Government rights in computer software under this License.

9. MISCELLANEOUS.

This License represents the complete agreement concerning subject matter hereof. If any provision of this License is held to be unenforceable, such provision shall be reformed only to the extent necessary to make it enforceable. This License shall be governed by the law of the jurisdiction specified in a notice con-tained within the Original Software (except to the extent applicable law, if any, provides otherwise), excluding such jurisdiction's conflict-of-law provisions. Any litigation relating to this License shall be sub-ject to the jurisdiction of the courts located in the jurisdiction and venue specified in a notice contained within the Original Software, with the losing party responsible for costs, including, without limitation, court costs and reasonable attorneys' fees and expenses. The application of the United Nations Conven-tion on Contracts for the International Sale of Goods is expressly excluded. Any law or regulation which provides that the language of a contract shall be construed against the drafter shall not apply to this License. You agree that You alone are responsible for compliance with the United States export adminis-tration regulations (and the export control laws and regulation of any other countries) when You use, dis-tribute or otherwise make available any Covered Software.

10. RESPONSIBILITY FOR CLAIMS.

As between Initial Developer and the Contributors, each party is responsible for claims and damages arising, directly or indirectly, out of its utilization of rights under this License and You agree to work with Initial Developer and Contributors to distribute such responsibility on an equitable basis. Nothing herein is intended or shall be deemed to constitute any admission of liability.

NOTICE PURSUANT TO SECTION 9 OF THE COMMON DEVELOPMENT AND DISTRIBUTION LICENSE (CDDL)

The code released under the CDDL shall be governed by the laws of the State of California (excluding con-flict-of-law provisions). Any litigation relating to this License shall be subject to the jurisdiction of the Fed-eral Courts of the Northern District of California and the state courts of the State of California, with venue lying in Santa Clara County, California.