

GWAVA Inc.

GWAVA Retain

Administration and User Guide

Version 3

“GWAVA” is a registered trade mark of GWAVA Inc, 100 Alexis Nihon, Suite 500, Saint Laurent, Quebec H4M 2P1

Retain” is a trade mark of GWAVA Inc, 100 Alexis Nihon, Suite 500, Saint Laurent, Quebec H4M 2P1

GroupWise is a trademark of Novell, Inc.

Exchange and Windows are trademarks of Microsoft Inc.

Facebook and Twitter are registered trademarks of their respective owners.

Contents

Important Notes and New Features	6
Post-Install Tasks.....	7
Initial Setup Overview	7
GWAVA Retain Initial Setup.....	8
<i>About Storage Paths</i>	9
<i>Advanced Storage Options</i>	9
Using GWAVA Retain	13
Logging In.....	13
<i>Status and Updates</i>	14
<i>Saving Changes</i>	15
<i>Notifications</i>	16
<i>Modules</i>	17
<i>Mailbox Mapping Options</i>	18
Exchange Module.....	19
Social Messaging.....	28
GroupWise Module.....	29
<i>Retain Job components:</i>	32
<i>Jobs</i>	32
Schedules	33
Profiles	34
<i>Exchange Profile</i>	34
<i>GroupWise Profile</i>	39
<i>About GroupWise's Smart Purge and Retention Flags</i>	43
<i>Setting up SmartPurge and Retention flags in GroupWise 7</i>	45
<i>Social Messaging Profile</i>	50
Workers.....	51
Jobs	59
<i>Exchange Jobs</i>	59
<i>GroupWise Jobs</i>	64
<i>Social Messaging Jobs</i>	68
Management.....	70
Users	70
<i>Offline Password</i>	70
Creating a new user	71
<i>Core Settings</i>	71
<i>User Rights</i>	72
<i>User-Level Rights Summary</i>	73
<i>User-level rights</i>	73
<i>The Read Configuration right (Redline Integration)</i>	74
<i>Mailboxes</i>	74
Address Book Selector	68
<i>Adding Mailboxes</i>	68
<i>About "Show only recently cached items"</i>	68
<i>Miscellaneous Tab</i>	68
<i>GroupWise Proxy Support</i>	70
Groups.....	71
<i>Creating a New Group</i>	71

<i>Group Rights</i>	72
Deletion Management	75
<i>Core Settings</i>	75
<i>Basic Options</i>	75
<i>Data Scope</i>	76
<i>Job Members</i>	77
<i>Notification</i>	77
<i>Schedule</i>	77
<i>Litigation Hold</i>	78
Configuration	79
Server Configuration	79
<i>Communications</i>	79
<i>Index</i>	80
<i>Storage</i>	82
<i>Accounts</i>	82
<i>Intruder Lockout</i>	83
<i>Maintenance</i>	84
<i>Ignore Domains</i>	85
<i>Logging</i>	86
<i>Auditing Records</i>	86
<i>Miscellaneous</i>	87
<i>Export XML</i>	87
Module Configuration.....	87
Storage Manager / Storage Engine	88
<i>Partitions</i>	90
Stubbing Server	91
<i>ConsoleOne GroupWise settings</i>	91
<i>Retain Stubbing Settings</i>	93
Statistics Server Configuration.....	96
Licensing.....	103
Overview	104
Audit Log	104
System Log	105
Searching the Retain Archives	106
<i>Quick View</i>	107
<i>Using Online Help</i>	107
Browse	108
<i>Language Selection</i>	108
<i>Merged Mailboxes</i>	109
<i>Browsing Tools (Widgets)</i>	109
<i>Browse Window Defaults</i>	109
<i>New Mailbox</i>	110
<i>Toolbar Options</i>	111
<i>Specifying a Date Range</i>	111
<i>Litigation Hold</i>	111
<i>Forwarding a Message</i>	111
<i>PDF Export</i>	112
<i>Restore</i>	115

Search.....	116
<i>Opening/Viewing a Message</i>	117
Search in More Detail	119
<i>Search Window</i>	119
<i>Core Tab:</i>	119
<i>Selecting Mailboxes</i>	120
<i>Loading, Saving and Deleting Searches</i>	120
<i>Scope Tab:</i>	121
Item Type	121
<i>Sort Tab:</i>	121
<i>Misc. Tab:</i>	121
<i>Litigation hold</i>	122
<i>Shared Queries</i>	122
Search Window	125
<i>Core Settings</i>	125
<i>User Rights Tab</i>	126
<i>Mailboxes Tab</i>	126
<i>Miscellaneous Tab</i>	126
Using the Retain Help System	127
About and Diagnostics Page	128
Diagnostic Buttons	129
<i>Mailbox Merger / Splitter</i>	130
<i>GW to Exchange Scope</i>	131
<i>GW Exchange Migrator</i>	131
Appendix A: Backing up Retain	134
Fine tuning GWAVA Retain	135
<i>Tomcat Memory Usage Tuning Overview</i>	135
<i>Two Critical Issues</i>	135
General Recommendations	135
<i>Retain Server Component</i>	135
<i>Retain Worker Component(s)</i>	135
Tomcat Configuration Parameters	136
How to change memory values	137
<i>Checking memory:</i>	137
UPGRADED SYSTEMS	137
Appendix B – Troubleshooting for Common GroupWise Problems	138
EA04	138
D712/D714	138
C05D	138
<i>Other notes regarding GroupWise 7, SP 3 (7.0.3):</i>	138
Appendix C – The Tools Folder	139
GroupWise tools	139
<i>GroupWise Archive Migration Tool</i>	139
<i>Archive Migration Tool</i>	139
<i>The Archive Plug-in</i>	146
<i>Blob Extraction Tool</i>	146
<i>Single Sign On</i>	146
<i>WebAccess</i>	147

<i>Exchange PST Migrator</i>	149
<i>Single Sign On</i>	153
Recover Admin	155
Appendix D – Reload Integration	155
<i>A Brief Review on How Reload Works</i>	156
<i>How Retain Takes Advantage of Reload's features</i>	157
<i>Multiple Workers on One Server</i>	157
<i>Timing</i>	157
Retain Settings	158
<i>How to Setup Reload to work with Retain</i>	159
IMPORTANT Notes for the Integration	166
<i>Retain</i>	166
<i>Reload</i>	167
Appendix E – Migrating Retain to a new Server or Platform	168
Appendix F – Switching the Indexing Engine	170
Appendix G – Exchange Permissions required for Retain	177
Exchange settings	177
<i>Grant Impersonation Permissions to the Retain user</i>	177
<i>Room and Equipment Resources:</i>	178
<i>Authentication Methods</i>	179
Appendix H: Trusted Application Key Creation	180
Trusted Application Key creation through Console One for GroupWise 8	180
GroupWise Trusted Application Key generator	181

Important Notes and New Features

Retain adds support for archiving Social Messaging. With the Retain for Social Messaging, activity in Facebook and Twitter is recorded and archived into Retain for full compliance. Retain for Social Messaging is designed as a virtual appliance, with full support for VMware ESX and ESI 5.

Retain's message store now has improved searching capability, with multiple simultaneous search fields and extended operator support. Lucene also has improved search abilities with a near real-time search support.

Retain 3 also boasts a new enhanced data store, full Office 365 support, and improvements to the deletion manager which improves performance, even while under maintenance lock.

Retain has also added a new reporting and statistics server. The Statistics Server gathers pertinent information on performance and messages and can generate reports on jobs, processes, and resources for the Retain Server and connected components. The reports can either be delivered immediately or through email.

Systems that are upgraded from previous versions of Retain must migrate the database. A database migration wizard is provided. Due to size and time considerations with the migration, a benchmark tool is provided to offer an estimate on how long the migration process will take. In addition, the migration temporarily requires substantial free disk space; up to twice the size of the current Retain database. Please look to the *Upgrading from Previous Versions of Retain* guide for detailed information.

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 regular 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.

Initial Setup Overview

- 1) Choose your storage path.
- 2) Retain prompts you for access information to your SQL Database Server.
 - You must create this database beforehand. (See the MySQL Quick Start section in the install guide, or documentation for the SQL server of your choice.)
 - 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.
- 3) Create an administrator account with the default username, admin. It is the first account to be created.
 - This 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).
 - Additional accounts can be created as desired with as many rights as you wish. 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.
- 4) Retain prompts you for the following SMTP information for notification and forwarding.

DNS host names are generally recommended instead of IP Addresses in almost all cases due to ease in central management and repair if a server moves or is replaced.

 - SMTP 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.
- 5) Select the indexing engine. There are two different indexing engines available for use with Retain: Lucene and Exalead. Exalead is to be used with larger systems and requires a separate dedicated server. Lucene is a high performance embedded indexer, which requires fewer servers, but is more memory intensive and does not offer high availability features. Please see the install guides and system requirements for more information on which is appropriate and should be used. Exalead is priced separately; see your appropriate sales representative.

After initial setup, configuration of a module and creation of a profile, schedule, worker, and job are required before messages can be archived.

GWAVA Retain Initial Setup

- 1) Open your web browser.
- 2) Go to
http: //<your- Retai nServer- i p>/Retai nServer
- 3) Follow the wizard.

You will see this screen initially:

Retain

File Paths
Message Store Database
Configuration Database
Administrator Account
SMTP Notification
Indexing Engine

Welcome to GWAVA Retain

This wizard will guide you through the initial configuration of Retain.

- Storage Path
- Message Store Database
- Configuration Database
- Set Administrator Password
- SMTP Notification
- Indexing Engine

During each step, details will be provided to help guide you.

File Path

A storage path is required for Retain, under which indexes, attachments, backups is stored. The wizard will attempt to create the directory if it does not exist. If you receive an error, check the rights you granted to your Tomcat application.

Storage Path: /var/opt/beginfinite/retain

[Advanced Settings](#)

If you are planning to use the standard engine or the NetApp engine, this path should point to a storage device with sufficient (significant) disk space. You may use mapped drives, UNC format, etc.

For all engines except the standard engine (including NetApp), you'll need to perform additional configuration steps in the Storage Engine section of the main administrative interface, after completing the wizard.

[Continue](#)

Please pay special attention to the Storage Path and the Advanced Settings (see [next page](#)).

For descriptions of the different engines, see the [Storage Manager](#) section. If you select a storage system that is not local, be sure to visit the [Storage Manager](#) after the initial setup to ensure that all information has been provided.

- 4) 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.
- 5) Click “Continue”.

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 operating 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 situations like these, a SAN might be a good choice for consideration. (Please visit the [Storage Manager](#) section to view options for different SAN systems.) Either way, it's imperative that you take the time to investigate your storage needs BEFORE you begin.

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

- 2 TB limit. Will storage needs ever exceed this?
- 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.

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.

The storage path is where indices and archives are stored. It should be on a highly reliable storage device, with plenty of disk space.

Advanced Settings

Advanced Settings

For more advanced scenarios, some paths may be assigned to one device and some to another.

Derive all file locations from above base path? ☒

Encryption

Compression (always on)

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 Index/Storage tab. If you wish to specify your storage locations in more detail, click on the “Advanced Settings”.

The storage path is where indices and archives are stored. It should be on a highly reliable storage device, with plenty of disk space.

Advanced Settings

Advanced Settings

For more advanced scenarios, some paths may be assigned to one device and some to another.

Derive all file locations from above base path? ☐

Store attachments under this path

Store xml mappings under this path

Store index files under this path

Store Embedded DB under this path

Store backups under this path

Store license under this path

Encryption

Compression (always on)

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 consider Block level encryption options (such as BitLocker) that have escrow keys and are processed directly at the operating system level.

Encryption

Compression

aes

Database Setup

Welcome to GWAVA Retain

This wizard will guide you through the initial configuration of Retain.

- Storage Path
- Message Store Database
- Configuration Database
- Set Administrator Password
- SMTP Notification
- Indexing Engine

During each step, details will be provided to help guide you.

Message Store Database Connection Parameters

To successfully create the tables the database server must be running with the *retain* database created, and permissions for that database granted to a user. Be sure to specify:

- Driver: Type of Database Server. MySQL 5.0/5.1/5.5, Oracle 10 and 11, MS SQL Server 2005/2008, and PostGres 9.0+ are supported. MySQL 4.x will NOT work correctly. If using MySQL 5, make sure you install the latest version - even minor revisions contain significant bug fixes and performance enhancements.
- Database Server: **If the database server is not running on the same machine as Retain, change localhost to the IP address or host name of the machine.** (Preferably a DNS host name.) Including port is legal syntax (e.g.: 192.168.2.31:3306); if not supplied, default ports will be used.
- Database Name: Name of the database you have created for Retain to use.
- Username: Username with full rights to the database. We recommend that you create a specific user with rights only to the "retain" database. Also, remember that many database servers require specific settings or rights to connect locally (same box setup) vs. REMOTELY (multibox setup).
- Password: Password for access.

Warning: If you have not setup the database according to the documentation, irreparable corruption to your data is likely to occur.

Driver:

Database Server:

Database Name:

DB Instance (MSSQL only):

DB Username:

DB Password:

If you are running MySQL, we also need the database server root username/password for the initial setup. This will not be stored.

DB ROOT Username:

DB ROOT Password:

- 6) Next, this screen appears for entering the database connection parameters. This is the same information from the database created during the preparation phase. (See Install guide) The connection information must be supplied for both the storage and Configuration databases, though they can be housed in the same database. Assumed default is shown.

Welcome to GWAVA Retain

This wizard will guide you through the initial configuration of Retain.

- Storage Path
- Message Store Database
- Configuration Database
- Set Administrator Password
- SMTP Notification
- Indexing Engine

During each step, details will be provided to help guide you.

Message Store Database Connection Parameters

To successfully create the tables the database server must be running with the *retain* database created, and permissions for that database granted to a user. Be sure to specify:

- Driver: Type of Database Server. MySQL 5.0/5.1/5.5, Oracle 10 and 11, MS SQL Server 2005/2008, and PostGres 9.0+ are supported. MySQL 4.x will NOT work correctly. If using MySQL 5, make sure you install the latest version - even minor revisions contain significant bug fixes and performance enhancements.
- Database Server: **If the database server is not running on the same machine as Retain, change localhost to the IP address or host name of the machine.** (Preferably a DNS host name.) Including port is legal syntax (e.g.: 192.168.2.31:3306); if not supplied, default ports will be used.
- Database Name: Name of the database you have created for Retain to use.
- Username: Username with full rights to the database. We recommend that you create a specific user with rights only to the "retain" database. Also, remember that many database servers require specific settings or rights to connect locally (same box setup) vs. REMOTELY (multibox setup).
- Password: Password for access.

Warning: If you have not setup the database according to the documentation, irreparable corruption to your data is likely to occur.

Driver:

Database Server:

Database Name:

DB Instance (MSSQL only):

DB Username:


DB Password:

If you are running MySQL, we also need the database server root username/password for the initial setup. This will not be stored.

DB ROOT Username:

DB ROOT Password:


- 7) Now, supply a password for Retain's administrator account (admin). **Do not forget what you choose here!** (Use the [Recover Admin](#) utility if you have lost the administrator password.)



The screenshot shows the 'Welcome to GWAVA Retain' wizard. On the left, a list of steps is shown with progress indicators: File Paths (green check), Message Store Database (green check), Configuration Database (green check), Administrator Account (green check), SMTP Notification (grey circle), Indexing Engine (grey circle), Trusted Application Setup (grey circle), and SOAP Connection (grey circle). The main content area is titled 'Administrator Account' and contains instructions: 'This administrator account will be needed for full access to Retain. You may also create additional administrators once this setup wizard has completed. Keep this password secure and safe, and consider creating a backup administrator account. (If you do lose the account password, a utility is available to create a new account, in the Tools folder of your installation package)'. Below the text are three input fields: 'Login Name' with 'admin', 'Password' with six dots, and 'Confirm Password' with six dots. A 'Continue' button is at the bottom.

- 8) 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 successfully delivered. If you see an error, check firewall settings, (including local desktop/server firewalls such as McAfee, Symantec, and SUSE).

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 of the Administrative rights granted. DO NOT delete this account. If the admin account is deleted, please contact support or refer to the [Recover Admin](#) tool to restore a full administrator account.



The screenshot shows the 'Welcome to GWAVA Retain' wizard at the 'SMTP Notification' step. The left sidebar shows progress: File Paths (green check), Message Store Database (green check), Configuration Database (green check), Administrator Account (green check), SMTP Notification (green check), Indexing Engine (grey circle), Trusted Application Setup (grey circle), and SOAP Connection (grey circle). The main content area is titled 'SMTP Notification' and contains instructions: 'SMTP Notification is used for sending you progress reports, and also is used whenever a user forwards a message from Retain. Below, you configure the defaults. These can be overridden on a per-job basis. Most of these entries are self-explanatory. The TO is used for job notifications, not for forwarding. The SMTP user and password are used for relaying—which you may not need for job notifications, but you will need it if you want to permit forwarding outside your Internet domain.' Below the text are five input fields: 'SMTP Mail Server' with '192.168.1.104', 'SMTP Mail From Address (e.g. retain@yourcompany.com):' with 'retain@gwava.com', 'SMTP To Address:' with 'admin@gwava.com', 'SMTP Username:' with 'chris', and 'SMTP Password:' with four dots. A 'Continue' button is at the bottom.

9) Index Engine Selection

The screenshot shows the 'Welcome to GWAVA Retain' wizard. On the left, a list of configuration steps is shown with green checkmarks for 'File Paths', 'Message Store Database', 'Configuration Database', 'Administrator Account', and 'SMTP Notification', and a grey circle for 'Indexing Engine'. The main content area is titled 'Welcome to GWAVA Retain' and explains that the wizard will guide through the initial configuration. It lists the steps: Storage Path, Message Store Database, Configuration Database, Set Administrator Password, SMTP Notification, and Indexing Engine. Below this, it states 'During each step, details will be provided to help guide you.' The 'Step 5: Indexing Engine' section explains the importance of the choice and provides a summary of the two engines: Lucene (fast, easy, no additional software) and Exalead (powerful, scalable, requires hardware and license). At the bottom, there is a dropdown menu for 'Indexing Engine' with 'Choose one of the listed engines' selected, and a 'Continue' button.

- Lucene is the default indexing engine, and is included in Retain. It is well suited for smaller installations and requires no further configuration or management.
- Exalead is a more robust and powerful search engine, providing high availability and clustering features, as well as more advanced search operations. It requires additional hardware, as well as an additional license, which you may obtain from your sales representative. **Exalead should not be installed on the same server where Retain server is installed.**

If you wish to use Exalead at this point, it **MUST** be installed and configured before this step is completed. If selecting Exalead, provide the IP address or DNS name of the Exalead server and the active base port.


This screenshot shows the same 'Welcome to GWAVA Retain' wizard, but now all configuration steps on the left have green checkmarks, including 'Indexing Engine'. The main content area is the same, but the 'Continue' button has been replaced by a 'Complete' button at the bottom.

Initial setup is now complete. Select the 'Complete' button to be presented with the login prompt.

Using GWAVA Retain

Logging In

- 1) From your web browser, type:
http://<YourServerIP>/RetainServer
- 2) You'll be brought to the login screen:



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

Server

Index Migration in progress	False
SQL Server	OK
Server	RetainServer (UTLW\INCKLSITTPSJZFOWNSQG)
Stored EMails	0
Storage Engine	datastore_process
Server started at:	23-Jul-2012 11:10:41
Archived since server started (total/dupes/new) :	0 / 0 / 0
Server errors	0

Installation time	23-Jul-2012 11:21:35
Enable Index Optimization	Yes
Last Index Optimization	Never

Total documents ever indexed	0
Total emails in memory queue awaiting indexing	0
Total documents indexed so far since uptime	0
Total emails indexed so far since uptime	0
Total emails with indexing errors since uptime	0
Total emails with items skipped (*) since uptime	0
Last item indexed at	Never
Time to process last batch (ms)	0
Indexing thread alive	True

(*) Items may be skipped in indexing for perfectly benign reasons, such as there being no handler for that specific type of file or the file being too large according to your settings in Server Configuration. Also, the ratio of documents to emails is indeterminate, but usually ranges from 1-2

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 on login. Other users will be forwarded to the [Search Interface](#) (covered later).

Status and Updates

The status page has four different tabs which show the status of your system. The status page is set to display 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.

The Server tab shows the basic status of the Server with the messages archived and maintenance information. The Workers tab shows the workers connected, and the status of each.

Server			
Index Migration in progress	False		
SQL Server	OK		
Server	RetainServer (GRZFZEOEQRNSETOAJSZYDVJY)		
Stored EMails	0		
Storage Engine	standard		
Server started at:	15-Mar-2011 19:51:32		
Archived since server started (total/dupes/new):	0 / 0 / 0		
Server errors	0		
Installation time	12-Jan-2011 09:02:14		
Enable Index Optimization	Yes		
Last Index Optimization	11-Mar-2011 01:02:00		
Total documents ever indexed	45363		
Total emails in memory queue awaiting indexing	0		
Total documents indexed so far since uptime	0		
Total emails indexed so far since uptime	0		
Total emails with indexing errors since uptime	0		

Workers			
Worker	migrator		
Last contact	Never		
Last Job:	Never		
Last job start:	Never		
Last job end:	Never		
Registered:	Never		
Total messages stored:	0		
Total mailboxes scanned:	0		
Worker	new worker		
Last contact	15-Mar-2011 22:21:26		
Last Job:	new exchange		
Last job start:	15-Mar-2011 22:21:27		
Last job end:	15-Mar-2011 22:21:30		
Registered:	26-Jan-2011 05:07:25		
Total messages stored:	18590		
Total mailboxes scanned:	7		


The Jobs page shows the jobs that are assigned to the server and their status. You also can select the “show only active jobs” checkbox, and only jobs that are currently running will be displayed.

Jobs			
<input type="checkbox"/> Show only active jobs			
Job	new exchange		
Job running	No		
Job began	15-Mar-2011 22:21:27		
Job ended	15-Mar-2011 22:21:30		
Time elapsed	0d : 0h : 0m		
Next job time	Find		
Last PO / Mailbox scanned	-- / --		
Total errors	0		
Total mailboxes (job)	2		
Total messages (job)	0		
Total messages (lifetime)	0		
Job	new job		
Job running	No		
Job began	03-Mar-2011 20:45:00		

Updates			
You have the latest build.			
Announcements			
The following releases are the newest for your system:			
Current Release >>>			

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.

If you are using Exalead as your indexing engine, you will see a 'Trigger Exalead Indexing' button on the 'Server' tab of the Status and Updates page. Selecting this button will force an immediate indexing run by Exalead.




Jobs	Workers	Server	Updates
Installation time 03-Jun-2010 10:26:22			
Enable Address Book Caching Yes			
Last Address Book Cache 03-Jun-2010 11:30:45			
Enable Index Optimization Yes			
Last Index Optimization Never			
Total documents ever indexed 8423			
Exalead internal queue 8423 Trigger Exalead Indexing			
Total emails in memory queue awaiting indexing 0			
Total documents indexed so far since uptime 8423			
Total emails indexed so far since uptime 1821			
Total emails with indexing errors since uptime 0			
Total emails with items skipped (*) since uptime 772			
Last item indexed at 03-Jun-2010 14:10:29			
Time to process last batch (ms) 660			

After many messages are sent to the Exalead server for indexing, these items are queued. In some cases, there may be a noticeable delay in the processing of that queue, as other processes (such as searching) take high priority. To force an immediate indexing run, press this button.

Warning: Repeatedly using this button can overwhelm your Exalead server, and should NOT be necessary in normal operations.

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 interface, 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:

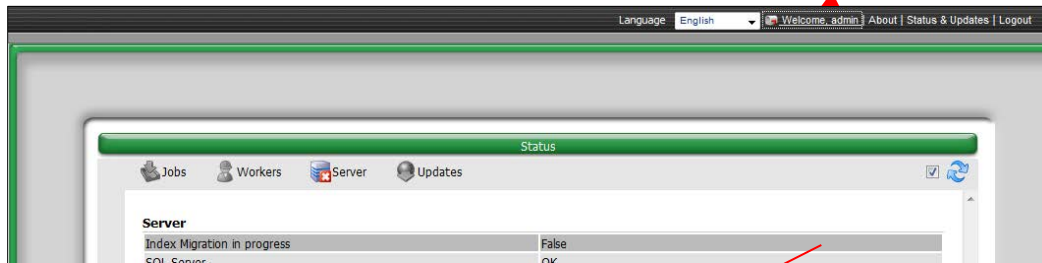


When the "Save Changes" lights up, click on the disk icon to save any changes you have made on this screen otherwise your changes will be lost.


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 servers activation.

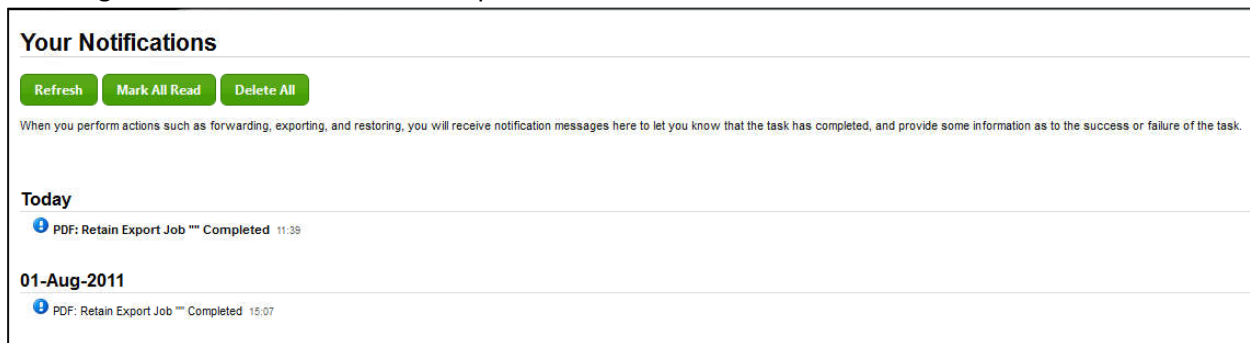
To spawn the notifications page, click on the 'Welcome' link at the top of the page.



When notifications are pending, the red flag on the mailbox is lifted.

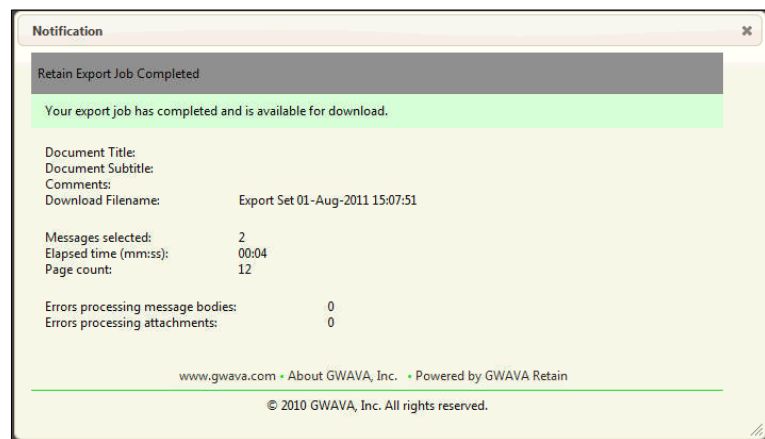


Selecting the mailbox or welcome link opens a new notifications window.



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.

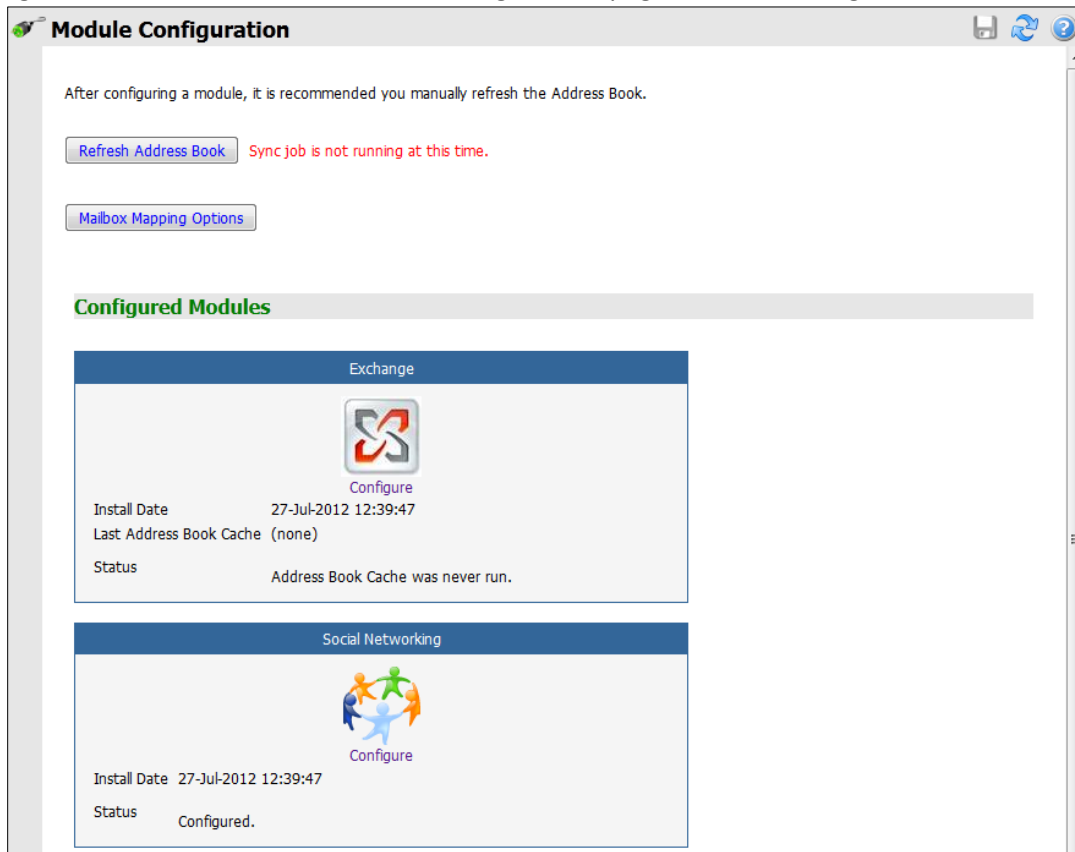


Modules

Retain integrates and works with several different message systems. Data collection will not be available until one or more modules are configured for the Retain system.

Note: Depending on what license you have purchased for your Retain system, different modules may require an additional license. Check the licensing [link it] page to see what licenses you own. An Evaluation license, issued upon initial install, lasts 30 days and grants licenses for all modules

To configure a module, select the 'Module Configuration' page from the Configuration menu.



Each module link will open a new configuration window or tab. If the modules are not configured correctly, messages will not be collected. Profiles and Jobs are separated into module specific sections and no jobs can be configured without first configuring a module. Select the appropriate module(s) for your system by clicking on the module icon.

Certain settings or actions for both Exchange and GroupWise message systems must be configured before Retain can communicate with them. Make sure to complete or check the following changes and settings for each system before continuing. (Links lead to sections in the appendix.)

Mailbox Mapping Options

Mailbox mapping automates the merging of users' mailboxes that have both GroupWise and Exchange accounts, combining the GroupWise and Exchange archives under the same user in Retain. Retain merges the internal UID'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, "un-merging" must be performed manually via the [Mailbox Merger / Splitter tool](#) found on the 'bug' or 'about' screen.** If a refresh of the mailboxes is not performed before users attempt to login or other activity 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.

Mailbox Mapping Options

Mailbox Mapping Options

When running more than one email system, it is often desirable to automatically "map" two mailboxes from different email systems together. Effectively, these two separate mailboxes become one virtual mailbox. Switching between the two is a simple matter of the user choosing the mailbox from the current browse view. Also, rights to mailboxes become shared between the two, so the user need only log on as one mailbox to have access to them all.

To achieve this, during address book caching, the administrator needs to have enabled at least one mailbox mapping option. This **must** be enabled before the newly added email system is cached for the first time.

You may enable multiple address book mappings if you wish.

It is usually sufficient to enable mapping only for address book caching and only as a one-time procedure, unless both email systems will remain active. (For the latter case, enabling for all processes, and for constant activity, is advised)

Mapping Enabled
Never

Last time mappings were used
never
Reset

If you enable the XML File mapping option, you must provide the full path to the XML file, including the filename.

XML File Path

Active Mappings

Match the E-Mail Addresses
Add

Custom
Description
Class

To map one mail system to another, both message systems, Exchange and GroupWise, must be functioning and have the desired accounts created on the system which allows for identification of the different users.

Accounts may be matched via email address, display names, or specified by an administrator supplied XML file. The Address book **MUST** be refreshed immediately in order to merge the users. The merge will not complete 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.

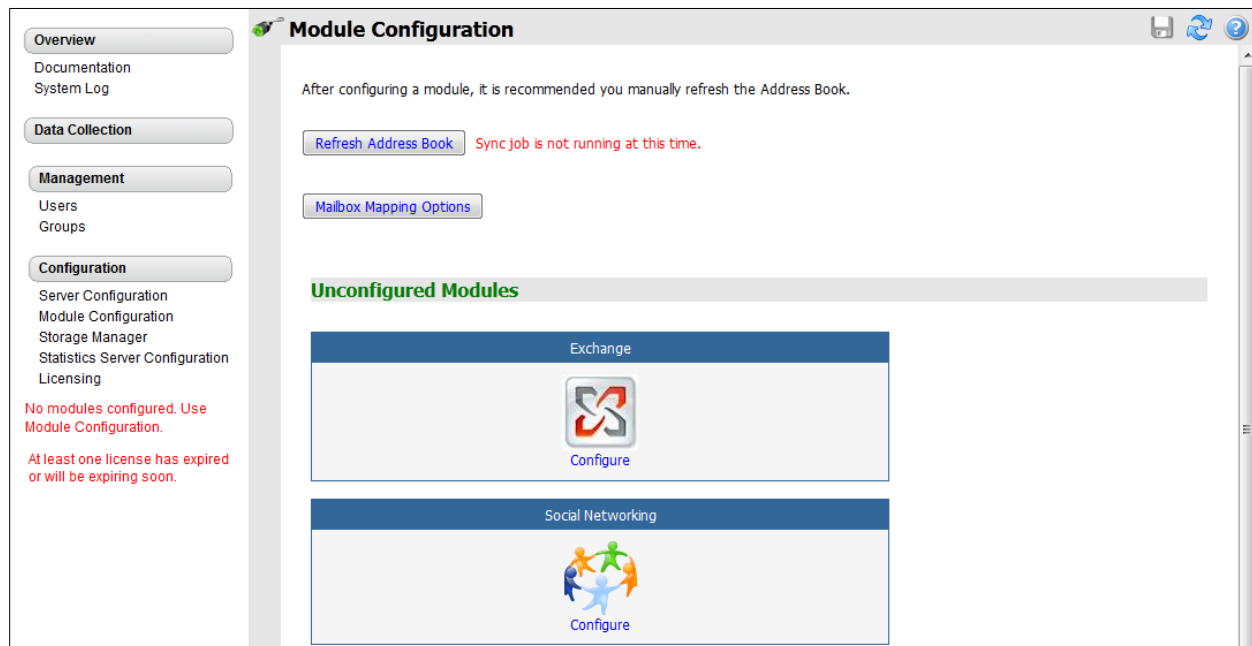
Exchange Module

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)
- Office 365

Retain does NOT support multiple linked Exchange Forests. Ensure that the [Exchange settings](#) have been configured correctly before continuing the Exchange module setup.

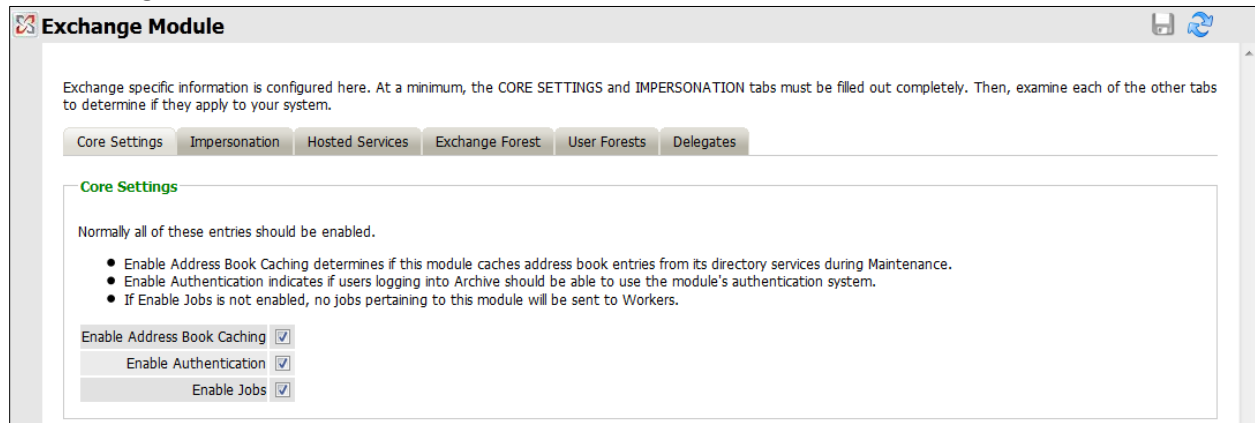
The Exchange module must be configured in the Retain Server before any communication between Retain and an existing Exchangemessagesystem can occur. Open the Retain management page on the Retain Server, and select Module Configuration.



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.

Core Settings



Normally all of 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 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 system 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 configured for this module will be processed.

Impersonation

If the Impersonation and Core Settings tabs are not completely configured with the correct information, the hosted system will not be archived correctly.




The screenshot shows the 'Exchange Module' window with the 'Impersonation' tab selected. The window title is 'Exchange Module'. Below the title bar, there is a message: 'Exchange specific information is configured here. At a minimum, the CORE SETTINGS and IMPERSONATION tabs must be filled out completely. Then, examine each of the other tabs to determine if they apply to your system.' Below this message is a tab bar with six tabs: 'Core Settings', 'Impersonation', 'Hosted Services', 'Exchange Forest', 'User Forests', and 'Delegates'. The 'Impersonation' tab is active. The main content area of the 'Impersonation' tab has a green header 'Impersonation'. Below the header, it says: 'You will need to provide this information regardless of whether you are archiving a hosted or on premise Exchange system.' and 'The impersonation user MUST fulfill the following requirements:'. There are three bullet points: 'Is mail enabled.', 'Has impersonation rights granted to all Exchange servers.', and 'Delegation rights must be granted to any Room or Equipment accounts that you wish to archive.' Below the bullet points are two input fields: 'Global Catalog User' with the value 'user@fqdn.com' and 'Global Catalog Password' which is empty. To the right of the password field is the text '(Password is not set)'.

Hosted Services

If the Exchange system is a hosted service system, Retain must be configured to support the hosted system setup.

Setting up the system is no more taxing than setting up a normal system, though a hosted system requires that the Impersonation and Core Settings tabs are completely and correctly filled out. Other tabs are only applicable depending on the settings of the hosted system, configure as appropriate.



The screenshot shows the 'Exchange Module' window with the 'Hosted Services' tab selected. The window title is 'Exchange Module'. Below the title bar, there is a message: 'Exchange specific information is configured here. At a minimum, the CORE SETTINGS and IMPERSONATION tabs must be filled out completely. Then, examine each of the other tabs to determine if they apply to your system.' Below this message is a tab bar with six tabs: 'Core Settings', 'Impersonation', 'Hosted Services', 'Exchange Forest', 'User Forests', and 'Delegates'. The 'Hosted Services' tab is active. The main content area of the 'Hosted Services' tab has a green header 'Hosted Services'. Below the header, it says: 'Archive supports hosted Exchange services such as Live@Edu and Office365. You may ignore this tab if you are not using such services.' There are three bullet points: 'These services generally do NOT provide LDAP access to their user store, therefore you must maintain a list of user accounts manually, using a Powershell script (see the user manual before proceeding).', 'If your service DOES provide LDAP access, uncheck the checkbox on this tab, and configure the items in the Exchange Forest and User Forest tabs.', and 'If your service DOES NOT provide LDAP access, check the checkbox on this tab. Do NOT configure the Exchange Forest, User Forest, and Delegates tabs - they will be ignored.' Below the bullet points is a checkbox labeled 'I am using a hosted Exchange system' which is currently unchecked.

Mark the checkbox to enable Retain support for hosted Exchange services. No further configuration is required on this page.

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.

Installation

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

1. Install PowerShell 2.0

The appropriate version for the Windows system can be found here:

<http://www.microsoft.com/powershell>

Windows 7 and 2008 R2 already come with PowerShell 2.0
(Even if the version says 1.0, really it's 2.0)

2. Install the Office 365 PowerShell cmdlets

The cmdlets can be found here:

<http://onlinehelp.microsoft.com/en-us/office365-enterprises/hh124998.aspx>

Two setup files are needed from that page:

- Microsoft Online Services Sign-In Assistant
- Microsoft Online Services Module for Windows PowerShell

Execute them both in above order.

3. 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 GWAVA must be run in PowerShell while **running under Administrator rights**.

Set-ExecutionPolicy RemoteSigned

4. Adapt and execute the “sync365.ps1” script provided by GWAVA

Select the “sync365.ps1” script provided by GWAVA in Tools/Exchange/Office 365 folder.

Right click on the script and select “modify”.

This will open the script in the Microsoft Integrated Scripting Environment (ISE) editor.

At the top you can 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.

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

Example:

```
$ExportBasePath="C:\\Temp"
```

Note: If it is unacceptable to set a password in clear text, please view the following solution: <http://bsonposh.com/archives/338> -presents an alternative, which loads and saves encrypted credentials from file.

Execute the script by clicking the play button. This process can take a while for many users. Once the script has completed, a message indicating that the script is done should be displayed along the bottom status bar.

5. Automate it: to run automatically with Task Scheduler at least once a day

The method of automation is up to the administrator. However, the script needs to be run at least once a day, and, due to step 6, the files need to be copied to the Retain Server.

6. Get the files over to Retain Server

In Step 4 the ExportBasePath was set and should contain the two csv files. These files must be automatically copied to the RetainServer/WEB-INF/cfg directory once the script completes. That might require setting up a mapped network drive, or the acceptable solution for the local system. While the ExportBasePath may be mapped directly to the Retain Server, it is recommended to have the csv files copied afterwards.

Retain needs to know login information and existing domains before any archiving can be accomplished. Open the “Exchange Forest” tab and enter the login information.

The screenshot shows the 'Exchange Module' configuration window. At the top, there's a title bar with 'Exchange Module' and a 'Save Changes' button. Below the title bar, a message states: 'Exchange specific information is configured here. At a minimum, the CORE SETTINGS and IMPERSONATION tabs must be filled out completely. Then, examine each of the other tabs to determine if they apply to your system.' A tabbed interface is shown with tabs for 'Core Settings', 'Impersonation', 'Hosted Services', 'Exchange Forest' (selected), 'User Forests', and 'Delegates'. The 'Exchange Forest' tab contains the following content:

Exchange Forest

Ignore this tab if you are running a hosted system.

You MUST fill out all of the entries on this tab if you are running an on-premise Exchange system. Hosted services do not normally use this information.

Typically, you provide the connectivity information for a global catalog, and a user that has full access to Active Directory.

Here are the specific requirements for this user:

- The user account must fulfill all the requirements specified on the Impersonation tab. In fact you enter the username and password on that tab, not here.
- Has read-only access to all parts of Active Directory involving the Exchange System and users.

You also provide a list of Active Directory DNS to search for users and groups.

Global Catalog Host: 192.168.1.120

Global Catalog Port: 3268

Global Catalog Security: Plain Text

+ Add Search Base (e.g. dc=users,dc=company,dc=com)

✗ dc=b, dc=company, dc=com

Test Connection

Resource Forest

Select this checkbox if Exchange is running in a dedicated resource forest. You must then fill out information on the User Forests tab. If this checkbox is not selected, information on the User Forests tab is ignored.

Enable Resource Forest: ☐

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 settings](#) section in the appendix.

On the Exchange Forest tab, you configure all of 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.

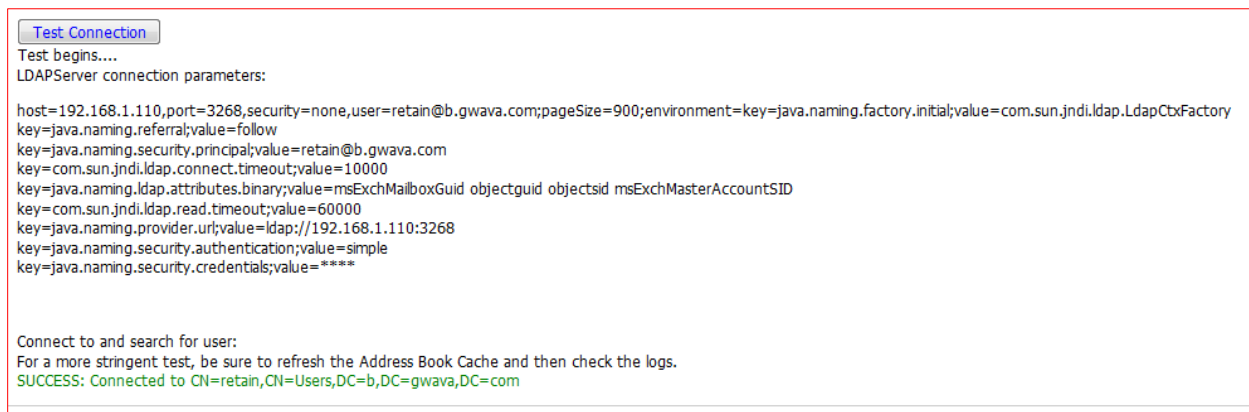
On the Exchange Forest tab, specify whether to use SSL or not for the Global Catalog Security, (strongly recommended), and the "search base". The search base is the LDAP path to the base of where Retain will start searching for valid Exchange users.

The Global Catalog Port defaults depend on whether SSL is used for security or not. SSL is **STRONGLY** recommended. (Default ports are 3268 for plain text, and 3269 for SSL.) Adjust as appropriate for your system.

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 settings](#) section in the appendix. 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 necessary for the Service Account.



Test Connection

Test begins....


LDAPServer connection parameters:

```
host=192.168.1.110,port=3268,security=none,user=retain@b.gwava.com;pageSize=900;environment=key=java.naming.factory.initial;value=com.sun.jndi.ldap.LdapCtxFactory
key=java.naming.referral;value=follow
key=java.naming.security.principal;value=retain@b.gwava.com
key=com.sun.jndi.ldap.connect.timeout;value=10000
key=java.naming.ldap.attributes.binary;value=msExchMailboxGuid objectguid objectsid msExchMasterAccountSID
key=com.sun.jndi.ldap.read.timeout;value=60000
key=java.naming.provider.url;value=ldap://192.168.1.110:3268
key=java.naming.security.authentication;value=simple
key=java.naming.security.credentials;value=****
```

Connect to and search for user:
For a more stringent test, be sure to refresh the Address Book Cache and then check the logs.
SUCCESS: Connected to CN=retain,CN=Users,DC=b,DC=gwava,DC=com

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.

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 forest, to allow the end user to log into Retain.





Resource Forest

Select this checkbox if Exchange is running in a dedicated resource forest. You must then fill out information on the User Forests tab. If this checkbox is not selected, information on the User Forests tab is ignored.

Enable Resource Forest ☐

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.


Exchange Module Save Changes  


Exchange specific information is configured here. At a minimum, the CORE SETTINGS and IMPERSONATION tabs must be filled out completely. Then, examine each of the other tabs to determine if they apply to your system.


User Forests

Ignore this tab if you are running a hosted system.

You need to list every User Forest (dns name) that you will archive in Archive, along with a Global Catalog of that forest, and search base(s) capable of resolving all users. Information on this tab is used only for end-user authentication.



 Add User Forest



 Add Search Base (e.g. dc=users,dc=company,dc=com)

The User Forest must have an entry for each user forest attached to the system. Select the green '+' button and input the LDAP information required by the Forests' Global Catalog server: IP address or hostname, 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.

Save all changes before closing the Exchange Module page.

Exchange Module Save Changes  

Exchange specific information is configured here. At a minimum, the CORE SETTINGS and IMPERSONATION tabs must be filled out completely. Then, examine each of the other tabs to determine if they apply to your system.

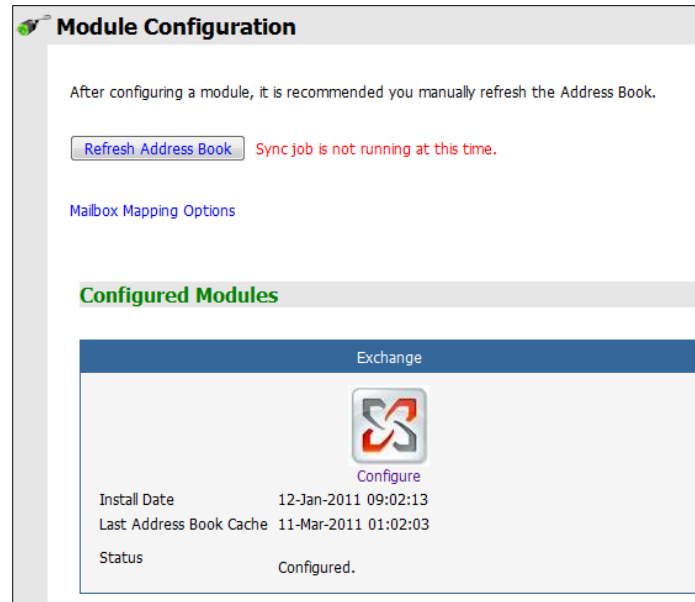
Delegates

Ignore this tab if you are running a hosted system.

Archive can partially support the Delegates feature in Exchange. If a user uses Outlook to grant delegation rights to their mailbox to another user, the second user can have access to the Archive mailbox of the first user. Please note the rights are not granular in this case - the second user either has full rights or no rights to the Archive mailbox.

Enable delegates support ☐

After saving changes, return to the Retain server 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.

Social Messaging

The Social Messaging module allows Retain to archive Social Messaging communications from Facebook and Twitter. To configure the Social Messaging module, all that is required is to provide the connection information and address to the Retain for Social Messaging Gateway. If the RSM Gateway is not yet installed or available, a direct download link is provided on the last tab.




Connections to Social Messaging Gateway are configured here.

Core Settings Gateway Download RSM OVA

Download RSM OVA

You need to [download](#) , install and configure the Retain for Social Messaging OVA, which contains the proxy services needed to capture social media information.

The “Enable Jobs” box must be checked to have any jobs function against the RSM Gateway. See the Retain for Social Media installation guide for more information.



Social Messaging

Connections to Social Messaging Gateway are configured here.

Core Settings Gateway Download RSM OVA


Core Settings

Normally all of these entries should be enabled.

- If Enable Jobs is not enabled, no jobs pertaining to this module will be sent to Workers.

Enable Jobs ☒

On the Gateway tab, the connection information must be filled-out. Either place the host name or the IP Address of the RSM appliance and the connection port for the administration interface. Default is shown.



Social Messaging

Connections to Social Messaging Gateway are configured here.

Core Settings Gateway Download RSM OVA

Connection

You must fill out all of the entries on this tab.

Typically you provide the information necessary to access the Social Messaging proxy.

Host

Port

Security

User

Password (Password is not set)

Security and user login credentials are required to connect to the RSM. The user must have administration rights to gather data from the Gateway. By default the user account is ‘tech’ with the password ‘retain’.

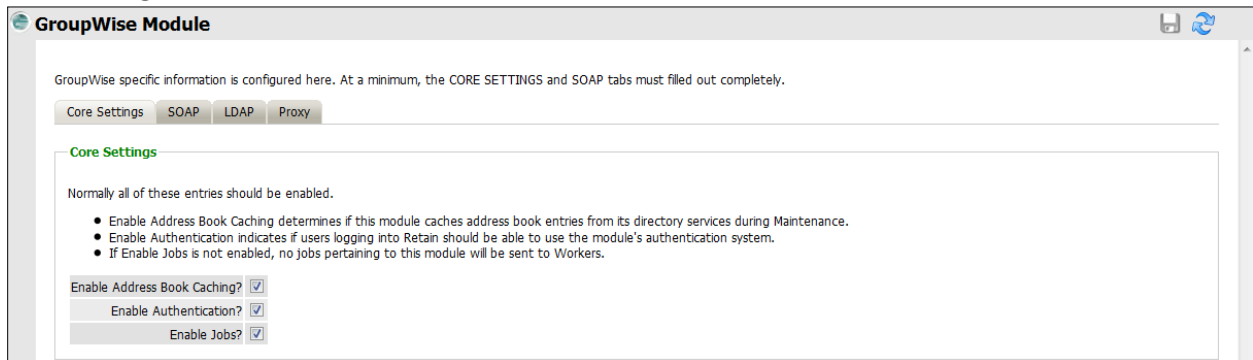
Save changes and the module is configured.

GroupWise Module

The GroupWise module must be fully configured before Retain can archive or communicate with the GroupWise mail system. Retain must be provided with an IP address and port of a readily accessible Post Office (fast access), appropriate login credentials, and a trusted application key.

The GroupWise module page opens first with the Core Settings.

Core Settings



Normally all of 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 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 system 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 configured for this module will be processed.

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 Console One for GroupWise 8.x and higher, or via the provided Key Generator provided and linked from the page. For instructions, see the [trusted app key section in the appendix](#).

The screenshot shows the 'GroupWise Module' configuration window with the 'SOAP' tab selected. The window has a title bar with a save icon and a refresh icon. Below the title bar, a message states: 'GroupWise specific information is configured here. At a minimum, the CORE SETTINGS and SOAP tabs must be filled out completely.' There are four tabs: 'Core Settings', 'SOAP', 'LDAP', and 'Proxy'. The 'SOAP' tab is active and contains two main sections. The first section, 'Trusted Application Key Generation', explains that Retain uses the Trusted Application feature of GroupWise to access user mailboxes and provides instructions on how to generate a key. It includes a list of steps: 'Download the key generation application', 'Run the key generation application', and 'Paste the key generated below'. Below this, there are input fields for 'Trusted Key Name' (containing 'Retain') and 'Trusted Application Key' (containing a long alphanumeric string). The second section, 'GroupWise SOAP Access', explains that Retain accesses the GroupWise message store via the SOAP protocol and provides input fields for 'POA Host Name' (192.168.1.104), 'SOAP Port' (7191), and a checkbox for 'Enable SSL?' (unchecked). Below these fields is a 'Test Connection' button. The third section, 'Post Office Redirection Tables', explains that users can view and edit redirection tables and includes a link for 'Mail Servers'.

GroupWise Module

GroupWise specific information is configured here. At a minimum, the CORE SETTINGS and SOAP tabs must be filled out completely.

Core Settings SOAP LDAP Proxy

Trusted Application Key Generation

Retain uses the Trusted Application feature of GroupWise to access the user mailboxes. You must generate a trusted application key to continue.

The Trusted Application key generation requires a Windows machine with a connection to the GroupWise Primary Domain.

Alternatively, GroupWise 8 SP1 and above include the ability to generate trusted applications on any platform running ConsoleOne.

- Download the key generation application.
- Run the key generation application
- Paste the key generated below

Trusted Key Name: Retain

Trusted Application Key: 1274F00111E00000A5BFEDA2F08107161274F00211E00000ADDCE493233A1E86

GroupWise SOAP Access

Retain accesses the GroupWise message store via the SOAP protocol which needs to be enabled at all Post Offices in ConsoleOne. Retain needs a initial default POA to gather the information from.

POA Host Name: 192.168.1.104

SOAP Port: 7191

Enable SSL? ☐

Test Connection

Post Office Redirection Tables

You may view the cached Post Offices below, and edit the redirection tables used internally by Retain. This may be necessary for complex systems spanning WANs.

Mail Servers

The SOAP access information must be provided, and the connection tested and verified before the system 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.

The screenshot shows the 'Test Connection' button clicked, resulting in a red text message: 'Initiate Soap Connection. Use the following parameters: host:port= 192.168.1.104:7191 trustname=Retain, trustkey=1274F00111E00000A5BFEDA2F08107161274F00211E00000ADDCE493233A1E86 ssl=false'. Below this, a green text message states: 'SUCCESS: Connected to admin'. A note follows: 'For a more stringent test, be sure to refresh the Address Book Cache and then check the logs.' The 'Post Office Redirection Tables' section is also visible, with the same explanatory text and 'Mail Servers' link as in the previous screenshot.

Test Connection

Initiate Soap Connection.
Use the following parameters:
host:port= 192.168.1.104:7191
trustname=Retain, trustkey=1274F00111E00000A5BFEDA2F08107161274F00211E00000ADDCE493233A1E86
ssl=false

SUCCESS: Connected to admin
For a more stringent test, be sure to refresh the Address Book Cache and then check the logs.

Post Office Redirection Tables

You may view the cached Post Offices below, and edit the redirection tables used internally by Retain. This may be necessary for complex systems spanning WANs.

Mail Servers

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.

LDAP

GroupWise Module

GroupWise specific information is configured here. At a minimum, the CORE SETTINGS and SOAP tabs must filled out completely.

Core Settings SOAP **LDAP** Proxy

E-Mail Address Based Login

If enabled, Retain can use the email address of the user as an alternative to the mailbox login name. An LDAP query will be performed, and the password will be matched against the LDAP directory password. Most administrators do not need this functionality - the LDAP Authentication option in GroupWise at the Post Office level largely duplicates this functionality. However, if you want to allow users to login using their e-mail addresses it is required.

- The email address attribute in eDirectory is not indexed by default. For large trees, you may want to index this field. Consult the user manual for details on how to do this.
- The Admin DN, and Search DN need to be in LDAP form, not eDirectory form. The Admin DN should specify a user with rights to all properties in the LDAP Server.

Enable EMail Address lookup? ☐

LDAP Server (DNS/IP) 127.0.0.1

LDAP Port (usually 389 or 636) 389

Use SSL? ☐

LDAP Admin User (e.g. cn=admin,o=beginfinite) cn=admin,o=beginfinite

Password for LDAP Admin User

Top Search Context (e.g. o=beginfinite) o=beginfinite

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 ConsoleOne.

Proxy

GroupWise Module

GroupWise specific information is configured here. At a minimum, the CORE SETTINGS and SOAP tabs must filled out completely.

Core Settings SOAP LDAP **Proxy**

Proxy

You may also enable or disable Proxy Support, which allows users to access mailboxes which they have proxy access to. For performance reasons, proxy access is verified when the user accesses the account, and the verification status is cached, for a period of days that you specify. This does mean that if the proxy access was revoked in GroupWise, several days may pass before access is revoked in Retain. A value of zero disables caching, but is not recommended.

Enable support for GroupWise Proxy ☐

Cache Proxy verifications (days) 7

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.

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.

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.

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.

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 module 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.

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 mailbox server, selected user list, or distribution list. The Worker will not start archiving anything until a Job is created is bound to a Worker, Schedule, and Profile. There is no limit on the amount of Jobs, Schedules, Profiles, or Workers that a Retain Server can be configured to use

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, select the box to do not resend a job that has an elapsed start time. 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 interrupted. For cases where you do not want the job to be interrupted – for 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".

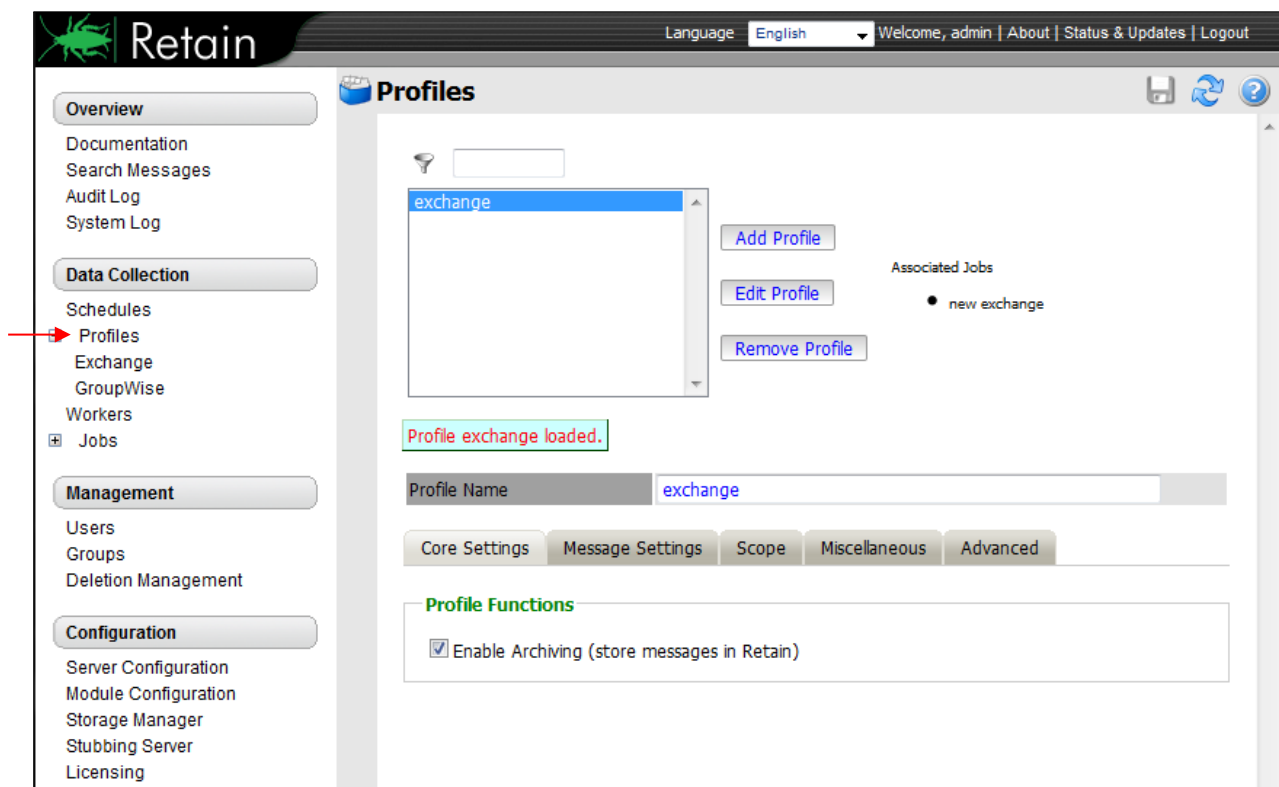
Profiles

A profile assigned to a job tells the system what to archive. Message types, status, date ranges, excluding or including specific users, private or public folders or both, and redundancy check settings are all configured under the profile.

The profile defaults to collecting all data, with no date restrictions. After Retain has finished collecting data the first time, you'll want to narrow the date range and criteria considerably, or pay a high price in performance. Each module has a specific profile page, but different profile pages will not display unless multiple modules have been configured. Each module's profile is explained below.

Exchange Profile

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 configuration 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

The core settings consist of an enabled/disabled option which must be enabled for any jobs based on this profile to archive anything.

Manage Settings

Retain can archive and select specific types of mail and Exchange system items to be archived. The Manage Settings tab provides access to manage those settings.

The screenshot shows the 'Manage Settings' interface with five tabs: 'Core Settings', 'Message Settings', 'Scope', 'Miscellaneous', and 'Advanced'. The 'Message Settings' tab is selected. Below the tabs are four sections, each with a title and a description:

- Mailbox Type:** "You can choose to store items only from User mailboxes or from Room / Equipment resources." It includes two checked checkboxes: ☒ Users and ☒ Room / Equipment.
- Item Type:** "By default, Retain stores every type of message. You can restrict this." It includes five checked checkboxes: ☒ Mail, ☒ Appointment, ☒ Note, ☒ Task, and ☒ Voice Message.
- Item Source:** "Messages are classified in one of these categories. You may restrict the storage of messages if desired." It includes four checked checkboxes: ☒ Received, ☒ Sent, ☒ Draft, and ☒ Posted.
- Message Status:** "You can restrict the storage of messages based upon the setting of various status flags." It includes four dropdown menus: "Read is", "Private is", "Personal is", and "Confidential is". Each dropdown currently shows "Doesn't matter".

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.

A close-up of a dropdown menu. The menu is open, showing three options: "Doesn't matter" (highlighted in blue), "True", and "False".

Scope

The Scope tab dictates how pervasive the attached archiving jobs are, and what kind of duplicate checking is done, or not done, in the Retain system.

The Date Range 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.

The duplicate check either scans for all messages, or instructs Retain to only look for messages which are newer than the Item Store Flag, which is a time-stamp on the system. All messages younger than the time stamp will be archived by Retain if the item store flag is utilized, but Retain will ignore all messages which are older.

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 Journaling and archiving a journaling mailboxes, guarantees all items have been archived.

Core Settings Message Settings **Scope** Miscellaneous Advanced

Date Range to Scan

All Messages (ignore date) ▼

Duplicate Check

Retain only stores a single instance of each message and attachment. Defining how the Worker determines new items (so it may skip sending them to the Server) is an important performance factor. Using the item store flag is recommended. Never publish all messages unless you are priming the system for the first time.

Try to publish all messages (SLOW) ▼

Set Storage Flags

You normally want to advance the Item Store flag after a job is run. However if you run overlapping jobs, you might not want to.

☒ Item Store Flag

It is recommended that the first job be allowed to scan for all mail that is desired to be archived, with the duplicate check set to 'publish all messages'. Afterwards, it is recommended to have the duplicate check look for the Item Store Flag, and only check for new mail in the system.

Miscellaneous

The Miscellaneous tab allows access to particular 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.

The screenshot shows the 'Miscellaneous' tab selected in a settings window. The window has tabs for 'Core Settings', 'Message Settings', 'Scope', 'Miscellaneous', and 'Advanced'. The 'Miscellaneous' section is titled 'Miscellaneous' in green. It contains the following options:

- You can save disk space by restricting attachment storage.**
 - ☒ Store all attachments
 - ☐ Don't store any attachments, other than the message
- You may choose to store and index the Internet Headers. This exacts a performance penalty. However, it may be necessary for your compliance policy.**
 - ☐ Store/index Internet Headers
- Note: The following options can dramatically slow down archiving and increase disk space utilization.**
 - ☐ Include user's archive mailbox
 - ☐ Include user's recoverable items
 - ☐ Include Public Folders(exhaustive)
- You may store Journaling Mailboxes in one flat folder, or partition them by year and/or month.**
 - ☒ Store in one folder
 - ☐ Store by year (yyyy)
 - ☐ Store by year and month (yyyyMM)

A tooltip is visible over the 'Store all attachments' option, showing a list of folder selection options: 'Items from All Folders', 'ONLY items from folders listed below', and 'All folders EXCEPT those listed below'.

Advanced

The screenshot shows the 'Advanced' tab selected in the settings window. The 'Advanced Criteria' section is titled in green and contains the following options:

- You may add advanced criteria such as specific attachment names, sizes, subject here.**
 - Add** button
 - Delete** button
 - A dropdown menu with options: Subject, Sender, Recipient, Size, Attachment Name.
 - A text input field with the value 'is'.
 - A text input field.
- Folder Scope** section:
 - Normally, you want to include all folders in your scan. However, you may have specific reasons to include. Choose a folder root, an optional subfolder (you may have multiple subfolders with the \ delimiter, and whether to include all subfolders in the pattern).**
 - A dropdown menu with the value 'Items from All Folders'.

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 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 you specify.

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

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 ‘Add’ to open a new selection.

- 2) Specify a System Folder (mandatory). Example: Calendar.

- 3) You specify a subfolder of that folder (optional).

Example: entering “old” would mean the folder “old” under “Calendar”.

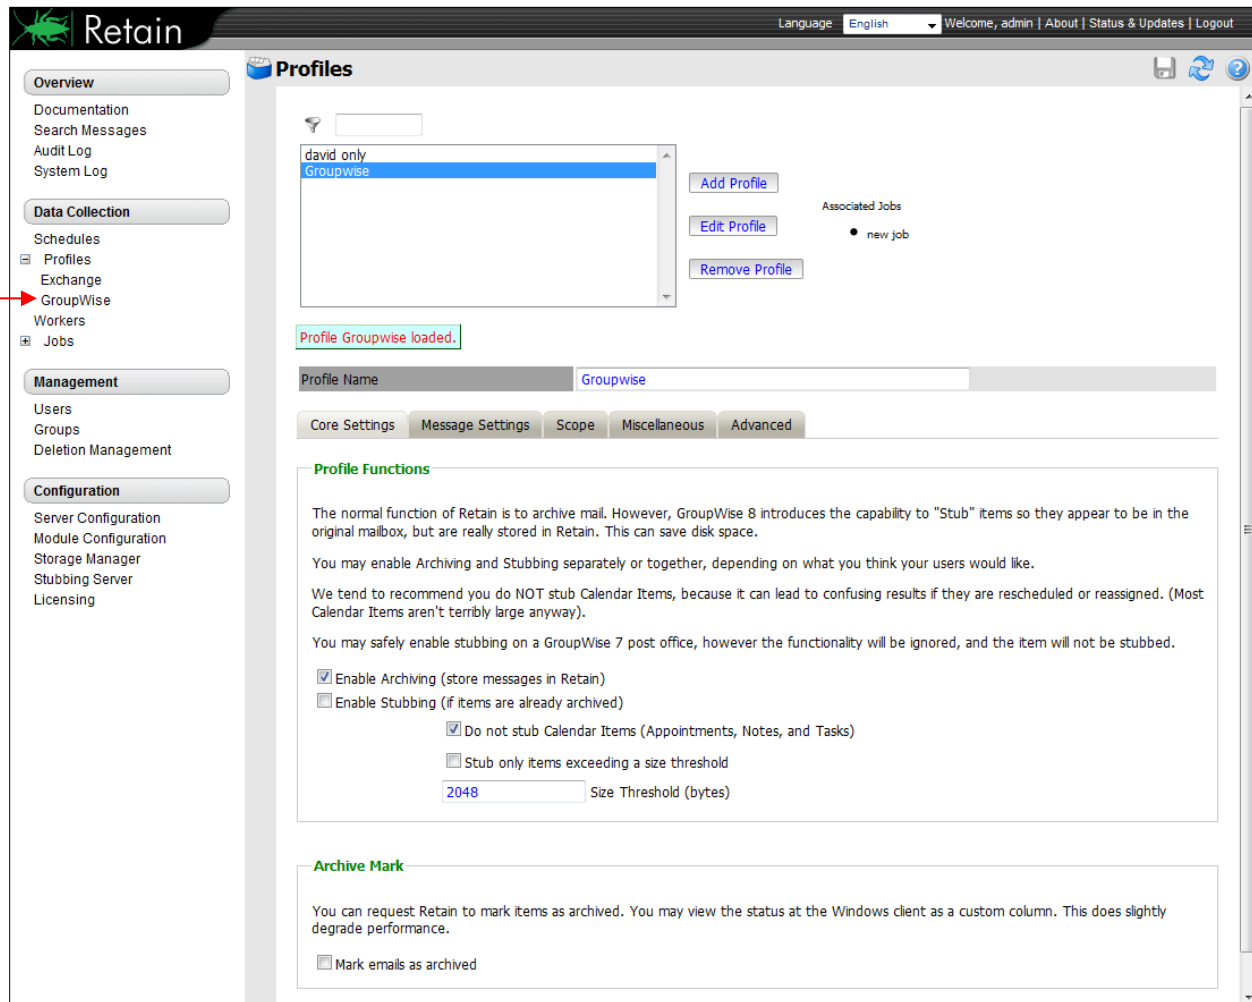
- 4) You can have multiple hierarchies under that with the / delimiter.

Example “old/mail” would mean the subfolder “mail” under “old” under “Calendar”.

- 5) 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.

GroupWise Profile



Core Settings

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

Stubbing

Before enabling stubbing, make sure you understand the way stubbing works and how that will work in your system because **once enabled, you cannot reverse the stubbing process**. Stubbing, (requires GroupWise 8), 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 difference 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:

- 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.
- Mobile, Web, IMAP, and POP3 clients will only see posted empty items, not actual contents.



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](#) section for configuration. 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 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.

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.

	From	Subject	Date ▾	RetainArchived
	Suzie Wathen	cool lightning	8/25/2008 4:21 PM	
	Steve Orchard	Old Ladies' Noggins	8/25/2008 4:17 PM	1

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**.

Message Settings:

- **Mailbox Type:** You can restrict the type of mailbox to users, resources or both.
- **Message Item Type:** 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.
 - Default is to dredge all items, regardless of flag status.

Core Settings
Message Settings
Scope
Miscellaneous
Advanced

Mailbox Type
You can choose to store items only from Users or from Resources.
☒ Users ☒ Resources

Item Type
By default, Retain stores every type of message. You can restrict this.
☒ Mail ☒ Appointment ☒ Note ☒ Task ☒ Phone Message

Item Source
Messages are classified in one of these categories. You may restrict the storage of messages if desired.
☒ Received ☒ Sent ☒ Draft ☒ Posted

Message Status
You can restrict the storage of messages based upon the setting of various status flags in GroupWise.
Opened is Read is Private is

By default, all items are selected for collection.

Scope

Core Settings
Message Settings
Scope
Miscellaneous
Advanced

Date Range to Scan

Duplicate Check
Retain only stores a single instance of each message and attachment. Defining how the Worker determines new items (so it may skip sending them to the Server) is an important performance factor. Using the retention or purge flag choices are the fastest choices, if these are options. Never publish all messages unless you are priming the system for the first time.

Set Storage Flags
If you are using either the Purge or Retention features in GroupWise, you probably want these to be advanced automatically as items are stored so users may delete messages in their mailbox that have been stored by Retain. The Item Store flag is of similar function, stored in Retain itself, but cannot prevent users from deleting item. It is most useful in conjunction with GWAVA Reload or in use with multiple overlapping jobs.
☒ Retention Flag ☐ Purge Flag ☐ Item Store Flag

Date Range

The Date Range determines which message items are collected, depending on the date of the message. While the default is to collect all messages, regardless of dates, the administrator will **want and need** to change this. Dredging all messages makes sense **only** when you are collecting data for the very first run (priming the archive). On a day-to-day basis, however, the performance toll will be intolerable.

For instance, you may want to dredge everything once and then make a daily or weekly dredge for items that are 'n' days old. These limited-scope dredges will run much more quickly and they will save you bandwidth. See below under "duplicate check."

Sometimes, you might just want to dredge items that fall within a specific date range. Specify that here.

Duplicate Check

Duplicate Check

Retain only stores a single instance of each message and attachment. Defining how the Worker determines new items (so it may skip sending them to the Server) is an important performance factor. Using the retention or purge flag choices are the fastest choices, if these are options. Never publish all messages unless you are priming the system for the first time.

Try to publish all messages (SLOW) ▼

Retain is a single-instance-storage system. In other words, any given message is stored only once, to eliminate wasted disk space. Therefore, there are mechanisms in place to check for duplicates as messages are dredged and stored. The Retain Worker dredges the messages, sends them across a network link to the Retain Server and the Retain Server stores them in the database.

Eliminating duplicates will save network bandwidth and it will eliminate wasted disk space.

The Retain Server will check for duplicates before a message is stored. However, sending megabytes of messages and attachments to the Retain Server only to have them discarded because "they are already here" wastes processor time and the network bandwidth between the Worker and the Server. Therefore, you want to be able to do some checking at the Worker side.

If you are using the Reload™ integration, the item store flag is the only duplicate check flag that will work. The Retention and Purge flags will not work against a Reload backup.

In this section, you will tell the worker how to decide whether to dredge a message. Here are your options:

Try to publish all messages (SLOW)
Ignore all messages older than retention flag (fastest)
Ignore all messages older than purge flag (fastest)
Publish all messages newer than last stored message (fast)

"Reload" is a trade mark of GWAVA Inc., 100 Alexis Nihon, Suite 500, Saint Laurent, Quebec H4M 2P1

- **Publish All Messages:** No duplicate checking.
 - Every message is sent to the Server.
 - The Server is responsible for duplicate checking.
 - If there are duplicates, the Server will discard them.
 - This is the slowest option.
- **Use SmartPurge Flag:** Uses the GroupWise flags to determine the newest message.
 - Fastest option.
 - Compares time stamp of the flag to run time.
 - Ignores messages older than flag's time stamp.
 - It is NOT a usable option when dredging from Reload.
- **Use Retention Flag** functions very similarly to the SmartPurge flag.
See below for a comparison
- ☐☐☐ **Use Item Store Flag** Use an internally stored timestamp.
 - Similar in function to GroupWise Retention/Purge flags.
 - Determines last message stored.
 - Fast option but not as fast as Retention/Purge flags.
 - Ignores messages older than flag's time stamp.

Set Flags:

As discussed previously, Retain can use the SmartPurge, Retention, and/or the Item Store flags. Select which flags should be “advanced” (in their timestamp values) as items are archived by this profile. Normally, the Retention flag is the one most administrators will wish to use.

NOTE: *Generally, if you're using a backup system that leverages SmartPurge, like most backup software, (such as GWAVA Reload), does, you will not want to touch that flag.*

About GroupWise's Smart Purge and Retention Flags

The SmartPurge, Retention, and Item Store flags all perform similar functions. The SmartPurge and Retention flags are stored internally in the GroupWise post office on a per-mailbox basis. The Item Store flag is stored by the Retain Server internally in Retain's own configuration database.

Each can be thought of as simply a “timestamp” indicating the last time data was retrieved by a 3rd party application. There are a few subtle differences between the three flags, which are discussed below.

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

With **SmartPurge** enabled in GroupWise,

- The flag is a timestamp, stored in the user's mailbox.
- Items older than the timestamp can be deleted freely.
- Items newer than the timestamp cannot be deleted.
- A 3rd party application, such as backup software can retrieve and/or set this timestamp value. GroupWise never changes the timestamp 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 timestamp once all items are backed up. The backup software can also use this timestamp to “know” that all items older than this timestamp 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 GWAVA 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 SmartPurge flag, the Retention flag was introduced by Novell in GroupWise 6.5.1.

In almost all respects, the Retention flag is exactly identical in purpose and function to the SmartPurge flag. Like the SmartPurge 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 SmartPurge 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 SmartPurge 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 compelling reason not to. A few possibilities are discussed below.

Retain can use the SmartPurge and Retention flags. In addition, a third flag, the Item Store flag is also available for use. The Item Store flag is similar to the other flags, but has some important differences:

- It's stored in the Retain Server's configuration database, instead of in GroupWise. This rarely matters much, but it does mean that the existence and storage of the flag is not dependent on the existence of the mailbox.
- Just like the SmartPurge and Retention flags, it can be used by Retain to skip previously archived messages, enormously increasing performance.
- Unlike the SmartPurge and Retention flags, it cannot be used to prevent users from prematurely deleting message items before they are stored by Retain. Given this severe limitation, the administrator might well wonder why this flag is useful...

It's true that you cannot prevent premature deletion using the Item Store flag. However, there are two scenarios where the Item Store flag becomes useful or even essential:

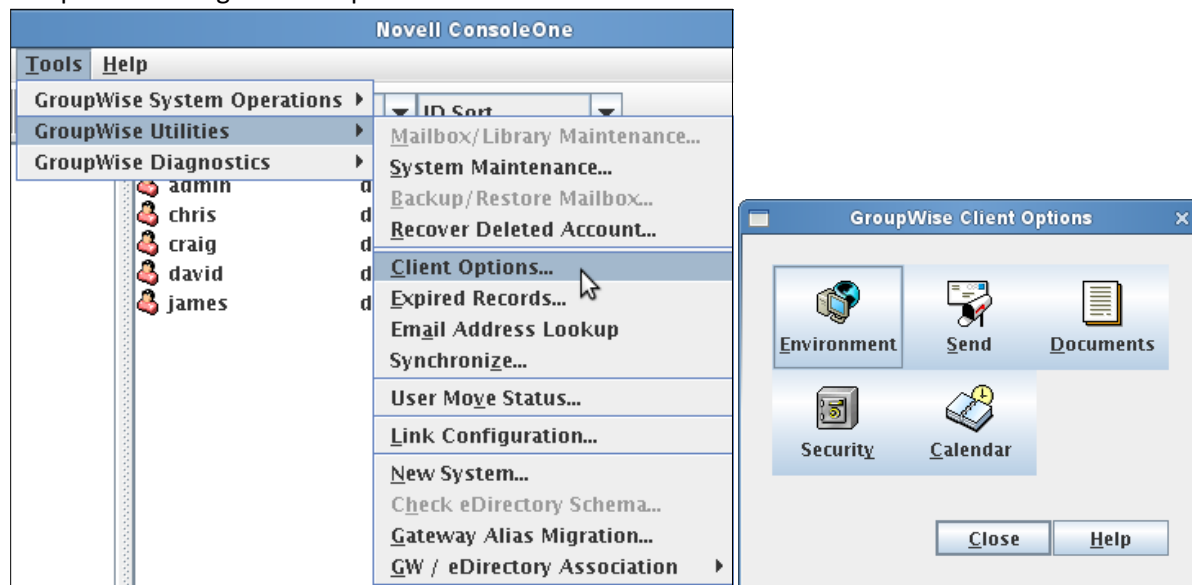
- The first is if for some reason the administrator needs to run multiple jobs or profiles on mailboxes that overlap. In this case, two separate flags are needed so the two different jobs can keep track of where (or more accurately WHEN) to continue archiving from. Admittedly, this is a rare scenario, but without a separate flag, your options would be severely limited.
- More commonly, consider the case of GWAVA Reload. Customers have expressed interest in using the backups created by GWAVA Reload to feed Retain, avoiding unnecessary network bandwidth utilization. But you cannot do this with the built in SmartPurge or Retention flags!
 - GWAVA Reload already uses the SmartPurge flags, and date ranges and profiles between Reload and Retain don't necessarily (and in fact rarely do) match.

- The Retention flag would be perfect, but Retain isn't able to change the Retention flag. Or more accurately, it is changing it on the Reload box, not on the live GroupWise server. Hence, the changes will never be synchronized with GroupWise, leading to the disastrous results of:
 - Users will never again be able to delete e-mail from the live GroupWise system.
 - Retain will rescan the entire post office every night, because Reload will make a new backup, with an unchanged Retention flag!
- The situation is resolved by using the internal Item Store flag.


Setting up SmartPurge and Retention flags in GroupWise 7

(Smart Purge Setup under GroupWise 8 is similar.)

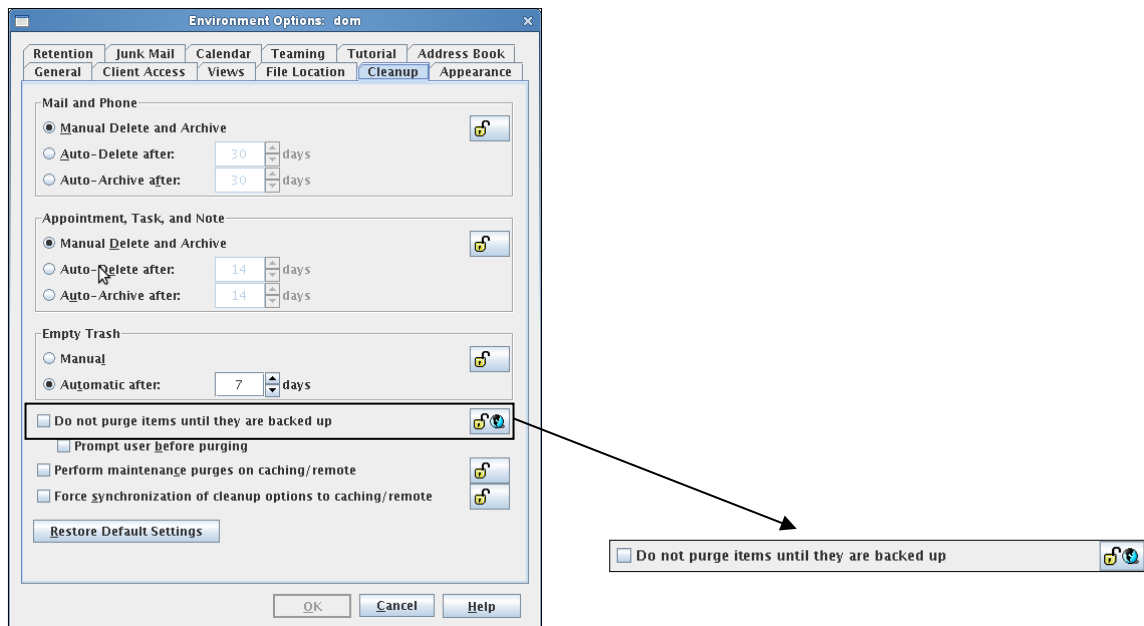
Setup for both flags is accomplished in ConsoleOne:



- 1) Go into ConsoleOne
- 2) Select your domain or post office
- 3) Go to Tools | GroupWise Utilities | Client Options
- 4) Select Environment and the 'Cleanup' tab



Once here, SmartPurge is set up like this:



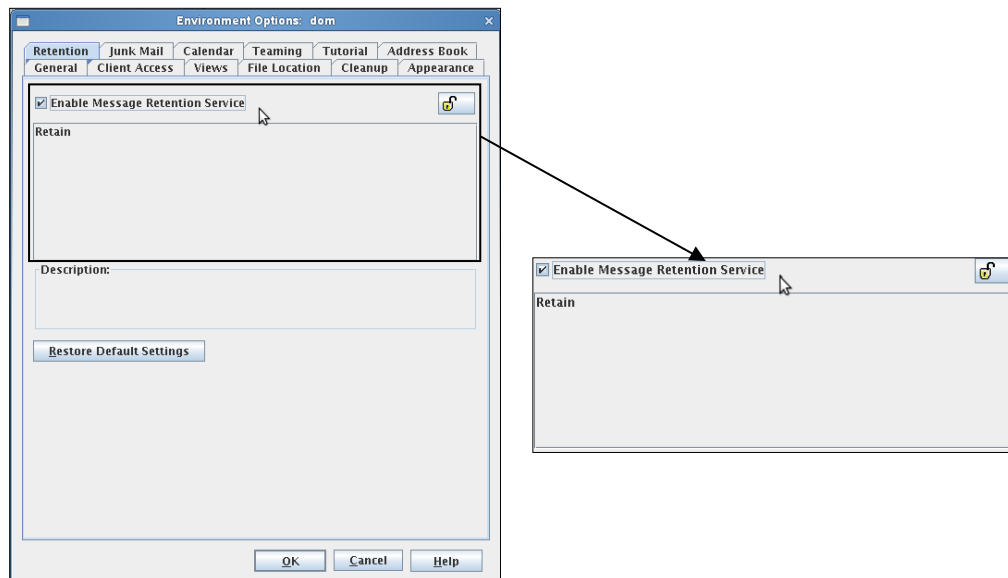
In the Cleanup tab, UNCHECK the box on “Allow purge of items not backed up” to ENABLE SmartPurge. Lock it to prevent users from disabling it.
(GroupWise 8 uses reverse logic with the purge option. GroupWise 8 lists to “Do not purge items until they are backed up.” Reverse the setting for the desired effect.)

Several different backup-retention services may currently use this same timestamp. Plan your system accordingly as Retain also uses this flag. The SmartPurge feature ONLY uses the backup / purge flag set as shown above. If another flag is desired, use of the Item Store Flag is recommended.

Retention Flag

The Retention flag is set up a little differently than SmartPurge. (This is only used when Stubbing.)


- 1) A Trusted App requesting Retention Services must be created. (Done at Retain setup time).
- 2) The Retention flag is enabled on a per domain/po/user basis.
- 3) From the 'Retention' tab under Environment Options, the setting should look like this:



Notice in this example that Retain is listed as a trusted application. This was done during our post-install tasks when we ran the Trusted Application Key program.

Next, CHECK the box to ENABLE the Retention Flag service. Lock it on.

Miscellaneous



The screenshot shows a settings window with five tabs: 'Core Settings', 'Message Settings', 'Scope', 'Miscellaneous', and 'Advanced'. The 'Miscellaneous' tab is selected and highlighted. Below the tabs, the 'Miscellaneous' section is titled in green. The text reads: 'You can save disk space by restricting attachment storage. You can also attach the Routing Properties to the message, if desired.' There are three radio button options: 'Store all attachments' (selected), 'Don't store MIME.822 attachments', and 'Don't store any attachments, other than the message'. Below this, there is explanatory text: 'You may choose to store Sent Items properties, and/or store and index the Internet Headers. Both of these (particularly Sent Items properties) exact a performance penalty. However, it may be worthwhile to store and index the Internet Header if your compliance policy then allows Retain to choose to skip storing MIME.822 (which consumes lots of disk space since it's not usually single-instanceable).' At the bottom, there are two checkbox options: 'Store/index Internet Headers' and 'Include Routing Properties', both of which are currently unchecked.

Miscellaneous: What to do with attachments

Most of the time, you will want to store all attachments. That is the default.

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 store 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.

Advanced

The screenshot shows a software interface with two main sections. The top section, titled 'Advanced Criteria', contains a text box with the instruction: 'You may add advanced criteria such as specific attachment names, sizes, subject here.' Below this is an 'Add' button and a 'Delete' button. A dropdown menu is open, showing options: 'Subject', 'Sender', 'Recipient', 'Size', and 'Attachment Name'. The 'Subject' option is selected. To the right of the dropdown is a text input field and a dropdown menu with 'is' selected. The bottom section, titled 'Folder Scope', contains a text box with the instruction: 'Normally, you want to include all folders in your scan. However, you may have specific reasons to include. Choose a folder root, an optional subfolder (you may have multiple subfolders with the \ delimiter, and whether to include all subfolders in the pattern.' Below this is a dropdown menu with three options: 'Items from All Folders', 'Items from All Folders ONLY items from folders listed below', and 'All folders EXCEPT those listed below'. The 'Items from All Folders' option is selected.

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 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 you specify.

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

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


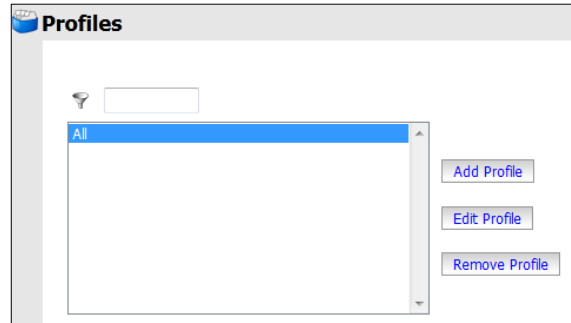
How to specify the list of folders to dredge/exclude:

- 6) Specify a System Folder (mandatory). Example: Calendar.
- 7) You specify a subfolder of that folder (optional).
Example: entering “old” would mean the folder “old” under “Calendar”.
- 8) You can have multiple hierarchies under that with the / delimiter.
Example “old/mail” would mean the subfolder “mail” under “old” under “Calendar”.
- 9) 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.

Social Messaging Profile

To create a new profile to archive social messaging data, first select the 'add profile' button and name the profile. Then continue configuring as desired.

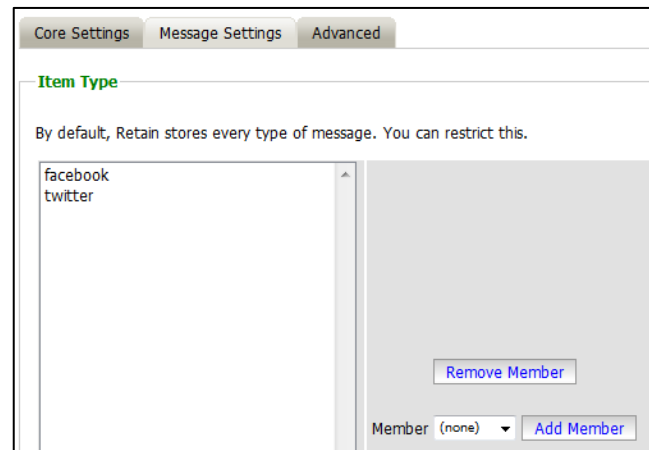
The **Core Settings** tab holds the enabling setting, which allows jobs connected to this profile to run. This must be checked before jobs will run.



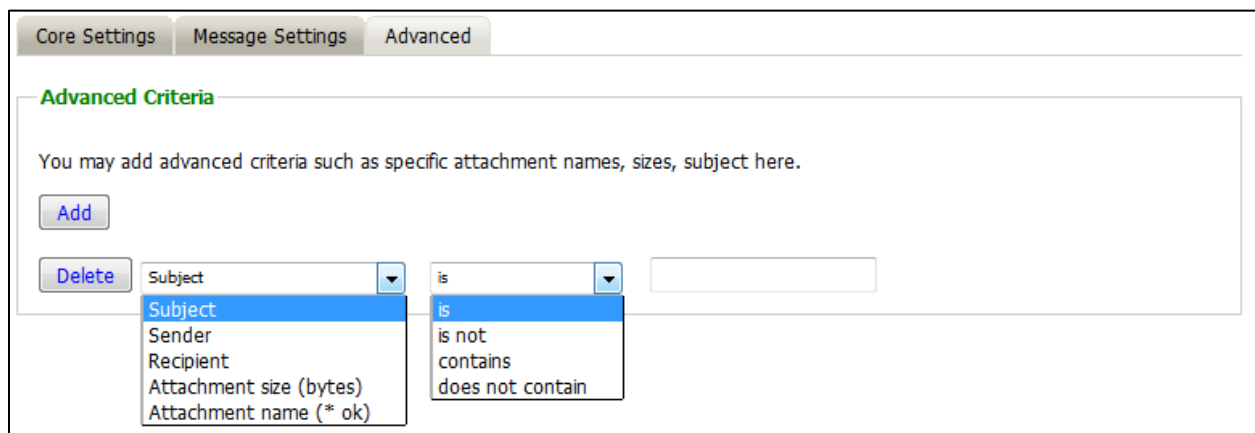
The **Message Settings** tab can be used to exclude or include specific data types for this profile. By default all types of messages will be archived. If left blank, the default will persist.

If only specific data types are to be included in a profile, they can be selected from the Member drop-down list, and added to the inclusion list.

All types of data included in the list will be archived. If the list is blank, all users and data will be archived.



Under the **Advanced** tab, specific criteria can be used to restrict the archived data. Any combination of the variables listed may be used to restrict the data selection. Advanced criteria may be added up to 7 fields.



Save changes and the profile is configured.

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 purposes. 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.

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.
 - Default 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 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.
- **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.

Worker Configuration

From the Retain home page, click Workers. You will see this screen:



Workers

migrator
new worker

Add Worker
Edit Worker
Remove Worker

No jobs are associated with this worker.

Enter new worker below

Worker: new worker


Polling Logging Connection Module specific

Interval

How often should the worker poll for configuration changes, and how many messages should the worker send before updating status?

Poll for configuration every 10 minutes
Update status every 500 messages.

- 1) Click “Add Worker”.
- 2) Put the name of the worker in the “Worker” box. *Example:* “Linux Worker 1”.
- 3) 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.



Polling Logging Connection Module specific

Server Connection

The worker will need to know how to connect to the Server. It is highly recommended that you use a DNS host name for the Host Name. It's important to note that you are specifying the URL from the point of view of the **Worker contacting the Server**, which may or may not match the URL you'd use from your browser to reach the Worker or Server.

For administrators using the Reload interface, note the worker password assigned must be the same for each worker that contacts the Reload Server. As part of setting of the Reload Server, you must know this password.

Server Protocol: http
Server Host Name: 192.168.1.106
Server Port: 48080
Server Path: /RetainServer
Worker Password:

Note: Changing any of these parameters will require recopying the worker bootstrap file!

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 setting 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.

- 4) 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.

Depending on the modules you have configured, you will need to provide additional configuration information here.

GroupWise SOAP Access

When using GroupWise, the Retain Worker needs to connect via SOAP to a initial default POA (preferably one local to the Worker) to gather job membership information. Initially this is set to the default you specified in the Groupwise Module Configuration, but it can be changed on a per worker basis.

POA Host Name:

SOAP Port:

Enable SSL: ☐

Exchange

As the worker processes each mailbox, it performs an auto-discovery process to find the "best" server to connect to. Providing an Active Directory Site improves the accuracy of the algorithm, although this is optional.

Active Directory Site (case sensitive):

- 5) After you have checked and set your settings, click **"Save Changes"** before browsing to another page to finish creating and to save your worker.

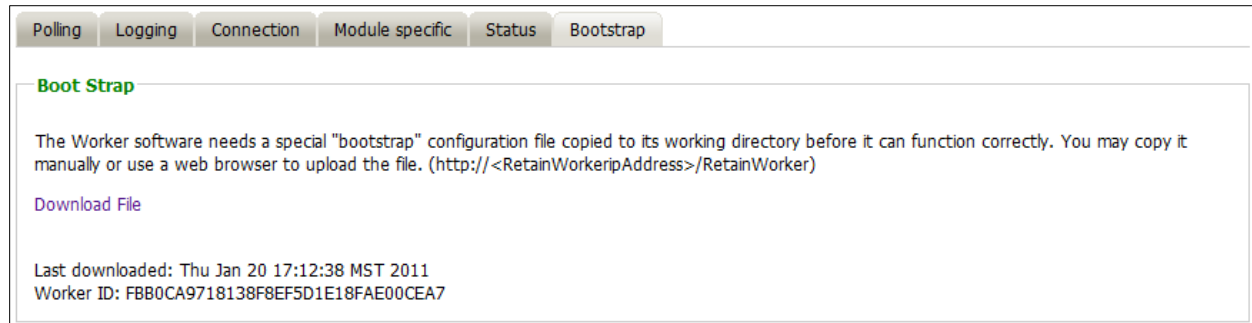
After the worker has been saved, two additional tabs become visible for settings on the worker. You MUST save changes on a newly created worker before the 'Bootstrap' tab is available to be used.

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.


Polling	Logging	Connection	Module specific	Status	Bootstrap
Worker Status					
Registered	26-Jan-2011 05:07:25				
Last contact	09-Mar-2011 20:35:39				
Last Job	new exchange {D69530CFFBACFFE9F4A32EE560D13257}				
Last job start	09-Mar-2011 20:35:39				
Last job end	09-Mar-2011 20:35:42				
Total mailboxes scanned	7				
Total messages stored	18590				

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.

- 6) After you have saved the changes, the Bootstrap tab appears (see below). Click “Download File” to save the configuration to a local workstation.



- 7) Save this file to your hard drive in a place where you can find it. Click “Save”.



- 8) Open a new browser window and type:

http://< your-RetainWorker-ip >/RetainWorker

This screen appears when a worker has not been initialized.

9) Upload the bootstrap file.

Worker

Each worker needs to be configured in the Archive Server web interface. Once a worker has been configured a "bootstrap" file needs to be downloaded and applied to the worker here.

Without a bootstrap file the Worker will not know how to find the Archive server. After applying the bootstrap file any configuration changes made in the server will be polled by this worker. You will not have to reconfigure the Worker manually unless you change the Connectivity settings configured in the Archive Server.

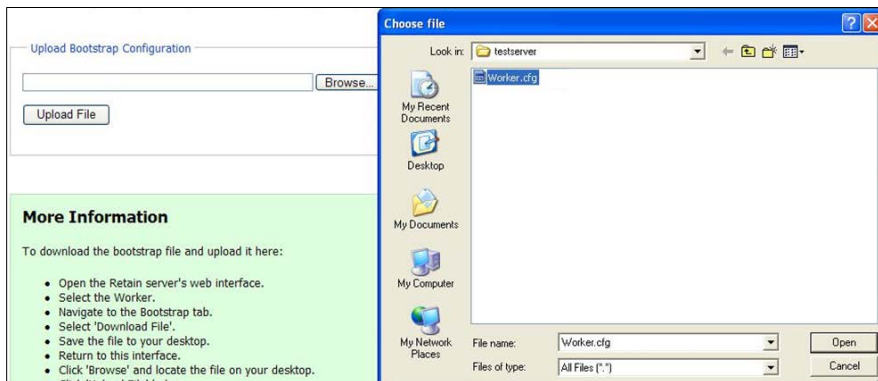
Upload Bootstrap Configuration

More Information

To download the bootstrap file and upload it here:

- Open the Archive server's web interface.
- Select the Worker.
- Navigate to the Bootstrap tab.
- Select 'Download File'.
- Save the file to your desktop.
- Return to this interface.
- Click 'Browse' and locate the file on your desktop.
- Click 'Upload File' below.

- Click "Browse" to find the file you just downloaded to select it.




- Click "Open".
- Click "Upload File".
- If the upload went well, you will see this screen:

Configuration Complete. From now on, you'll have to log in as a user (such as the initially created admin user) with the Manage Workers right.

After you click the "Continue" button, you will be returned to the login page.

10) Login as Admin.




Login

Password

Login

On all subsequent logins, you should be shown the configuration page first.



Configuration Server Connection Status Scheduler

Worker Configuration

Friendly Name	linux worker
Worker Unique ID	A08A76358BD4025F0A0E7E76800C978F
URI to Server	http://10.1.1.103/RetainServer
Worker Version	1.6.0
Memory Free/Max/Total	59551/1048576/262654

[Check for New Version](#)

[Refresh job cache now](#)

Click "Refresh job cache now" to prompt the worker to poll for configuration changes right now.

Bootstrap information has been applied to this server. If the configuration must be changed in the future, replace the configuration file at /opt/beginfinite/retain/RetainWorker/WEB-INF/cfg/ and restart the Worker.

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.

11) Verify the Retain Server connection success by selecting the Server Connection tab.




Configuration Server Connection Status Scheduler

Server Connection

Successfully connected to server

- Go back to the Retain Server.
- Click “Workers”.
- 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. You may use “Diagnostic” at first, until the system is fully operational and tested. 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.

Jobs

Jobs are separated into different modules. A module, schedule, profile, and worker must be configured before a Job can be created. Module specific job menu's only display when multiple modules are configured. Jobs bring the entire configuration together to archive mail from the mail server, according to the configuration. Jobs are where data collection is achieved.

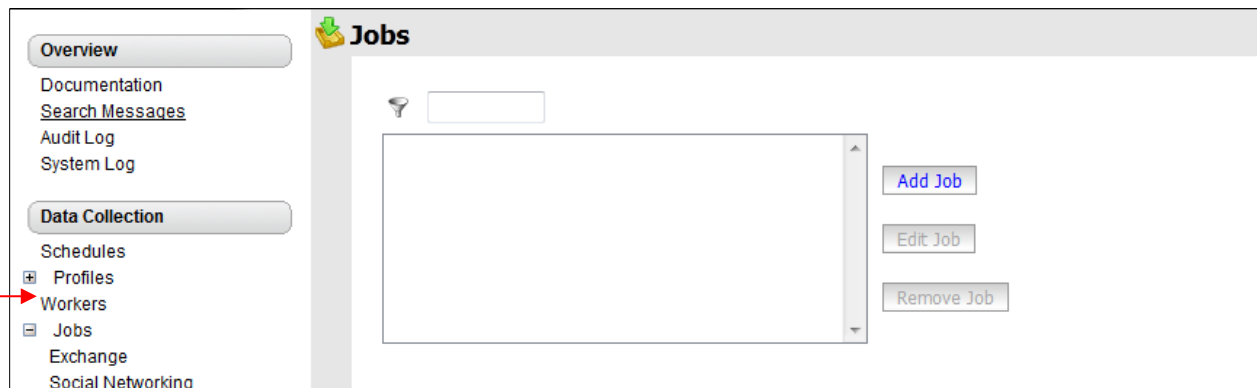
A job:

- Starts according to the selected **SCHEDULE**.
- Collects data according to the selected **PROFILE**.
- Utilizes the selected **WORKER**, and only the selected worker according to the schedule and profile.
- Collects data from the specified mail source, (messaging server, distribution list, group, and individual user)

Jobs look and act differently depending on which module they are configured under. Select the desired module from the Job menu, (or simply 'Jobs' if only module is configured), and configure according to job needs.

Exchange Jobs

All mail archiving is accomplished through this interface because the methods to interact with the Exchange system are tied into this job interface. Only Exchange profiles will be able to be viewed or selected through this interface, and no other modules profile will be selectable here. However, schedules and workers are universal and can be selected in any modules' job interface.



To create a new job, select the 'Add Job' button and provide a name for the job. To save a job, the job must have a Profile, Worker, Schedule, and a mailbox, distribution list, or server to be archived. In other words, a job must have a time to run, types of things to get, something to get it with, and a target or a place to get it, in order to be saved and run.

A job without an assigned profile, schedule, worker, or mailbox cannot be saved.

Core Settings

Under 'Core Settings' the profile, schedule, and worker are selected. Only existing profiles, workers, and schedules may be selected. The selected schedule will determine when the job is run; the Profile determines what kinds of mail are archived, and the worker is the agent which performs the work.

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, Caringo, 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.

This time stamp simply puts a date on which the mail is "set to expire", and can be in days, months, or years. Mail that has 'expired' due to this time stamp will remain in the database until removed by the deletion manager.

Job Expiration is not retroactive for mail in the database, and only applies to mail archived by the job that it is active for.

There may be some folders which it may not be desirable to place an expiration time stamp on. Such folders may be specified through the 'add' button at the bottom of the page.

The base folders and criteria specified here. All messages included in these folders will have a different, if specified, or exempted from the expiration date.

Journaling

Journaling **MUST** be enabled on the Exchange system, and Journaling mailboxes must be provided for legal archive compliance.

To archive messages contained in the Journaling mailbox, specify the Journaling mailbox by selecting it from the user/mailbox list. Open the mailbox selection window by selecting the 'Journaling mailbox' button and searching for the desired mailbox.

After searching for mailboxes and users in the cached address book, the users must be added to the list by selecting the checkbox next to the desired name in the search results box, then selecting the 'Add Selected' button at the top.

The red 'X' next to the names in the 'Currently Selected' list allows removal of selected items. Once the list has been configured as designed, selecting the 'Ok' button from the top of the window will add the selected names to the list.


First	Last	Mail Server	Mailbox	Dept.	Type	Domain
<input type="checkbox"/>	admin	po	admin		U	dom
<input type="checkbox"/>	David	Bassett	po	David	U	dom
<input type="checkbox"/>	Samuel	Clements	po	Sam	U	dom
<input type="checkbox"/>	Benjamin	Rush	po	Benjamin	U	dom

A 'Funnel' mailbox creates a single mailbox in Retain that 'funnels' all the mail from all the different journaling mailboxes into one. If this is not specified, then each Journaling mailbox will be created and archived in the Retain system as any normal mailbox would be archived and created.

It is **HIGHLY** recommended to delete archived items from the journaling mailbox to avoid bloating the existing message system or requiring manual maintenance and management of the Journaling mailboxes.

Mailboxes

The final setting that a Job must have in order to be saved and run, is a target: where to archive mail from. The Mailboxes tab dictates where a Job will look for mail to archive. Only the users or mail server included in a job will be archived by the job when it is run.



The screenshot shows the 'Mailboxes' tab selected in a navigation bar at the top, which also includes 'Core Settings', 'Journaling', 'Notification', and 'Status'. Below the tab, the 'Mailboxes' section is titled in green. It contains two instructional paragraphs: 'Assign the Mail Servers, Database stores, and/or Distribution Lists to be included in this job.' and 'Make sure all entries belong to the domain you specified in the Core Settings tab, or they will not be saved.' Below these instructions are three selection areas. The first is for 'Mail Servers', featuring a search box and a list with 'WINSERV0864' selected. The second is for 'Distribution Lists', also with a search box. The third is for 'Users', with a search box. Each selection area has a plus icon to the left and a funnel icon to the right.

Users may be specified separately, whole servers may be selected, and distribution lists spanning different servers may be selected. A job may span all the selected options, or only one of them, but jobs will be faster and more efficient if they are balanced in their scope, for instance, one separate job for each server each with its respective worker, to allow multiple jobs to be run simultaneously without overwhelming one server or worker.

The open text box next to the selection area is a search or filtering interface which allows the administrator to only display distribution lists or mail servers that match the criteria. A blank criteria window will result in all known servers or distribution lists being displayed.

If a user is desired to be included in or excluded from a job, the selection method is identical to that of the Journaling mailbox. Search for the user via the Select Mailbox window and add it to the list.

Mail servers in Exchange may also be expanded into their respective mailbox databases, and individual, or selective mailbox databases may be selected or excluded from any specific job.

Notification

All errors, summaries, and statistics may be sent in an email notification to a desired address on job completion. If it is desired to be notified, configure the notifications under this tab. At least one of the two notification options must be selected in order for an email notification to be created.

The screenshot shows the 'Notification' tab selected in the top navigation bar. Below the tab, there is a section titled 'Notification' with a description: 'You can have Archive e-mail a summary of operations and/or errors that occurred when the job completes.' Below this, there are several input fields for SMTP configuration: 'SMTP Mail Server' (192.168.1.121), 'SMTP Mail From Address' (admin@company.com), 'SMTP To Address' (admin@company.com), 'SMTP Username' (sam), and 'SMTP Password' (masked with dots, with a note '(Password is default)'). There are also two checkboxes: 'Mail when errors occur' and 'Mail summary when job complete', both of which are currently unchecked. At the bottom of the form is a 'Test Connection' button.

Status

On initial creation, there will be nothing of interest shown under the 'Status' tab, as the job is under creation 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.

The screenshot shows the 'Status' tab selected in the top navigation bar. Below the tab, there is a section titled 'Status [Refreshes in 0]' with a job ID: 'Job: new exchange (D69530CFFBACFFE9F4A32EE560D13257)'. Below this, there are two tables. The first table shows job statistics, and the second table shows mailbox statistics.

	Current	Last	Overall
Job began	---	04-Mar-2011 10:35:46	---
Job ended	---	04-Mar-2011 10:35:50	---
Mailboxes	0	2	2
Messages	0	0	0
Errors	0	0	0

	Current	Last	Error
Mailbox			
MS			
Error			---

GroupWise Jobs

All data collection for the GroupWise is configured through the GroupWise job interface. The jobs here combine the GroupWise 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.

Core Settings

The screenshot shows the 'Jobs' management interface. On the left is a sidebar with navigation links: Overview, Documentation, Search Messages, Audit Log, System Log, Data Collection (Schedules, Profiles, Workers, Jobs), Exchange, GroupWise, Management (Users, Groups, Deletion Management), and Configuration (Server Configuration, Module Configuration, Storage Manager, Stubbing Server, Licensing). The main area is titled 'Jobs' and contains a list of jobs with 'david only' and 'new job' visible. Below the list is a message 'Job new job loaded.' and a 'Job' dropdown set to 'new job'. A tabbed interface shows 'Core Settings' selected. The 'Core Settings' section includes: 'Job enabled' (checked), a note that each job must have a schedule, profile, and worker, and three dropdowns for 'Schedule' (full run job), 'Profile' (Groupwise), and 'Worker' (new worker). Below these are 'Enable data expiration' (unchecked), 'Expire in' (0 Days), and 'Base Expiration Time on' (Date Item is stored in Retain (recommended)). A final 'Add' button is at the bottom.

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 in order to run. Only previously configured Schedules, Profiles, and Workers can be selected.

The **Job 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, Caringo, 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.

Mailboxes

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 Post Office 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.

NOTE: Retain cannot access certain GroupWise accounts due to GroupWise account limitations. Expired and Disabled accounts are inaccessible to Retain. Prior to GroupWise 7.0.3, limited license accounts also cannot be accessed by Retain.

If an expired or disabled user account, (or limited license accounts under GroupWise 7.0.3), is selected for a job, or a post office containing such user accounts is selected, Retain will notify the administrator of a failure to archive those users. This is not an error, it is a limitation imposed by GroupWise. GroupWise and Retain are functioning correctly.

Reload

Selecting this option tells the job to use the Reload integration for systems utilizing Reload. See the [Reload integration](#) section for more information. The Server URL must be edited with the correct contact info for the reload server. The default Reload port is 5555.

The screenshot shows the 'Reload' tab in the configuration interface. At the top are tabs for 'Core Settings', 'Mailboxes', 'Reload', 'Notification', and 'Status'. The 'Reload' section has a green header. Below the header is a text block explaining that users can backup their GWAVA Reload server to save bandwidth, requiring Reload 2.5 or newer and configuration changes. There are two input fields: 'Enable Reload Integration' with a checkbox, and 'Reload Server URL' with the text 'http://myReloadBox:5'. A link for 'Advanced Reload Settings' is at the bottom.

Core Settings Mailboxes Reload Notification Status

Reload

You can backup your GWAVA Reload server, saving network bandwidth if desired. This requires Reload 2.5 or newer and some configuration changes on the Reload machine. Please consult the user manual for additional information.

Enable Reload Integration ☐

Reload Server URL

[Advanced Reload Settings](#)

Retain instructs Reload to load the corresponding backup set for each date range in the profile. If a different 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.

The screenshot shows the 'Advanced Reload Settings' page. It has a purple header. Below the header is a text block explaining that Reload sets correspond to the date range in the profile, but users can override this. There is a checkbox for 'Override Reload Date Range'. Below this are two rows of date pickers: 'Start' and 'End'. Each row has dropdowns for year (1995), month (Jan), and day (01), followed by a calendar icon.

[Advanced Reload Settings](#)

Normally the Reload Sets that are backed up correspond to the Date Range specified in the profile. Occasionally, you might want to override this.

Override Reload Date Range ☐

Start: 1995 Jan 01

End: 1995 Jan 01

Notification

When a job runs, you can request an error report or job summary to be sent by e-mail.

The screenshot shows the 'Notification' tab in the configuration interface. At the top are tabs for 'Core Settings', 'Mailboxes', 'Reload', 'Notification', and 'Status'. The 'Notification' section has a green header. Below the header is a text block explaining that users can have Retain e-mail a summary of operations and/or errors. There are six input fields: 'SMTP Mail Server' (192.168.1.104), 'SMTP Mail From Address' (retain@gwava.com), 'SMTP To Address' (admin@gwava.com), 'SMTP Username' (chris), 'SMTP Password' (masked with dots), 'Mail when errors occur' (checkbox), and 'Mail summary when job complete' (checkbox). A 'Test Connection' button is at the bottom.

Core Settings Mailboxes Reload Notification Status

Notification

You can have Retain e-mail a summary of operations and/or errors that occurred when the job completes.

SMTP Mail Server

SMTP Mail From Address

SMTP To Address

SMTP Username

SMTP Password

Mail when errors occur ☐

Mail summary when job complete ☐

[Test Connection](#)

Job Status

When a job is not running, the status tab displays information about the last time it ran:

Status [Refreshes in 1]			
Job: daily (582844007495D6705B173FFE5E4EC23B)			
	Current	Last	Overall
Job Began	Not running	12-Oct-2007 12:00:00	---
Job Ended	---	16-Oct-2007 12:42:04	---
Mailboxes	46	46	46
Messages	7596	238896	238896
Errors	0	0	0
	Current	Last	Error
Mailbox	steve	beta	
PO	ripo	ripo	
Error			---

When a job is running, you have the option of aborting the job:

Status [Refreshes in 7]			
Job: first backup (27318D8551ED008C24B8EABC6CDF8A36)			
	Current	Last	Overall
Job Began	19-Feb-2008 13:10:00	never	---
Job Ended	Abort Job	never	---
Completed: 12 / 13 (92%)			
Mailboxes	13	0	0
Messages	3568	0	0
Errors	0	0	0
	Current	Last	Error
Mailbox	lindy	craig	
PO	po-702	po-702	
Error			---

This window will refresh every ten seconds to keep you up to date as to the status of the selected job. Is it running, did it run and when, how long did it take, what did it accomplish?

In the title, it displays a refresh countdown in seconds.

The completed line is displayed during an active job. The completed status is a display of how many mailboxes 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 according to time.

Social Messaging Jobs

All data collection for the Social Messaging module is configured through the Social Messaging 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.

Core Settings

The screenshot shows the 'Core Settings' tab of a Social Messaging Job configuration window. At the top are four tabs: 'Core Settings' (selected), 'Mailboxes', 'Notification', and 'Status'. Below the tabs, the 'Core Settings' section is titled in green. It contains a 'Job enabled' checkbox which is checked. Below this, a text label states 'Each Job must have the following assigned to it:' followed by a bulleted list: '• schedule - How often to run', '• profile - What to do when running', and '• worker - Where to run the job'. Underneath the list are three dropdown menus labeled 'Schedule', 'Profile', and 'Worker', all of which currently show '(none)'. Below these dropdowns, a text label explains: 'Items archived with this job can optionally be set with an "expiration timestamp" which the Deletion Manager can use to remove obsolete items. The expiration timestamp can only be extended, never reduced.' This is followed by an 'Enable data expiration' checkbox which is unchecked. Below the checkbox is an 'Expire in' field with the value '0' and a 'Days' dropdown menu. At the bottom, there is a 'Base Expiration Time on' dropdown menu with the selected option being 'Date Item is stored in Retain (recommended)'.

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 in order to run. Only previously configured Schedules, Profiles, and Workers can be selected.

The **Job 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, Caringo, 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 data in the database, and only applies to data archived by the job that it is active for.

Mailboxes

The mailboxes tab contains the target users, distribution lists or groups to be archived by the job. By default, all users are included.

The job may be restricted to specific users, either on the 'Include Users' list. To add users, expand the User's option box and select the users to be archived. For users to be available for this dialog box, the RSM must have been configured with users. Users may be selected through the user search interface, spawned through selecting the 'add user' button.

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.

Notification

All errors, summaries, and statistics may be sent in an email notification to a desired address on job completion. If it is desired to be notified, configure the notifications under this tab. At least one of the two notification options must be selected in order for an email notification to be created.

Status

On initial creation, there will be nothing of interest shown under the 'Status' tab, as the job is under creation 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.

Job: first (78778DFEC4A3E6D22C03108FB693079A)			
	Current	Last	Overall
Job began	---	15-Sep-2012 15:38:09	---
Job ended	---	15-Sep-2012 15:42:09	---
Mailboxes	0	4	4
Messages	0	3321	3321
Errors	0	0	0
	Current	Last	Error
Mailbox		admin	
MS		po	
Error			---

Management

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 others 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.

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) in [Server Configuration](#).
- They 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*.

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.

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 *Edit User* or *Remove User* buttons.



Core Settings

Comment			
Authentication Method	Offline Authentication	<input checked="" type="checkbox"/> Use Exclusively	
Primary UID (links your rights)	2edf61cd-4b81-49bc-953a-f049f0e5be65		
Group Membership	default		
This account does not expire	<input checked="" type="checkbox"/>		
Change Internal Password	New	Confirm	
	<input type="checkbox"/> Prevent user from changing		
Language	English		
Disable Account	<input type="checkbox"/> Date:	2012	Jul 27



Users

admin

[Add User](#) [Edit User](#) [Remove User](#)

User admin loaded.

User ID: admin

Core Settings User Rights Mailboxes Miscellaneous

Core Settings

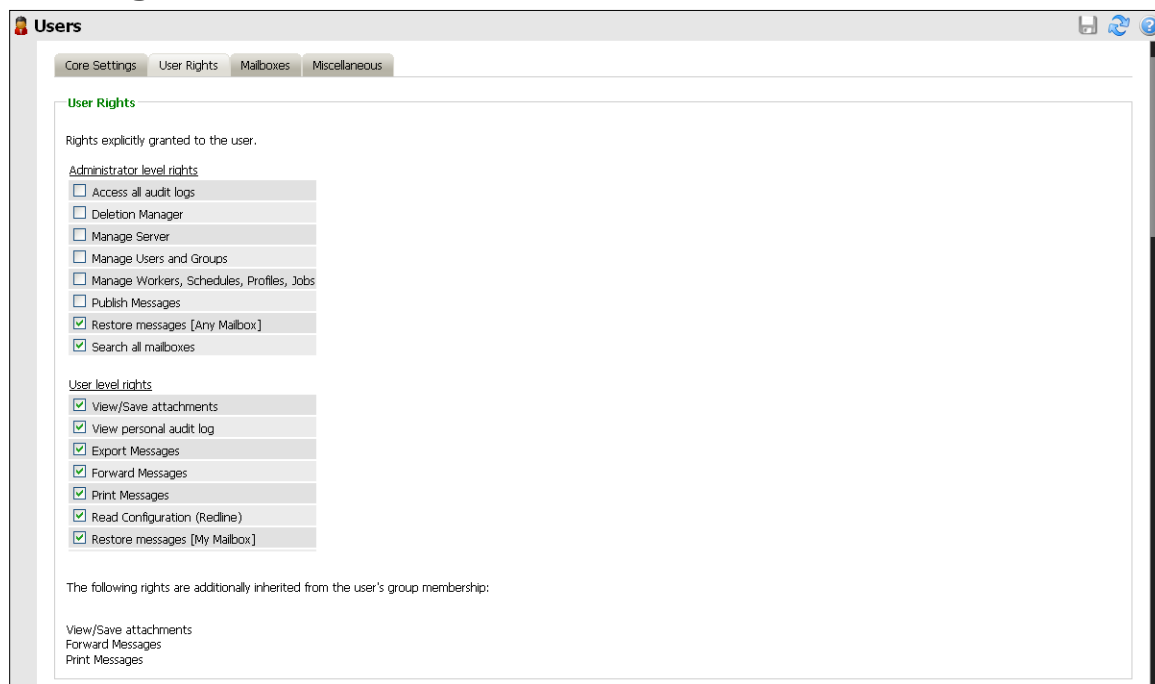
Authentication Method	Offline Authentication	<input checked="" type="checkbox"/> Use Exclusively?
Primary UID (links your rights)	int.cb/249a9-8bdd-66b6-144a-c28c764d1818	
Group Membership	default	
This account does not expire	<input checked="" type="checkbox"/>	
Change Internal Password	New	Confirm
	<input type="checkbox"/> Prevent user from changing?	
Language	English	
Disable Account?	<input type="checkbox"/> Date:	2008 Aug 15

Core Settings

- Authentication method
 - SOAP (for GroupWise users)
 - GroupWise users logging in are authenticated using SOAP.
 - These 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)
 - 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.
 - If “exclusive” is not checked, it will try one, then the other.
- Primary UID
 - *Offline* only accounts usually won't have one.
 - 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.
- Normally, accounts expire after 30 days or whatever you set 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.

User Rights



When an administrator level right is granted a user, that user will see that right in the management console 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.

User-Level 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.
- 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.

Note: Only users with administrative rights will see the administrator's screen on login. Non-admin users are simply forwarded to the Search Interface.

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.

Administrator level rights	
<input checked="" type="checkbox"/>	Access all audit logs
<input checked="" type="checkbox"/>	Deletion Manager
<input checked="" type="checkbox"/>	Apply or remove litigation hold
<input checked="" type="checkbox"/>	Manage Server
<input checked="" type="checkbox"/>	Manage Users and Groups
<input checked="" type="checkbox"/>	Manage Workers, Schedules, Profiles, Jobs
<input checked="" type="checkbox"/>	Publish Messages
<input checked="" type="checkbox"/>	Restore messages [Any Mailbox]
<input checked="" type="checkbox"/>	Search all mailboxes
User level rights	
<input checked="" type="checkbox"/>	View/Save attachments
<input checked="" type="checkbox"/>	View personal audit log
<input checked="" type="checkbox"/>	Export Messages
<input checked="" type="checkbox"/>	Forward Messages
<input checked="" type="checkbox"/>	Print Messages
<input checked="" type="checkbox"/>	Read Configuration (Redline)
<input checked="" type="checkbox"/>	Restore messages [My Mailbox]

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

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.

The screenshot shows the 'Mailboxes' configuration page. At the top, there are tabs for 'Core Settings', 'User Rights', 'Mailboxes', and 'Miscellaneous'. The 'Mailboxes' tab is selected. Below the tabs, there is a section titled 'Mailboxes' with a descriptive paragraph: 'These are additional mailboxes beyond the primary mailbox that the user may access. A non-admin user will only view messages addressed to these e-mail addresses (and to their primary mailbox), and might be able to forward/delete messages if given rights. An admin with Search All Mailboxes rights can access all mailboxes and perform any operation in the Search Interface.' Below this, there is a section titled 'Currently Selected' which contains a list of two mailboxes: 'David Bassett' and 'Benjamin Rush', each with a red 'X' icon next to it. Below the 'Currently Selected' section, there is a 'Criteria' section with a search form. The search form has a dropdown for 'First Name', a dropdown for 'begins with', a text input field, and a 'Search' button. Below the search form, there are checkboxes for 'Show only users' (checked) and 'Show only recently cached items' (unchecked), and an 'Add Selected' button. At the bottom, there is an 'Address Book' section with a message: 'No items found. Please try another search.'

Often, you will want some users to be able to search through more than just their own mailbox. Administrators 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.

Address Book Selector

This interface is utilized in various other areas, but is described here.

Currently Selected

- David Bassett
- Benjamin Rush

Criteria

First Name begins with

System Show only users ☒ Show only recently cached items ☐

Address Book

	First	Last	Mail Server	Mailbox	Dept.	Type	Domain
<input type="checkbox"/>		admin	po	admin		U	dom
<input type="checkbox"/>	David	Bassett	po	David		U	dom
<input type="checkbox"/>	Samuel	Clements	po	Sam		U	dom
<input type="checkbox"/>	Benjamin	Rush	po	Benjamin		U	dom

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. The user can then page back and forth among the first 5 pages of results.
- 5) Choose which of the results you want to add to the selected list
- 6) Click **Add Selected Items**

Notes: 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](#).

Core Settings User Rights Mailboxes 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-Year

Time 24 Hour

Example: 02-Aug-2012 17:20

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

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.
- **Message Age Display** Default date filter for searching. May be changed on the fly.
- **Session Timeout** Inactive session timeout. Can be between 5 and 60 minutes.

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.

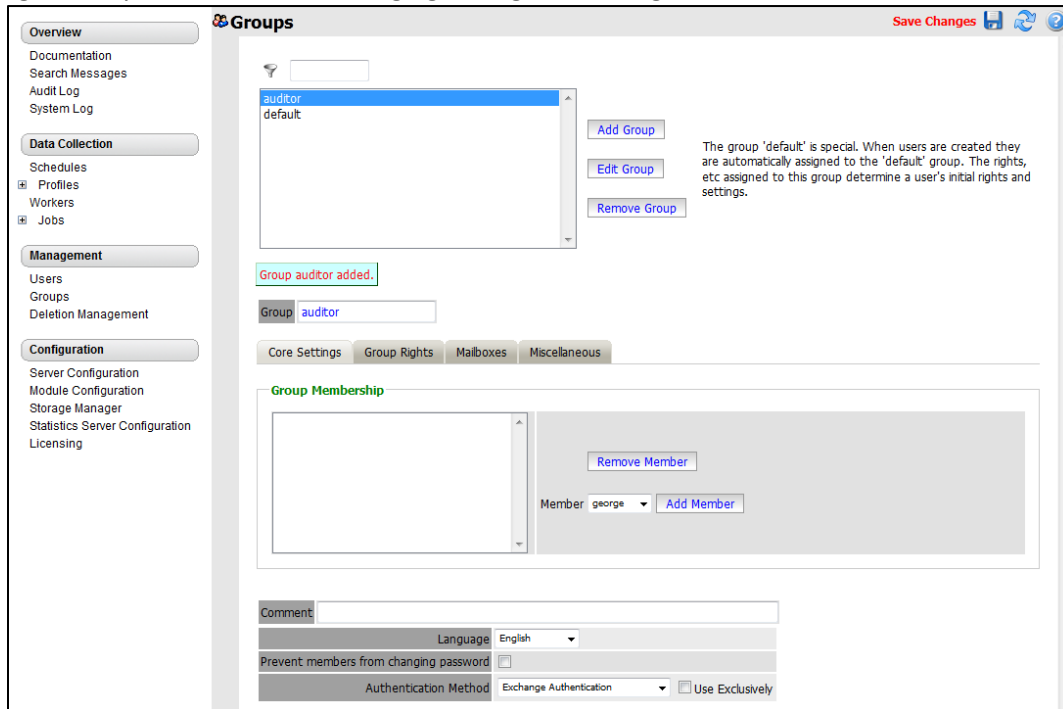
Retain checks these proxy rights the first time you access a proxy users mailbox, then caches the information for 7 days as configured in the server Configuration – Miscellaneous tab. (Default is 7 days.)

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.

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 explicitly 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.

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

The screenshot shows the 'Group Rights' tab selected in a navigation bar with other tabs: 'Core Settings', 'Group Rights', 'Mailboxes', and 'Miscellaneous'. Below the tab bar, the title 'Group Rights' is displayed in green. A descriptive text states: 'Rights granted to the group. Members of the group will inherit these rights, as well as ones explicitly assigned to them.' The configuration is divided into two sections: 'Administrator level rights' and 'User level rights'. Each section contains a list of rights with checkboxes. In the 'Administrator level rights' section, all checkboxes are checked. In the 'User level rights' section, some checkboxes are checked while others are unchecked.

Category	Right	Checked
Administrator level rights	Access all audit logs	Yes
	Deletion Manager	Yes
	Apply or remove litigation hold	Yes
	Manage Server	Yes
	Manage Users and Groups	Yes
	Manage Workers, Schedules, Profiles, Jobs	Yes
	Publish Messages	Yes
	Restore messages [Any Mailbox]	Yes
	Search all mailboxes	Yes
User level rights	View/Save attachments	Yes
	View personal audit log	No
	Export Messages	No
	Forward Messages	Yes
	Print Messages	Yes
	Read Configuration (Redline)	No
	Restore messages [My Mailbox]	Yes

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 function 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.

TIPS

Typically, you will not want a Redline™ group but in case you have more than one Redline control center and you want to see which one is monitoring, you can assign Redline rights to a group.

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.

"Redline" is a trade mark of GWAVA Inc., 100 Alexis Nihon, Suite 500, Saint Laurent, Quebec H4M 2P1

Mailboxes

The screenshot shows the 'Mailboxes' configuration screen. At the top, there are four tabs: 'Core Settings', 'Group Rights', 'Mailboxes' (which is selected), and 'Miscellaneous'. Below the tabs, the 'Mailboxes' section is titled in green. It contains a paragraph explaining that these are additional mailboxes beyond the primary mailbox that group members may access. A non-admin user will only view messages addressed to these e-mail addresses (and to their primary mailbox), and might be able to forward/delete messages if given rights. An admin with 'Search All Mailboxes' rights can access all mailboxes and perform any operation in the Search Interface.

Below the paragraph is a section titled 'Currently Selected' in green, which is currently empty. Below that is a 'Criteria' section. It contains a dropdown menu for 'First Name', a dropdown for 'begins with', and a text input field. To the right of these is a 'Search' button. Below the 'Search' button are two checkboxes: 'Show only users' (checked) and 'Show only recently cached items' (unchecked). Below these checkboxes is an 'Add Selected' button. Below the 'Criteria' section is an 'Address Book' section. It contains a text input field with the placeholder text 'Enter criteria to search the cached addressbook.'

This screen works exactly like the user's access to mailboxes. [Please see that section](#) on 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.

Groups – Miscellaneous

Core Settings

Group Rights

Mailboxes

Miscellaneous

Comment

The default comment appended to forwarded messages can be set here.

Comment

Date/Time Formats

Defines how dates and times are displayed.

Date

Day-ShortMonth-Year

Time

24 Hour

Example: 15-Mar-2011 23:45

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

Session Timeout

Time in minutes, to expire an inactive login session.

Time (minutes)

10

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](#).

Deletion Management

The Deletion Manager provides for the removal of mail from the archive according to the specified criteria. The Deletion management runs as a scheduled job in the archive, looking for, and processing or deleting messages that match the search terms. Mail removed from the archive is permanently deleted. Use this option with care. The Deletion Manager will not show up in your system menu if you have not granted the logged-in user the Deletion Management right, or have the litigation hold right. See [User Rights](#). The Litigation Hold right allows users to go to the deletion management section and add or remove the Litigation Hold right for other users. They cannot modify other settings. Users with the deletion management right can view the litigation hold tab, but they cannot grant rights; it is read only.

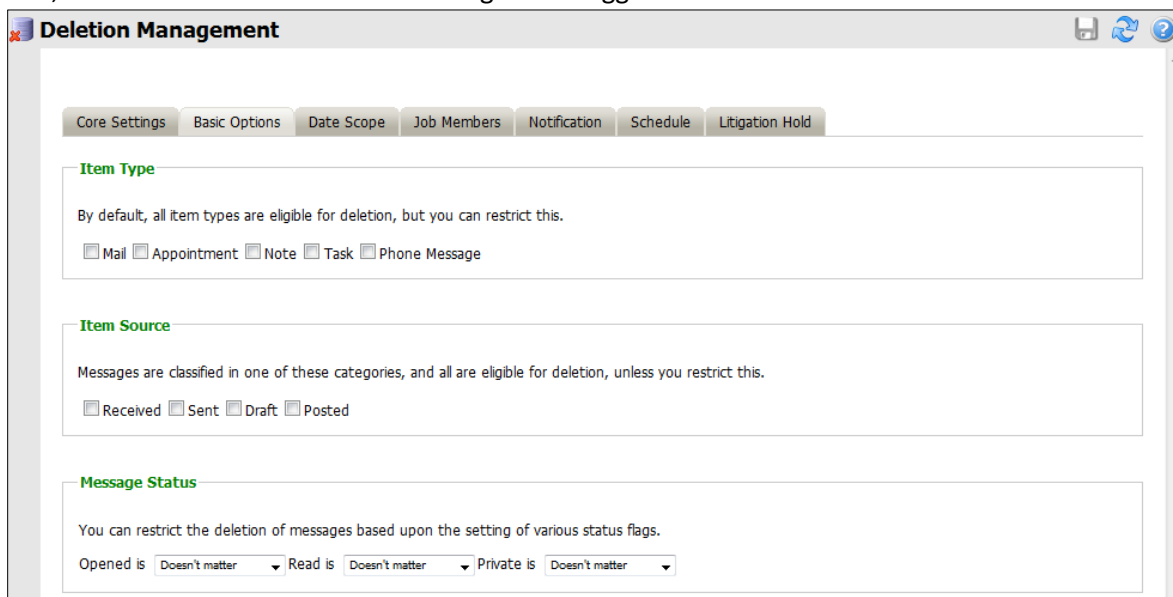
Core Settings

Here you enable and disable deletion jobs, and detail which actions they will take. When setting up a deletion job, you have the option to tell the job delete and report on the messages deleted, or to simply generate a report on the mail that will be removed from the database. The report function can be very handy to ensure that a job will not remove any mail that is needed, before the messages are actually deleted.




Basic Options

This tab provides the criteria that the deletion job will use to identify messages to be deleted. This should look nearly identical to the profile of an archive job. The functions are the same. The item type, source, and status determine which messages are flagged for deletion.



Data Scope

There are many dates that are contained in a mail system, and the deletion manager allows you to select different date ranges to identify the scope of the deletion manager. The setup is simple; the date range between the “Begin” and “End” dates will be targeted by the delete job.



Core Settings Basic Options **Date Scope** Job Members Notification Schedule Litigation Hold

Date Scope

Here the date range of messages to be deleted is defined. It's important to choose the correct range, and the correct date type. The creation and delivered date are internal dates of the message. The stored date is when the message was stored in Archive. The expiration date is related to the expiration settings active in the Job when the message was stored.

Delete messages where **Date Stored in Archive** is:

☒ Older than **Days**

☐ Between two specific dates

Start Date: Time:

End Date: Time:

The dates can be identified by the date filter. The Date filter allows you to specify the mail system or Retain message dates. The creation and delivered date are mail system dates. The date archived and expiration dates are set in Retain. The expiration date is tied to the job, and is set under the job section.

Date Stored in Retain
Delivered Date
Expiration Date
Created Date

The Job Expiration option allows you to set an ‘expiration date’ that the mail no longer needs to be archived after. (Different States have different laws and requirements. Usually it is between 5 and 10 years.) The Deletion Management interface can utilize this expiration date to identify messages that are due for removal.

The screenshot shows the 'Job Members' tab with the following options:

- ☐ Exclude Journaling Mailboxes
- Include these objects:
 - BESX (blackberry.gwava.com) [Add Mail Server]
 - [Add User]
- Exclude these objects:
 - [Add User]

Select the options as desired.

Schedule

The last tab is the schedule tab. This allows you to automate and run a deletion job automatically on mail that has passed its required archive duration. The options are to run this weekly or on a specific day of the month. The Deletion job is set to run at the same time as the rest of your scheduled maintenance. (Scheduled Maintenance is found under Server Configuration – [Maintenance](#) tab. Manual starting of a job is not currently supported.)

Job Members

A delete job will only be active for selected users or a selected mail server. The Job Members tab allows you to include an entire mail server or group of users, while excluding specific users from the delete job.

Use this in conjunction with the Generate Report option under Core Settings to pinpoint the mail that will be included in the delete job.

Notification

The reports, errors, and summaries of delete jobs can be sent to the listed address in the notification tab. Using the Generate Report option in Core Settings is useless unless you set this variable.

The screenshot shows the 'Notification' tab with the following settings:

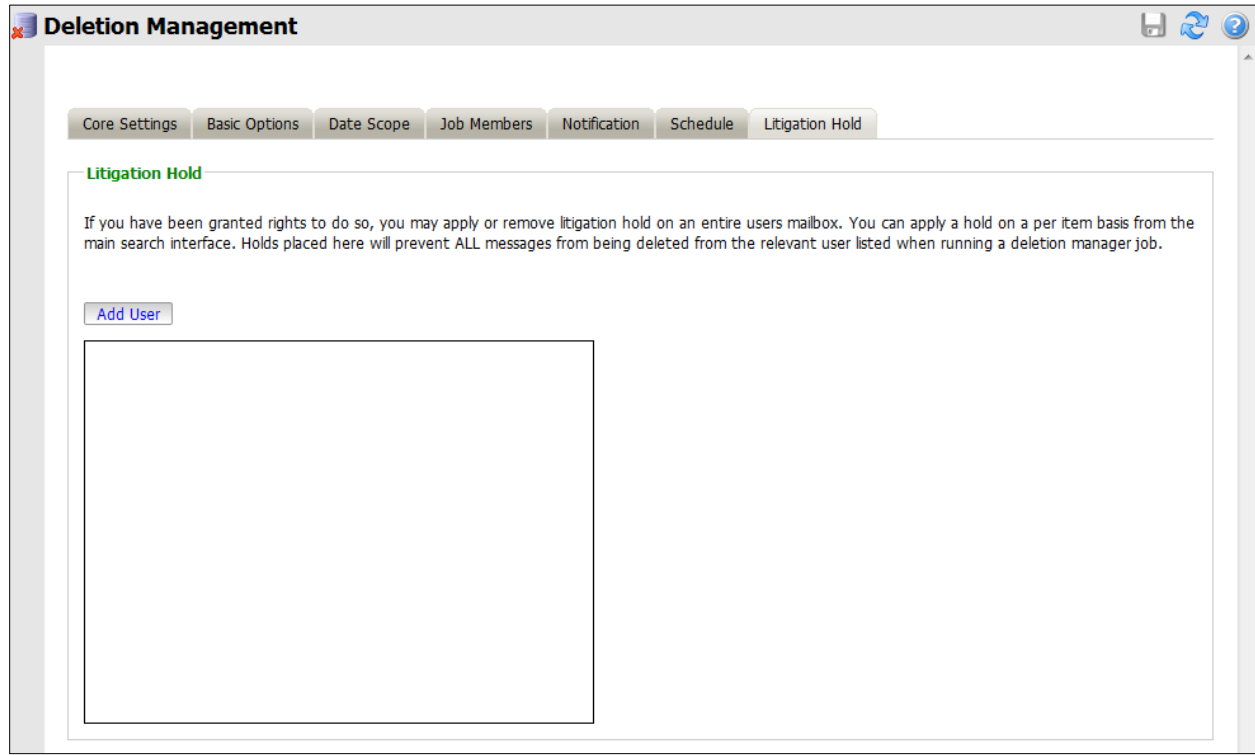
- SMTP Mail Server: 192.168.1.121
- SMTP Mail From Address: admin@company.com
- SMTP To Address: admin@company.com
- SMTP Username: sam
- SMTP Password: (Password is default)
- Mail when errors occur: ☐
- Mail summary when job complete: ☐
- [Test Connection]

The screenshot shows the 'Schedule' tab with the following options:

- Run Job when: weekly [Sunday]

Litigation Hold


The Litigation Hold tab grants the rights to users who will be able to set a hold on messages in the message archive.



Any official auditors, legal representatives, system administrators, or user may be added to this list. These accounts will be able to set and lift any legislative hold in the system, and therefore this is not a generally granted right and should be restricted to only specified users. Because of the power of this right it is granted separately from the usual rights for users.

To add a user to the legislative 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.



Configuration

Server Configuration

The overall operation of Retain is configured in this section.

Server Configuration Save Changes

Communications Index Storage Accounts Maintenance Logging Miscellaneous

Database Connectivity

The message store data is stored on a database server.

To successfully create the tables, the database server must already be running, with a database named archive created. Be sure to specify:

- The correct JDBC URL. If the database server is not running on the same machine as Archive, the IP address or host name of the machine must be included in the URL.
- You must provide a user and password for the database server which has full rights to the database you created.

	Message Store DB	Configuration DB
Database Driver	MySQL 5.0.5.1.5.5	MySQL 5.0.5.1.5.5
Database URL	jdbc:mysql://localhost/archivedb	jdbc:mysql://localhost/archivedb
Database User	archiveuser	archiveuser
Database Password (Password is set) (Password is set)

Notification

Archive uses these values for forwarding and for job notifications. If the mail server is not set, forwarding is disabled

SMTP Mail Server	192.168.1.121
SMTP Mail From Address	admin@company.com
SMTP To Address	admin@company.com
SMTP Username	sam
SMTP Password (Password is set)
Mail when errors occur	<input type="checkbox"/>

[Test Connection](#)

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.

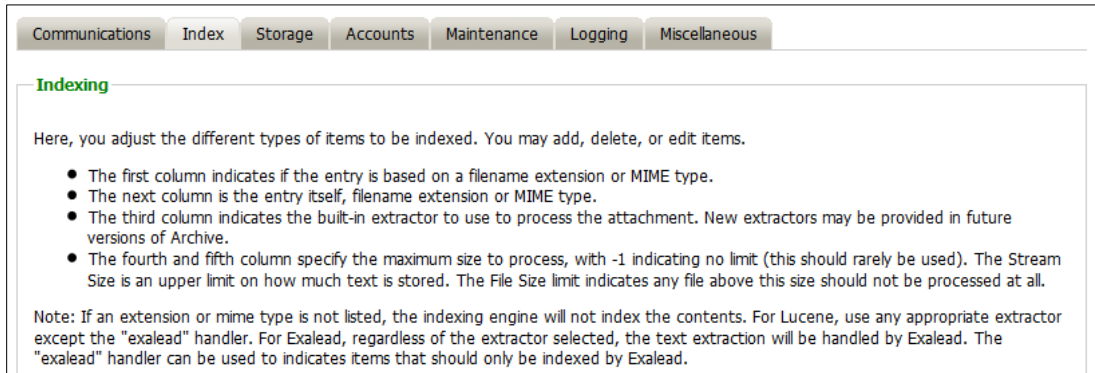
After changing any communications options, it is strongly recommended that Tomcat be restarted immediately.

Communications

The settings for the Communications tab are set in the initial setup of Retain. These include the connection to the SQL Database server, the Retain database, and the notification, SMTP information. If any changes in the system or corrections are needed, they should be configured here.

Index

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.



Because Exalead is a much more robust indexing engine it requires its own server and resources. As such, when Exalead is selected as the indexing engine, a connection address and starting base port are required. The default BASEPORT is 10000. To ensure that the connection to the Exalead server is working, the 'Test Connection' button may be selected, which triggers Retain to contact the Exalead server. The results should shortly appear as a small notification window in your browser.

Lucene is hosted locally on the same machine as the Retain Server, and requires no further configuration, but does not have the same options or the extent of the capabilities as the Exalead engine does.



The indexing engine can be changed between the two engines, but requires the index to be re-created. Recreating the indexes is a time consuming process and should not be done unless required. Searches of the Retain Archive during the index re-creation or migration process may not contain all results.

For details on how to change from one index to another, see the appendix section on [switching the index engine](#).

You can control what Retain indexes here. You may add as many items as you wish to the list of attachment 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. You 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. Retain supports HTML, RTF, TEXT, XML, OpenXML – (MS Office 2007 .docx), OpenOffice2, Word Perfect documents, Excel files, .DOC, and .PDF under the Lucene indexing engine, while Retain supports many more under Exalead. (Because of high CPU, memory, and performance requirements, MS Word and Adobe PDF are not indexed by default and must be enabled to be indexed. If you need to index these items, the allotted memory should be increased. Indexing these items will slow down the indexing process.) Select as many as you need. If an attachment type common in the system which needs to be indexed, but which does not already exist in the system, it may be added by using the 'add' row.

Force Indexing tells the server to index items that are not currently indexed. This queries the system for the top 500,000 items that are not currently indexed, and starts the indexer working if it is not currently working.

Indexing

Here, you adjust the different types of items to be indexed. You may add, delete, or edit items.

- The first column indicates if the entry is based on a filename extension or MIME type.
- The next column is the entry itself, filename extension or MIME type.
- The third column indicates the built-in extractor to use to process the attachment. New extractors may be provided in future versions of Retain.
- The fourth and fifth column specify the maximum size to process, with -1 indicating no limit (this should rarely be used). The Stream Size is an upper limit on how much text is stored. The File Size limit indicates any file above this size should not be processed at all.

Note: If an extension or mime type is not listed, the indexing engine will not index the contents. For Lucene, use any appropriate extractor except the "exalead" handler. For Exalead, regardless of the extractor selected, the text extraction will be handled by Exalead. The "exalead" handler can be used to indicates items that should only be indexed by Exalead.

<input checked="" type="checkbox"/>	Extension	Item:	exalead	Stream Size:	File Size:	Add
<input checked="" type="checkbox"/>	Extension	Item: log	text	Stream Size: 1048576	File Size: 1048576	Delete
<input checked="" type="checkbox"/>	Extension	Item: eml	text	Stream Size: 1048576	File Size: 1048576	Delete
<input checked="" type="checkbox"/>	Extension	Item: docx	openxml	Stream Size: 2048576	File Size: 2048576	Delete
<input checked="" type="checkbox"/>	Extension	Item: html	html	Stream Size: 1048576	File Size: 1048576	Delete
<input checked="" type="checkbox"/>	Extension	Item: cfg	text	Stream Size: 1048576	File Size: 1048576	Delete
<input checked="" type="checkbox"/>	Extension	Item: xml	xml	Stream Size: 1048576	File Size: 1048576	Delete
<input checked="" type="checkbox"/>	Extension	Item: rtf	rtf	Stream Size: 1048576	File Size: 1048576	Delete
<input checked="" type="checkbox"/>	Extension	Item: odt	ooffice	Stream Size: 2048576	File Size: 2048576	Delete
<input checked="" type="checkbox"/>	Extension	Item: txt	text	Stream Size: 1048576	File Size: 1048576	Delete
<input checked="" type="checkbox"/>	Extension	Item: odp	ooffice	Stream Size: 2048576	File Size: 2048576	Delete
<input checked="" type="checkbox"/>	Extension	Item: ods	ooffice	Stream Size: 2048576	File Size: 2048576	Delete
<input type="checkbox"/>	Extension	Item: doc	word	Stream Size: 2048576	File Size: 2048576	Delete
<input type="checkbox"/>	Extension	Item: pdf	exalead	Stream Size: 2048576	File Size: 2048576	Delete

Force Indexing Now

Storage

Communications Index **Storage** Accounts Maintenance Logging Miscellaneous

Storage

The configuration database, indices, XML files, and the archive files are stored under the storage path which should be on a highly reliable storage device with plenty of disk space.

The standard storage engine places archive files under the storage path. You may change the storage path on this screen; the storage engine can be changed from the Storage Configuration menu.

Storage Engine

Base Storage path

[Advanced Settings](#)

Encryption

Compression

The default storage path for your Retain archives is listed here. (See [About Storage Paths](#) and [Storage manager](#) sections.)

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 subdirectories). 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.

Accounts

Communications Index Storage **Accounts** Maintenance Logging Miscellaneous

Account Management

You may expire unused accounts, disable new accounts from being created automatically, and restrict some users from logging in.

Expire unused accounts after how many days (0=never)

Disable new accounts ☐

Prohibited Logins

[Remove selected address](#)

Address [Add](#)

You may turn on a password strength policy, if you want to enforce strong passwords in setting user passwords. Keep in mind, these passwords may be overridden by individual messaging systems.

Password strength:

Will accept any password

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:

Low

Must be between 5 and 15 characters in length, with at least 0 lower case characters, at least 0 upper case characters, at least 0 numerical characters, and at least 0 special characters.

Medium

Must be between 5 and 20 characters in length, with at least 1 lower case characters, at least 1 upper case characters, at least 1 numerical characters, and at least 0 special 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. . In addition the password will be checked against a dictionary.

Intruder Lockout

The screenshot shows a web interface titled "Intruder Lockout". It contains a text box explaining the feature: "Enable this feature to temporarily lockout an account that may be the subject of an illegal entry attempt. You may clear lock outs using the button below, or by restarting Tomcat". Below this are four configuration options, each with a label and a dropdown menu: "Enable Intruder Lockout?" with a checkbox, "Number of invalid login attempts" set to 3, "Time interval (minutes)" set to 5, and "If triggered, lock account for this period (minutes)" set to 10. At the bottom left is a button labeled "Clear lock outs".

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.

Maintenance

Maintaining a Retain system involves many tasks:

- Backup Configuration (Embedded) Database
- Compress Configuration Database
- indexing (All Indexes)

These items are all controlled here.

The screenshot shows the 'Maintenance' tab in the GWAVA Retain configuration interface. The 'Maintenance/Logging' section is active, displaying instructions and configuration options for database backup and index optimization. Below this, the 'Ignore Domains' section is visible, showing a list of ignored domains and an 'Add' button for new domains.

Maintenance/Logging

You may automatically backup and/or compress the embedded database at a specified time. The indices are also backed up. The Retain Server is offline at this time, so no jobs may be active at this point.

Note that with the Exalead engine, backup and optimization of the indices will always be disabled.

Enable Backup of Indices	<input checked="" type="checkbox"/>	(every day)	weekly
Enable Index Optimization	<input checked="" type="checkbox"/>	(every day)	weekly
Purge exported files when older than		12	hours
Run maintain procedure at		01	: 00
Last Backup of Indexes	12-Mar-2011 01:02:00		
Last Index Optimization	12-Mar-2011 01:02:00		

Ignore Domains

You might want specific Domains (and their child POs and users) to be completely ignored - not cached, not stored, etc.

Ignored Domains

Remove Selected Domain

Domain

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.

You can also configure how often indexes get backed up and how often the indexes are optimized for speed.

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. (see [index/storage](#) for location)

Ignore Domains

Ignore Domains

You might want specific Domains (and their child POs and users) to be completely ignored - not cached, not stored, etc. The reason? GroupWise has the capability to add non-GroupWise (external) domains to the system. This is commonly used, for example, to add internet addresses to the System Address Book or to Distribution Lists. However these objects cause problems with Retain, since it obviously cannot log into them.

Ignored Domains

Domain

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.

Logging

Configure the system logging here. Normally, you want verbose logging. Diagnostic is ordinarily done just for troubleshooting purposes.

The screenshot shows the 'Logging' configuration page. At the top, there is a navigation bar with tabs: Communications, Index, Storage, Accounts, Maintenance, Notification, Logging (selected), and Miscellaneous. Below the navigation bar, the 'Logging' section is titled in green. It contains a table with the following settings:

Logging Enabled?	<input checked="" type="checkbox"/>
Log level	Normal
Compress logs?	<input checked="" type="checkbox"/>
Remove old logs	<input checked="" type="checkbox"/>
When age exceeds (days)	10

You also specify how long to keep logs here. By default, logs are deleted after 10 days. Logs may be compressed to save disk space.

Auditing Records

Retain now will create an auditing record of all actions, specified by the user, which are taken on a specific 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 Configuration page, Logging tab.

The screenshot shows the 'Auditing' configuration page. At the top, there is a navigation bar with tabs: Communications, Index, Storage, Accounts, Maintenance, Notification, Logging, and Auditing (selected). Below the navigation bar, the 'Auditing' section is titled in green. It contains the following text:

Retain audits all access to messages by default, allowing you to trace when and who opened a message, forwarded a message etc.

Some advanced auditing options are disabled by default, because they would generate enormous audit logs. Enable them with care. Also, consider clearing your audit logs regularly (see below) or exporting your logs from the audit table on your SQL Server.

Days before audit records are removed (0=never remove)

Advanced auditing options

<input type="checkbox"/> Login	<input type="checkbox"/> Changed User/Group	
<input type="checkbox"/> Changed Data Collection	<input type="checkbox"/> Searched	<input type="checkbox"/> Mailbox Switch
<input type="checkbox"/> Audited	<input type="checkbox"/> Server Config Change	<input type="checkbox"/> Exported
<input type="checkbox"/> Published		

This logging option creates very detailed logs for the options selected, and it is very important to know that if every option is turned on, the logs can become extremely large. 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.

Miscellaneous

Communications	Index	Storage	Accounts	Maintenance	Logging	Miscellaneous
----------------	-------	---------	----------	-------------	---------	---------------

Miscellaneous

You may change the default Worker password, which is assigned to all Workers when they are created (but can be individually overridden).

You can enable or disable XML Export, which creates XML representations of each message as they are created. Generally this is not recommended unless you absolutely need this feature - it uses a lot of disk space, and slows down the Server.

Default Worker Password	(Password is set)
Export XML	<input type="checkbox"/>	
Server Friendly Name	Server	

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.

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.

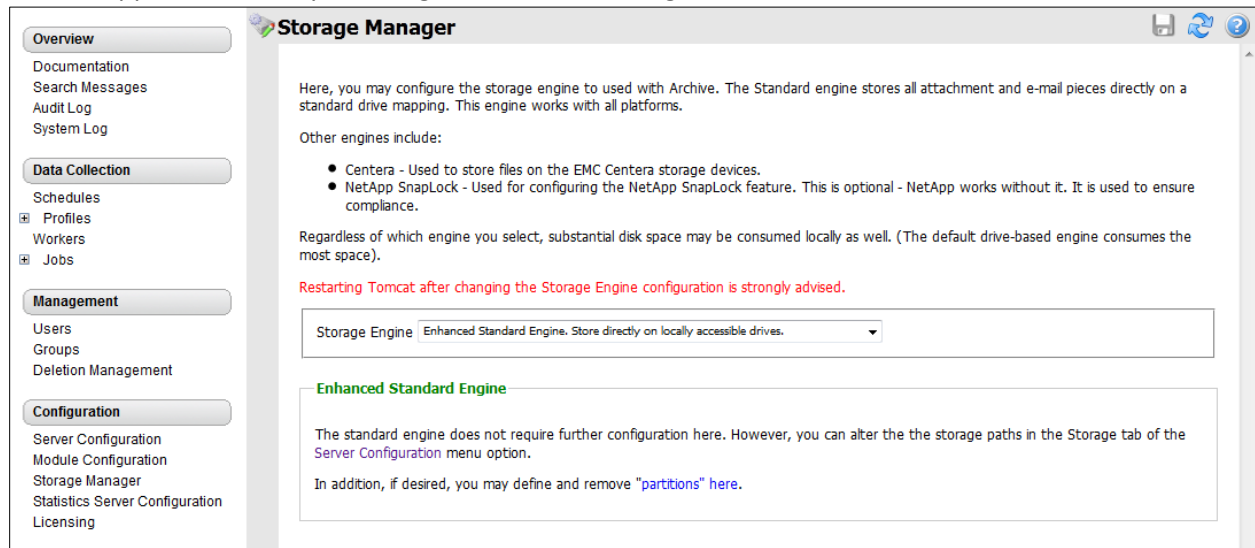
When you do NOT use the XML export function, you will benefit from Retain's single-instance message storage and data compression to save disk space and improve performance.

Module Configuration

Please see the [module configuration](#) section in the initial setup section of the manual.

Storage Manager / Storage Engine

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 three options are:

Legacy Engine
Centera Storage Device
NetApp SnapLock



Enhanced Standard Engine

The Enhanced 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.

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.

Centera

Server Connectivity Settings

Note: The hostname/ip may be suffixed with a :port entry, and you may enter more than one hostname. For example, mycenterbox.com:4256,152.62.65.11,152.62.65.12?name=2profile3,secret=2profile3 is a valid entry. At least one hostname must be entered.

Centera Server/IP

Server Access settings

With Centera devices, you have three options for authentication:

- Anonymous, which requires no password, but is insecure.
- Credentialed, in which you supply a traditional username and password.
- PEA, in which you supply a file system path accessible by Retain of the PEA file you encrypted to provide authentication information. A good place to put this is in the license directory, (/var/opt/beginfinite/retain/license).

Authentication Method Anonymous

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.

Authentication Method PEA

PEA Path (full path)

Authentication Method Credentialed

User

Password

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, username, password, and the correct path to the server. THE PATH MUST BE SET IN NETAPP FORMAT.

NetApp SnapLock

SnapLock settings

If you enable SnapLock, then as items are stored on the NetApp appliance with a Job specifying an expiry date, their retention date is set according to your Job options. This effectively enforces WORM rules. You may use NetApp without this setting, if preferred, in which case none of the parameters on this screen need to be filled out.

A note about the path requested below: it is not a Linux, Windows, or UNC path. There is no definite correlation for example between this path and a share path. It uses the NetApp format, and must point to the directory containing all of the Retain blob files. Consult your NetApp Technical Support for more information on this or on enabling SnapLock.

Enable SnapLock?	No <input type="button" value="v"/>
NetApp Server/IP	<input type="text"/>
Port	80 <input type="text"/>
Admin User	<input type="text"/>
Admin Password	<input type="text"/>
Base path to archive directory (example: /vol/myvolname/archive)	<input type="text"/>

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 storage between years. (I.e. /archive/2008; /archive/2009)

Partitions DS

Creating a partition allows you to direct all new attachment blobs to a new storage device. (Your old device might be running out of space, or you might prefer to split the storage on a yearly basis etc.). Retain retrieves and stores blobs based upon the partition's date threshold.

Deleting or adding a partition can easily result in the loss of data. For example, deleting a partition that contains data, will make all of that data completely inaccessible to Retain, unless it is moved to another partition. Carefully read the user manual and/or consult with Technical Support before proceeding.

Current path used for storage: /opt/archive/2012

Begin Date	Name	Path	Delete
02-Aug-2012 22:25:43	first partition	/opt/archive/2012	Delete

You may add a partition below. Just enter a name and the full path to the directory to store blobs under.

New Partition Name

New Partition Path

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 partition 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 support.**

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.

Partitions

Creating a partition allows you to direct all new attachment blobs to a new storage device. (Your old device might be running out of space, or you might prefer to split the storage on a yearly basis etc.). Retain retrieves and stores blobs based upon the partition's date threshold.

Deleting or adding a partition can easily result in the loss of data. For example, deleting a partition that contains data, will make all of that data completely inaccessible to Retain, unless it is moved to another partition. Carefully read the user manual and/or consult with Technical Support before proceeding.

Current path used for storage: /var/opt/beginfinite/retain/archive

There are no partitions defined currently. All blobs will be stored to the default storage path configured in Server Configuration.

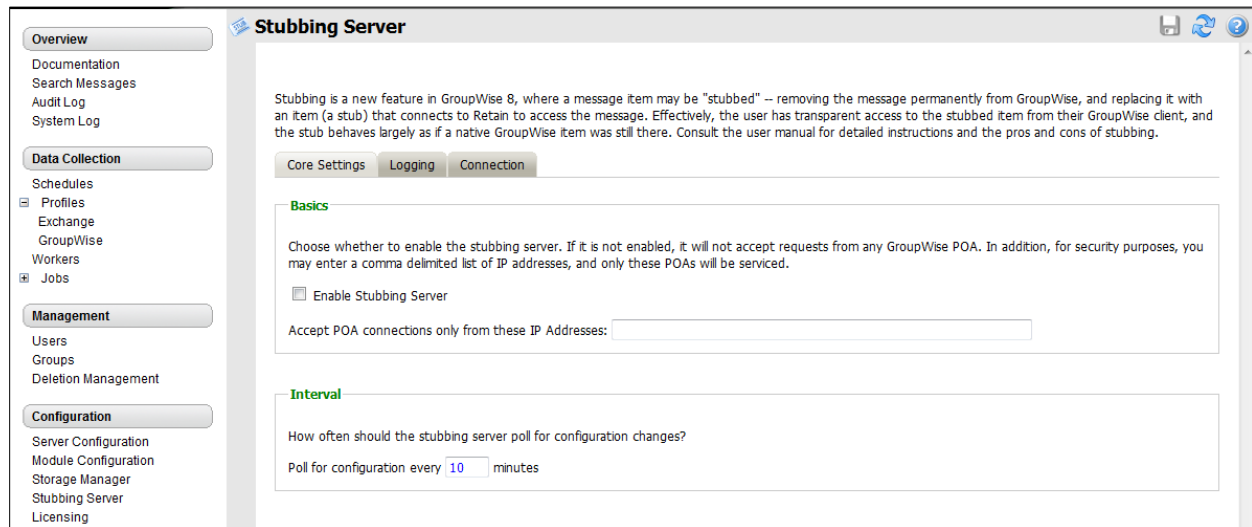
You may add a partition below. Just enter a name and the full path to the directory to store blobs under.

New Partition Name

New Partition Path

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)

Stubbing Server



To set up stubbing, you must perform the following tasks:


1. You **must** be running GroupWise 8.0HP1 Post Office Agent(s). **Once enabled, the stubbing process cannot be reversed. Stubbing requires GroupWise 8.0HP1. Previous versions of GroupWise do not contain critical stability fixes.**
2. ConsoleOne must be configured as discussed below.
3. Retain job profiles must have Stubbing turned on. (see the [Profiles section](#))
4. Stubbing server must be installed and setup as discussed below. (For install see your platform install guide.)

ConsoleOne GroupWise settings

For the Retain Server to function, GroupWise must be told where the Stub Server is located.


In ConsoleOne;

- Connect to the domain and highlight the domain or GroupWise system.
- Select Tools | GroupWise System Operations | Trusted Applications
- Select Retain and click Edit.
- Select "Allow Access to Archive Server"



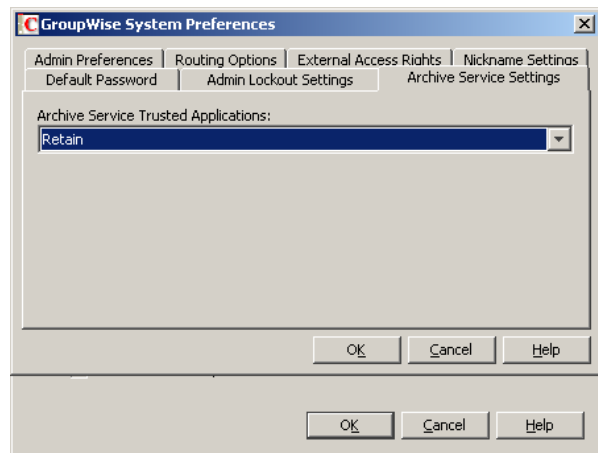
- Put the address or DNS name of the Stub Server and the correct port as shown. (48080)
- Click OK and exit the rest of the windows.

Your settings should look the same as the settings pictured here with your specific DNS or IP address for the Stub Server.



We also need to change the settings in the preferences menu:


- Connect to the domain and highlight the domain or GroupWise system.
- Select Tools | GroupWise System Operations | System Preferences...
- Ensure that Retain is set under the Archive Service Settings as a trusted Application. Select it if it is not already selected.



Set the Post Office Remote File Server Settings.

Enter the properties of the Post office. The settings should look as they do here.

Retain should be listed as the Default Archive Service Trusted Application. If it is not set as the default, override and select Retain.



Retain Stubbing Settings

Job profiles must have Stubbing enabled for the Server to function. See the [Profiles section](#) for details. The Stub Server must first be enabled in order to allow access to the bootstrap file required for communications. Set the interval as desired. (10 min. is lowest recommended setting. Anything lower may be unnecessary.)

The screenshot shows the 'Retain Stubbing Settings' window with tabs for 'Core Settings', 'Logging', 'Connection', and 'Bootstrap'. The 'Basics' section contains the text: 'Choose whether to enable the stubbing server. If it is not enabled, it will not accept requests from any GroupWise POA. In addition, for security purposes, you may enter a comma delimited list of IP addresses, and only these POAs will be serviced.' Below this is a checked checkbox for 'Enable Stubbing Server?' and a text input field for 'Accept POA connections only from these IP Addresses:'. The 'Interval' section contains the text: 'How often should the stubbing server poll for configuration changes?' and a text input field for 'Poll for configuration every 10 minutes'.

If you wish to restrict connections to the Stub Server, input the IP address(es) of all the Post Offices that the Stub Server will service. Use commas to separate multiple addresses. Save changes before continuing.

Logging

Log levels higher than Normal are not necessary unless troubleshooting or working with Support. It is strongly recommended to compress and remove logs after a time period to protect system space. Set time to desired level.

The screenshot shows the 'Retain Stubbing Settings' window with tabs for 'Core Settings', 'Logging', and 'Connection'. The 'Logging Options' section contains a table with the following settings:

Logging Enabled	<input checked="" type="checkbox"/>	
Log level	Normal	▼
Compress logs	<input checked="" type="checkbox"/>	
Remove old logs	<input checked="" type="checkbox"/>	
When age exceeds (days)	10	

Connection

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, change the Host Name to the DNS name or localhost (127.0.0.1).

The screenshot shows the 'Retain Stubbing Settings' window with tabs for 'Core Settings', 'Logging', and 'Connection'. The 'Retain Server Connection' section contains the text: 'The Stubbing Server will need to know how to connect to the Retain Server. It is highly recommended that you use a DNS host name for the Host Name. It's important to note that you are specifying the URL from the point of view of the **Stubbing Server contacting the Retain Server**, which may or may not match the URL you'd use from your browser to reach the Stubbing Server or Retain Server.' Below this is a table with the following settings:

Server Protocol	http ▼
Retain Server Host Name	64.34.151.106
Retain Server Port	48080
Retain Server Path	/RetainServer
Stubbing Server Password	••••••••

Note: Changing any of these parameters will require recopying the stubbing server bootstrap file!

When you have finished configuring the Stub Server, Save your changes and select the Bootstrap tab. (The bootstrap tab will not be available unless the Stubbing Server has been enabled and settings saved.)



Download and save the bootstrap file.


Open a new browser tab or window, and browse to the Stub Server. The Stub Server answers to the following address (The address is case sensitive):

http://<IP address of StubServer>:48080/RetainStubServer


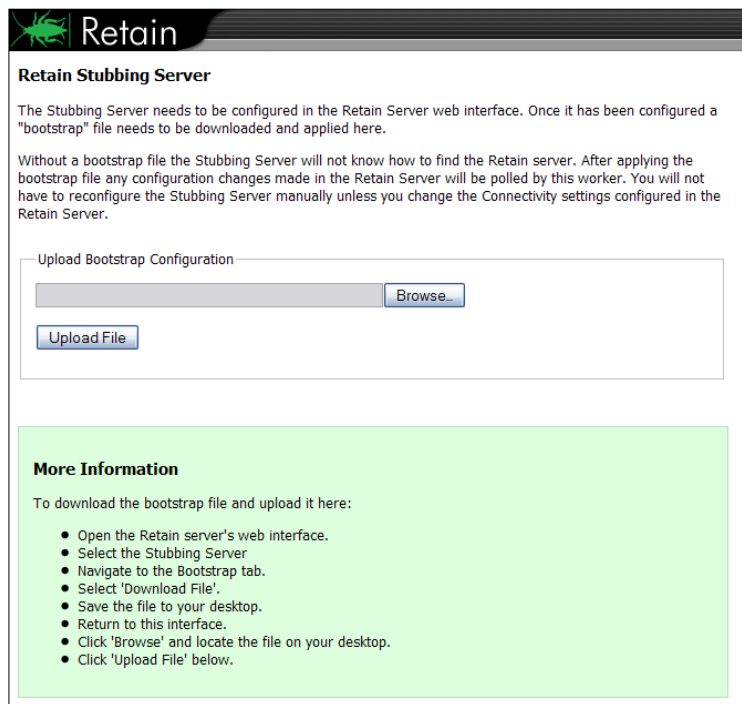
You should see a window like the one pictured here.

- Click 'Browse' and browse to the location where you saved the bootstrap file.
- Select the "StubConfig.cfg" file and select 'Open'.
- Select the 'Upload File' button.

After uploading the configuration file, the Stub Server will alert you that it requires authentication from a user with the "Manage Workers" right. Click 'Continue'.



We can now Login to the Stub Server. You must Login with a user that has sufficient rights. See the [User Rights](#) section to add or remove rights.



The Stub Server Interface is nearly identical to the Worker interface. You may check for updated versions of the Stub Server or manually trigger server polling by selecting the Refresh 'Configuration now' button.

The screenshot shows the 'Retain' web interface with the 'Configuration' tab selected. The 'Stub Server Configuration' section displays the following information:

URI to Server	http://192.168.0.101:48080/RetainServer/
Stub Server Version	1.7.0
Memory Free/Max/Total	233775/1040512/260224

Below the table are two buttons: 'Check for New Version' and 'Refresh Configuration now'. A note below the buttons states: 'Click "Refresh Configuration now" to prompt the Stubbing Server to poll for configuration changes right now.' At the bottom, a paragraph reads: 'Bootstrap information has been applied to this server. If the configuration must be changed in the future, replace the configuration file at /opt/beginfinite/retain/RetainStubServer/WEB-INF/cfg/ and restart the Stubbing Server.'

Check the **Server Connection** tab to verify the connection to the Server.

The screenshot shows the 'Server Connection' tab selected. The 'Server Connection' section displays the message: 'Successfully connected to Retain server' in green text.

Status

The Status tab gives statistics about up-time and function.

The screenshot shows the 'Status' tab selected. The 'Status 4' section displays a table with the following data:

Up since	Fri Nov 21 08:45:19 MST 2008
Successful Logins	0
Failed Logins	0
Message Served	0
Attachments Served	0
Communication Errors (POA)	0
Communication Errors (Retain)	0
Cached Sessions	0
Max Sessions	0
Expired Sessions	0


Scheduler

The scheduler tab contains the configuration details sent by the Server this tab is mainly used for Support.

The screenshot shows the 'Scheduler' tab selected. The 'Scheduler Status' section displays the message: 'To see advanced information regarding the thread scheduler, click [here](#).'

Statistics Server Configuration

Retain comes equipped with a separate statistics monitoring server. The Statistics Server may be installed on the same or a different physical or virtual machine as the Retain Server.



Overview

- Documentation
- Search Messages
- Audit Log
- System Log

Data Collection

- Schedules
- Profiles
- Workers
- Jobs

Management

- Users
- Groups
- Deletion Management

Configuration

- Server Configuration
- Module Configuration
- Storage Manager
- Statistics Server Configuration
- Licensing

Stats Server

The Stats Server, if enabled, can collect various statistics from the Workers and generate useful reports.

Core Settings | Logging | Connection | Notification

Basics

Choose whether to enable the Stats Server. If it is not enabled, statistics information will not be stored for Archive jobs. In addition, for security purposes, you may enter a comma delimited list of IP addresses, and only these Server will be allowed to update the statistics.

Enable Data Collection	<input type="checkbox"/>
Days to archive job information	14
Run maintenance at (hour)	02 :00
Update interval for Server Statistics (minutes)	10

Accept Server connections only from these IP Addresses:

Interval

How often should the stats server poll the Archive Server for configuration changes?

Poll for configuration every 10 minutes

The Statistics Server 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 connection, only server statistic information.

The Core Settings, Logging, Connection and Notification sections of the Statistics Server configuration 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.



Core Settings | Logging | Connection | Notification

Logging Options

Logging Enabled	<input checked="" type="checkbox"/>
Log level	Normal
Compress logs	<input checked="" type="checkbox"/>
Remove old logs	<input checked="" type="checkbox"/>
When age exceeds (days)	10

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.

Core Settings	Logging	Connection	Notification										
<h3>Connection from Stats Server to Archive Server</h3> <p>The Stats Server will need to know how to connect to the Archive Server. It is highly recommended that you use a DNS host name for the Host Name. It's important to note that you are specifying the URL from the point of view of the Stats Server contacting the Archive Server, which may or may not match the URL you'd use from your browser to reach the Stats Server or Archive Server.</p> <table border="1"> <tr> <td>Server Protocol</td> <td>http</td> </tr> <tr> <td>Archive Server Host Name</td> <td>192.168.1.114</td> </tr> <tr> <td>Archive Server Port</td> <td>80</td> </tr> <tr> <td>Archive Server Path</td> <td>/ArchiveServer</td> </tr> <tr> <td>Stats Server Password</td> <td>..... (Password is set)</td> </tr> </table> <p>Note: Changing any of these parameters will require recopying the stats server bootstrap file!</p>				Server Protocol	http	Archive Server Host Name	192.168.1.114	Archive Server Port	80	Archive Server Path	/ArchiveServer	Stats Server Password (Password is set)
Server Protocol	http												
Archive Server Host Name	192.168.1.114												
Archive Server Port	80												
Archive Server Path	/ArchiveServer												
Stats Server Password (Password is set)												
<h3>Connection from Archive Server to Stats Server</h3> <p>The Archive Server will need to know how to connect to the Stats Server. It is highly recommended that you use a DNS host name for the Host Name. It's important to note that you are specifying the URL from the point of view of the Archive Server contacting the Stats Server, which may or may not match the URL you'd use from your browser to reach the Stats Server or Archive Server.</p> <table border="1"> <tr> <td>Server Protocol</td> <td>http</td> </tr> <tr> <td>Archive Server Host Name</td> <td>192.168.1.114</td> </tr> <tr> <td>Archive Server Port</td> <td>80</td> </tr> <tr> <td>Archive Server Path</td> <td>/ArchiveStatsServer</td> </tr> </table>				Server Protocol	http	Archive Server Host Name	192.168.1.114	Archive Server Port	80	Archive Server Path	/ArchiveStatsServer		
Server Protocol	http												
Archive Server Host Name	192.168.1.114												
Archive Server Port	80												
Archive Server Path	/ArchiveStatsServer												

The Connection between the stats server and the <servername> should not be changed unless the Retain Server requires a modified connection path, (such as a custom port.)

The Statistics server is different from the Retain Worker in that the Retain Server is able to contact the statistics server as well. The connection address to the statistics server should not be modified unless custom connection settings have been created.

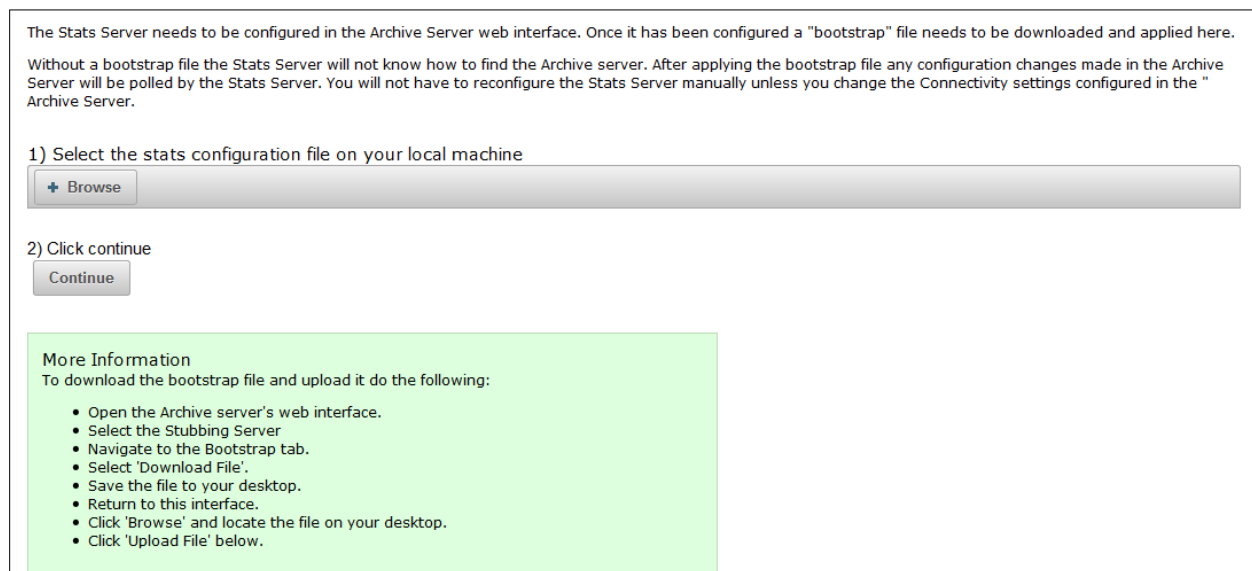
Core Settings	Logging	Connection	Notification										
<h3>Notification</h3> <p>You can have the Stat Server e-mail a summary of operations and/or errors that occur while the Stat Server runs.</p> <table border="1"> <tr> <td>SMTP Mail Server</td> <td>192.168.1.121</td> </tr> <tr> <td>SMTP Mail From Address</td> <td>archive@gwava.com</td> </tr> <tr> <td>SMTP To Address</td> <td>admin@gwava.com</td> </tr> <tr> <td>SMTP Username</td> <td>sam</td> </tr> <tr> <td>SMTP Password</td> <td>..... (Password is default)</td> </tr> </table> <p>Test Connection</p>				SMTP Mail Server	192.168.1.121	SMTP Mail From Address	archive@gwava.com	SMTP To Address	admin@gwava.com	SMTP Username	sam	SMTP Password (Password is default)
SMTP Mail Server	192.168.1.121												
SMTP Mail From Address	archive@gwava.com												
SMTP To Address	admin@gwava.com												
SMTP Username	sam												
SMTP Password (Password is default)												

Notification defaults to the standard notification address specified in the Retain Server activation and initial 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 Statistics 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 Statistics server webpage. The default connection address is shown on the connection tab.

ie. http://<ip_address_of_server/RetainStatsServer



From the Stats Server initial webpage, click on the "Browse" button and browse to the saved bootstrap file and select it.

2) Click continue

[Continue](#)

Click 'Continue' to upload the bootstrap file. As the bootstrap file is uploaded, the Stats Server will require an authenticated login. **Only users with the administrative "manage workers" rights will be able to login to the stat server.**

Login

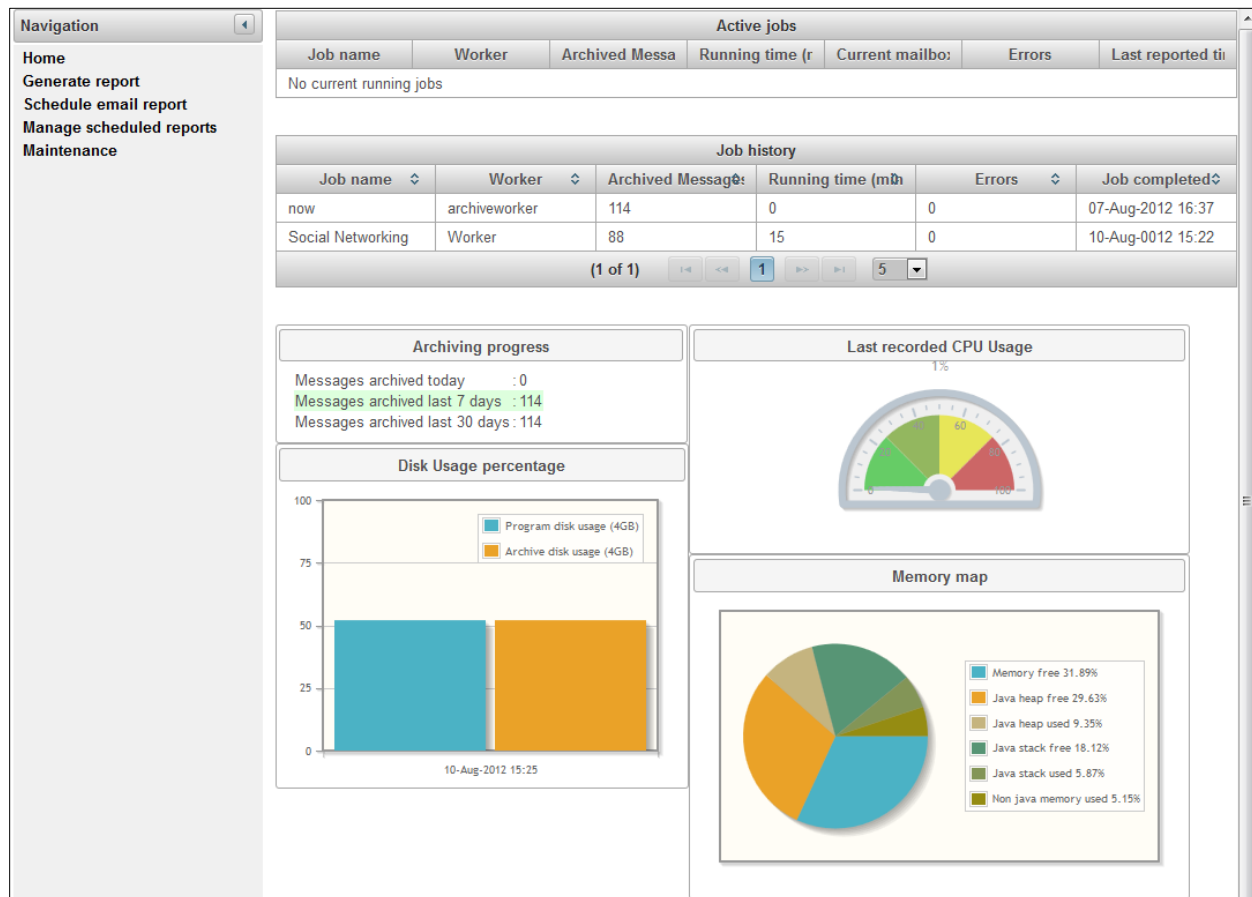
Password

[Login](#)

Home

On login, the Stats Server opens to the Home page. Navigation links to the different pages are listed in the left hand pane, and the navigation window can be collapsed or expanded for screen space.

The Home page displays the most pertinent statistics information: messages archived in last seven and 30 days, Job history, active job list, CPU usage, Disk Usage, and a memory map for the <<server name>>.



Generate report

On the Generate report page, different reports can be generated in one or several jobs.

Jobs can be sorted according to the job name, start time, or end time, and reports can be selected from a sorted list or from a drop-down alphabetical list.

Job name	Start time	End Time
Select a value		
now	07-Aug-2012 16:37	07-Aug-2012 16:37
Social Networking	10-Aug-0012 15:07	10-Aug-0012 15:22

To generate a report, select a report type from the 'Report to Create' dropdown menu. The default report type is 'Job summary'.

- Audit user activity
- Export audit table
- Job run averages
- Job summary**
- Archive summary by user
- Server utilization
- User timestamps
- Job mailbox summary

Also select the Report Output Format desired. The default output type is set to PDF, but HTML, XLS, and CSV format files are also available.

Each report must have a subject. Specify the desired report subject by selecting the appropriate item from the criteria menu. The one displayed is the job summary report, and displays the different jobs that have been run.

- PDF**
- HTML
- XLS
- CSV

Select one of the items, in this example, a job, and then select where the report will be delivered – immediately in the browser or delivered in Email.

- Browser**
- Email

Different reports have different criteria as appropriate. Select or specify the desired criteria for the desired report and then click on the "Generate" button to create the

Generate

report.

Schedule email reports

Reports can also be scheduled for automatic generation. Automatic reports are always delivered through email, and are created on the 'Schedule email report' page. The same jobs and criteria as found on the 'Generate report' page are found here, but the Schedule report page also includes a schedule for the recurring or single run report, and the dates effective for the report. When the report is configured as desired, click the 'Schedule' button to save the job to the job queue.

Navigation

[Home](#)
[Generate report](#)
[Schedule email report](#)
[Manage scheduled reports](#)
[Maintenance](#)

Generate report

Report to Create

Job summary

Report Output Format

PDF

Select the job to report on

Job name	Start time	End Time
Select a value		
now	07-Aug-2012 16:37	07-Aug-2012 16:37
Social Networking	10-Aug-0012 15:07	10-Aug-0012 15:22

Email details

Mail server hostname/IP

192.168.1.121

Sender email address

archive@gwava.com

Recipient email address(es)

admin@gwava.com

Email subject

Archive Report

SMTP username (optional)

sam

SMTP password (optional)

....

Report Schedule

Send Report

Daily

At Time (24hr format)

00 : 00

Schedule Start Date

Schedule End Date

Schedule

Manage scheduled reports

To view the scheduled job queue or to delete or modify existing reports, select the 'manage scheduled reports' page. The 'Manage scheduled reports' page is the only page which allows the removal or modification of any currently created or active report jobs. To remove a job select the job (s) from the display and select the 'Remove selected' button. To open a job for editing, select the 'View/change details' button. To initiate a job immediately, select the 'Run report now' button.

Navigation

[Home](#)
[Generate report](#)
[Schedule email report](#)
[Manage scheduled reports](#)
[Maintenance](#)

Currently Scheduled Jobs

Report name	Output format	Schedule type	Next fire time	Sender email address	Email subject
No records found.					


View/Change details

Run report now

✕ Remove selected

Maintenance

The Stats Server maintains a separate record of the address book and the job cache. The system automatically syncs this information with the Retain Server every 24 hours. Still, the different records may become out of sync. If a sync becomes required, the different records may be individually triggered to refresh immediately.



The screenshot shows a web interface with a 'Navigation' sidebar on the left and a 'Maintenance' panel on the right. The sidebar contains links: Home, Generate report, Schedule email report, Manage scheduled reports, and Maintenance (which is highlighted). The Maintenance panel has a title bar 'Maintenance' and contains the following text: 'Maintenance job(s) will be performed at 2:00 (24hr clock).', 'Currently configured to keep 14 days of statistical data.', and 'To configure these settings or other stats settings please do so via the [Archive Server](#).' Below this text are three buttons: 'Remove old statistical data now' followed by a text input field containing '14' and the word 'days', 'Synchronize address book cache now', and 'Synchronize jobs cache 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 statistics older than 30 days.

Licensing

The Retain license is a file with a PEM extension which should be uploaded through this screen. To acquire a license, contact your [GWAVA sales representative](#).

Licensing

Licensing

✖ Licensed to CN=pg-1384, EMAILADDRESS=info@gwava.com, O=GWAVA, L=Montreal, ST=QC, C=CA
License ID 108f3478-d038-4f15-92a8-ae7e9f6c617c
License State **Valid License file installed**
License Type base exchange gw socialnetworking
Install Date Fri Jul 27 20:50:36 MDT 2012
License expires **Thu May 30 21:59:59 MDT 2013**
Support expires Thu May 30 23:59:59 MDT 2013
Your license file is currently stored at: /var/opt/netboxblue/archive/license/rtn334351597124430935.pem
Make sure you keep a backup of this file.

Submit License

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

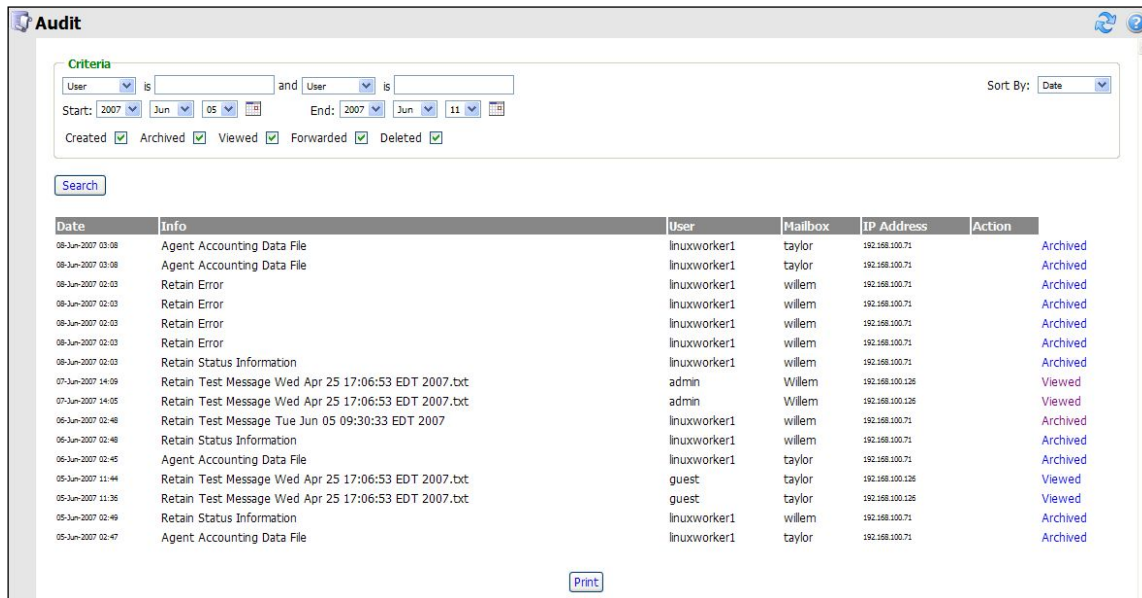
To submit your license, browse to your PEM file and click "Submit License".

Logout and re-login to the server to see the license in effect.

Overview

Audit Log

All operations involving the archives are logged. This is for auditing purposes so you may know who did what when.



- Accessing the audit log requires the "Access Audit Log" administrative right.
- Ordinary users cannot access their audit logs, unless you grant the "Access Personal Audit Log" right, in which case they can find the audit log on a per-message basis under the Properties tab.

To view the system audit trail,

- 1) Enter your search criteria (what part of the log do you want to see?)
- 2) Set a date range
- 3) Click Search.

The results appear below. This may be printed.

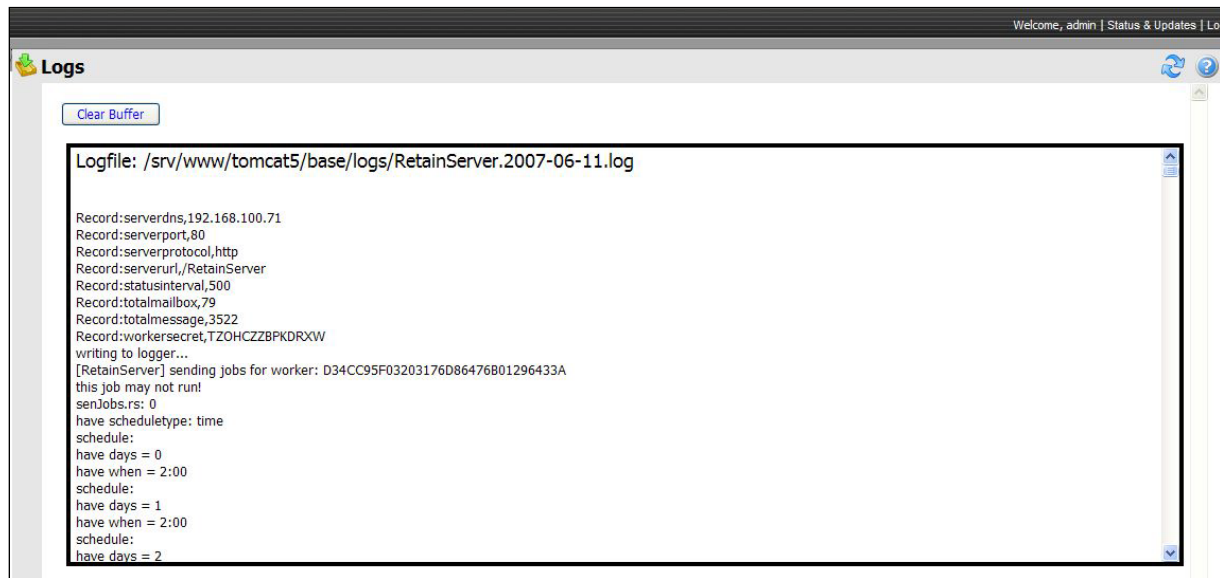
If you have rights, you may click on the Action link to open the message.

An auditor does not automatically get the rights to READ messages in mailboxes that are not assigned to them.

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.



Searching the Retain Archives

Now that there is data in your Retain archives, you will want to retrieve this information. This is the core of the information retrieval system in Retain. You search for what you want, the results appear and you can view, print, export, or forward the results. If the logged on user has at least one administrative right assigned, the user is brought to the administrative menus upon login, and must click on “Search Messages” to access the search interface.

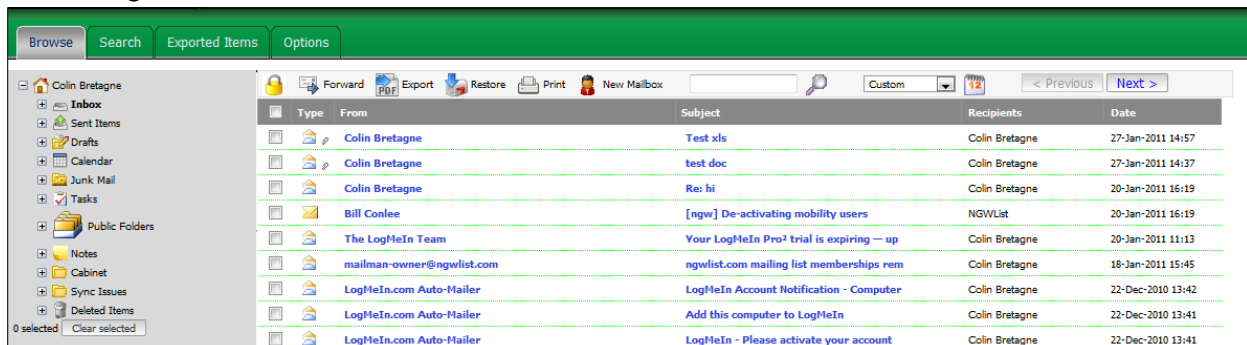
There are four tabs in the Search Interface:

- Browse- for casual browsing and filtering. You may only browse one mailbox at a time.
- Search- for in-depth search, possibly across many mailboxes.
- Exported Items- List of all export
- Options- for changing user settings. Similar to settings in the User Admin.


The browse screen shows you one mailbox at a time. The mailbox you’re viewing is listed at the top left and the folder structure is replicated as it is in the mail system. The view the screen displays is dependent on the source mail system for the selected or displayed mailbox. Accordingly, the icons and folder structure of different mailboxes will appear differently for a GroupWise account than it does for an Exchange account.

If you have the rights to access more than one mailbox in the mail system, (such as GroupWise proxy rights and the proxy option enabled in Retain), you will see the “New Mailbox” icon along the toolbar allowing you to switch between mailboxes you have access to. In addition, you may switch display languages using the drop-down language selector along the top.

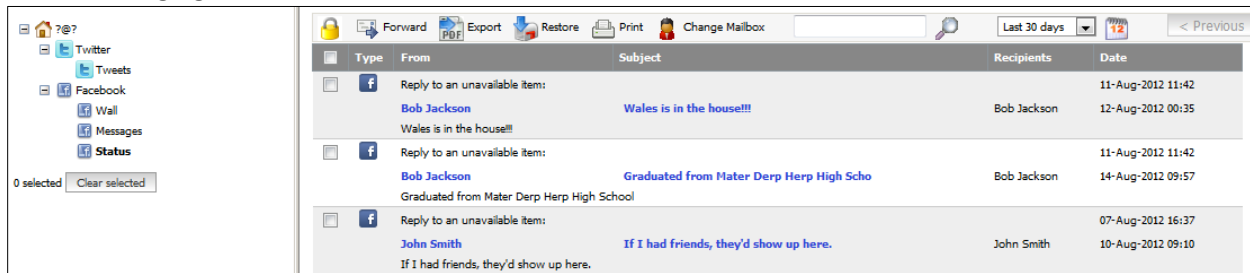
An Exchange mailbox:



GroupWise mailbox:



Social Messaging:

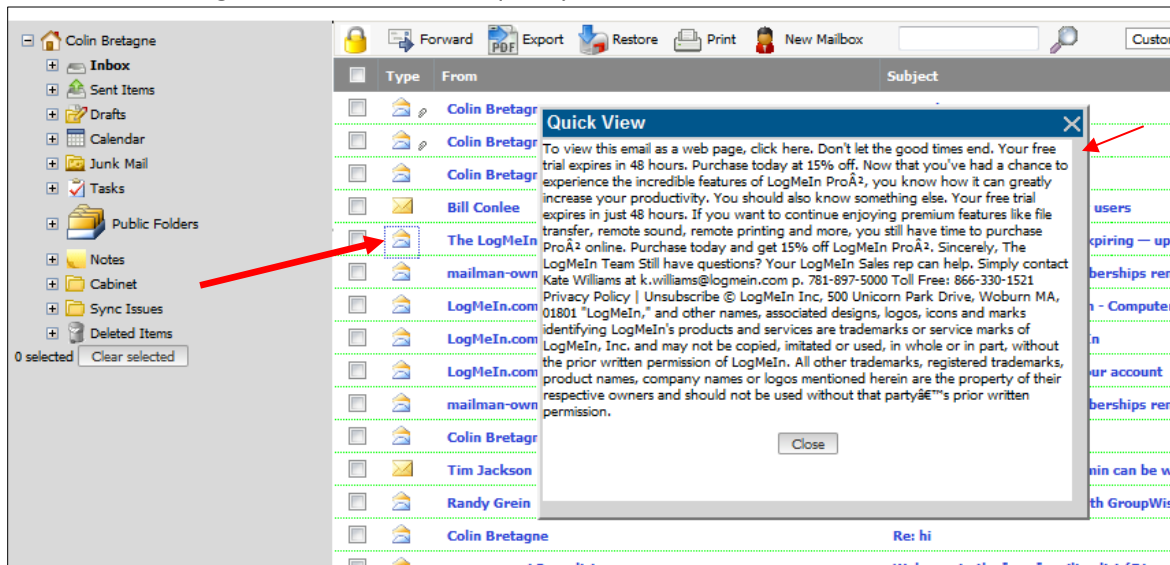


Cumulative Selections

In browse and search as you click on selections, these are preserved cumulatively from page to page or across searches, folders, etc. So you can forward, restore stub, PDF export from many places at once.

Quick View

The search messages interface allows to quickly view items.



Accessing the quick view is simple. Simply select the item icon in the 'Type' column, (mail envelope icon), to spawn the quick-view for that item. The quick view displays up to the first 2kb of text for that item in a small window.

NOTE: Date Range Filter



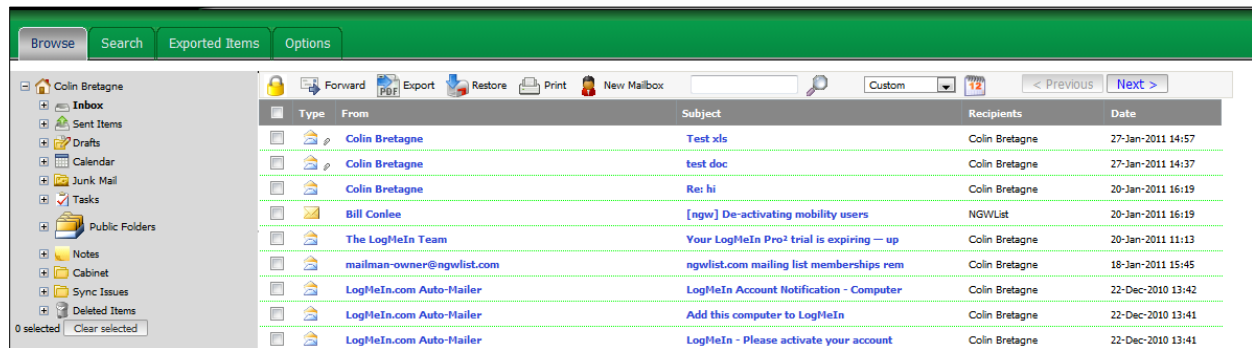
Please pay particular attention to the date range filter. Only messages that match the date range filter are displayed. If you don't see the results you expect, check the date range filter to be sure that the settings are correct. See [Specifying a Date Range](#).

Using Online Help

Retain has context-sensitive help screens that appear in the left column when you click the help icon.



Browse



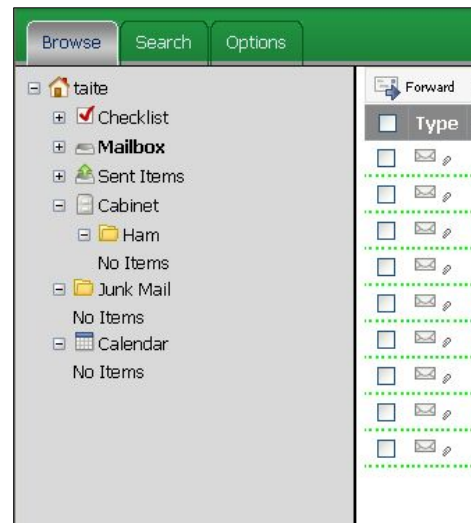
This is the default **Browse** screen.

The logged-in user's mailbox will be displayed by default along the top of the window. You may choose to browse through another mailbox if you wish (and if you have rights to do so).

If the user's mailbox is empty, or the user has no specifically assigned mailbox, then the user will be prompted to select a mailbox to browse. Only mailboxes which the user has rights to, will be available.


Notice that the folder structure of the mailbox is reproduced at left. You can expand the folder list until there are no more sub folders. When there are no more subfolders, you will see "No Items" marked.

You browse the folders in the current mailbox and when selected, their contents are displayed.




Language Selection

You may select what language the browse screen will use by selecting the drop-down language selector control and choosing your desired language. This selection applies to the search screen as well as the administration interface.




Merged Mailboxes


If a mailbox has been merged with its migrated counterpart or any other specified mailbox in Exchange, then the user will have the option to switch mailboxes with a drop-down selection box.



Mail from both systems will be viewable and manageable in the same window without having to logout or switch mailboxes.



Browsing Tools (Widgets)



Retain will display various options for the user. The specific options displayed depend upon the user's rights, if a user has not been granted the rights to a particular tool, that tool will not be present in their Search Messages interface. The rights are found in the Retain Server Management section, on the [Users](#) page. The full range of options includes:

- LITIGATION HOLD: Applies or removes the litigation hold, banning removal from archive. Litigation hold prohibits the deletion manager and jobs from removing applied and marked messages.
- FORWARD: forward selected items by e-mail.
- PDF EXPORT: opens the PDF export window to create a PDF of the selected items to be emailed to a specified mailbox
- RESTORE: Restores messages to the Mail system. See section for details.
- PRINT: prints a list of selected items.
- NEW MAILBOX: select a new mailbox to view.
- QUICK SEARCH: search the current folder for text in "from", "to", "subject" or body.
- DATE SELECTOR: select a custom date range. Retain will show items within this range. The date filter selection is saved automatically whenever it is changed in the browse and search windows and persists between logged-in sessions. Users should take care that the date range sufficiently covers the time range of interest – a common mistake is to assume Retain has not archived anything, when in fact the date range is simply too narrow.
- PREVIOUS/NEXT: move to the next or previous page of items.

Browse Window Defaults

- Retain remembers what the last browsed mailbox was.
- Non administrators go right to their mailbox when they log in.
- Administrators get a prompt to select the mailbox they want to view.

New Mailbox

If you click “Change Mailbox” (which is available if you have access to at least one other mailbox in the system), you will see the mailbox selector screen pop up. Select which mailbox you want to see by clicking on its radio button and clicking OK. This is just like the [mailbox selector used in the user rights administration section](#).

(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’ dialog. Otherwise, individual user accounts will be displayed.)

When selecting a single mailbox, the source 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.

First	Last	Mail Server	Mailbox	Dept.	Type	Domain
?@?	?@?	SocialNetwork	?@?		U	SocialNetwork

	ID	Name	Email	Department	Domain
<input type="radio"/>	d96	d96	BESX	d96	U blackberry.gwava.com
<input type="radio"/>	d97	d97	BESX	d97	U blackberry.gwava.com
<input type="radio"/>	d98	d98	BESX	d98	U blackberry.gwava.com
<input type="radio"/>	d99	d99	BESX	d99	U blackberry.gwava.com
<input type="radio"/>	Rick	Liljenquist	BESX	rickl	development U blackberry.gwava.com
<input type="radio"/>	MikeFirst	MikeLast	BESX	MikeAlias	mikedepartment U blackberry.gwava.com
<input type="radio"/>	Art	Purcell	BESX	art	U blackberry.gwava.com
<input type="radio"/>	u1	u1	BESX	u1	U blackberry.gwava.com

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.

Type	From	Subject	Recipients	Date
✉	<john@squirrel.com>		maston	03-Mar-2011 21:56
✉	<john@squirrel.com>		maston	03-Mar-2011 21:56
✉	<john@squirrel.com>		maston	03-Mar-2011 21:56

Toolbar Options:

Specifying a Date Range

You can narrow the scope of your browse by entering a date range. This is what it looks like.

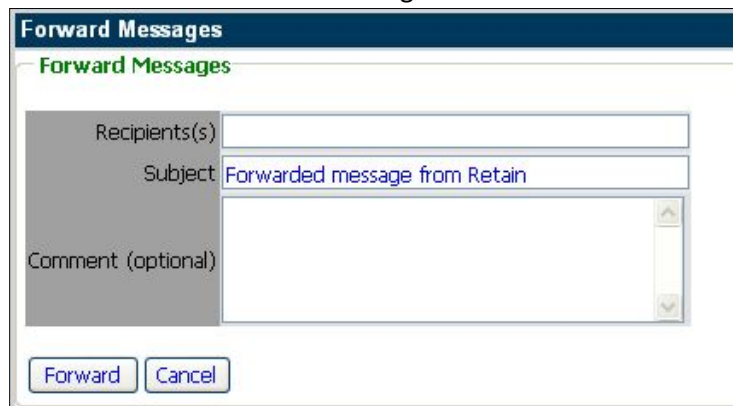


Litigation Hold

See the [Litigation Hold](#) section later in the document.

Forwarding a Message

In the main browser window, you can select a message by checking its checkbox and the click “Forward”. You will see this dialog:

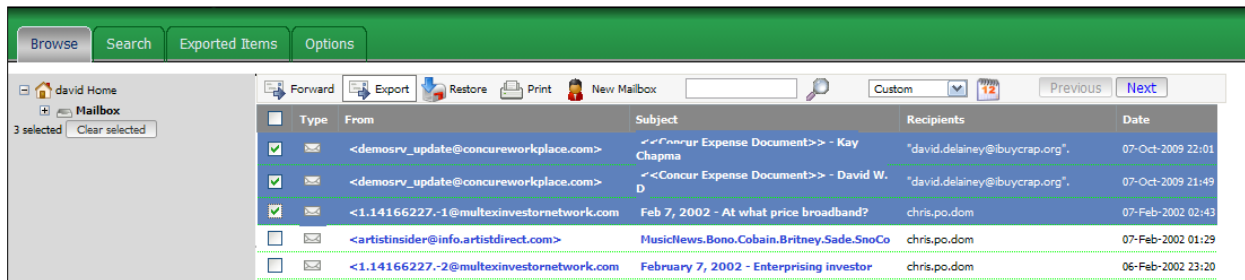


Enter the recipients (where you want the message(s) to be forwarded), separated by commas. Enter a comment if desired and click “**Forward**”. Retain will send the message(s) via SMTP.

PDF Export

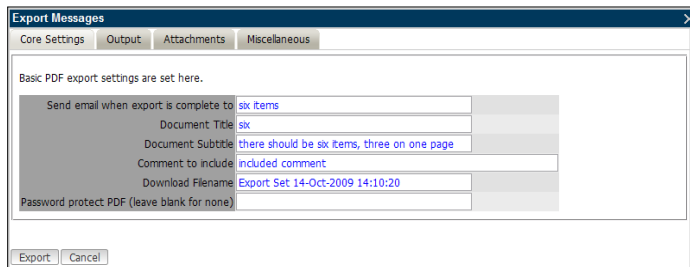
PDF Export is a new function which allows users, (with appropriate rights), to export single or multiple items into a PDF file. This PDF can be named and defined by the user exporting the document. Items will appear as navigable files in the PDF. The default view provides a list on top of the PDF viewer which works as a table of contents to provide access to the exported archive.

If you have rights to do so (new user right), can export browse or search results to PDF. Select the desired items and then select the “Export” button.

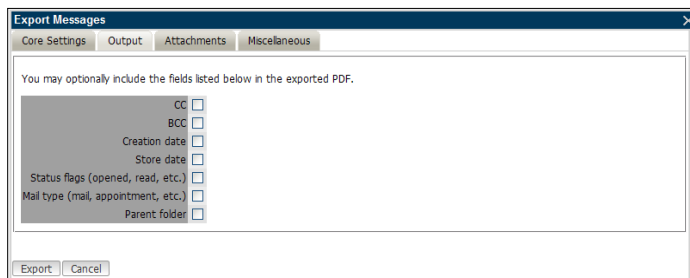


Fill out the various options as desired.

The different options include the file information.









The File output and included fields can be edited.




Any attachments to the email sent to the defined recipient.

...And finally, any restrictions to 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 larger, the archive will be exported into multiple PDFs will be created according to the maximum limit, until the archive export is complete. Time zone information will be used for formatting any time and date information in 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 finished, select 'Export'. Depending on the size of the PDF export list, the export may take some time. The Export will notify via email when it is complete, if a notification address was provided. Currently running export jobs are also listed under the 'Exported Items' tab. Once the Export has completed, the finished PDF is available under the 'Exported Items' tab. Locate the appropriate export list and select the disk icon to download the finished PDF.

Browse Search Exported Items Options				
A list of Export jobs is shown below. If you have any running (but not completed) jobs, they will be displayed, and you may request cancellation.				
For jobs that have completed, you may download and/or delete the exported file. Please note that the deletion cannot be undone.				
Running Jobs:				
No Items Found.				
Completed Exports:				
#	File	Size (KB)	Date	Action
0	Export Set 14-Oct-2009 14:10:20.pdf	39	14-Oct-2009 14:18:42	 
1	Export Set 13-Oct-2009 13:22:23.pdf	15	13-Oct-2009 13:23:02	 
2	Export Set 09-Oct-2009 14:19:41.pdf	39	09-Oct-2009 14:20:23	 

The finished PDF is structured with an email or item list on top as a working table of contents, and the selected item is displayed below. The title, subtitle, and comments are listed on the title page, as shown.

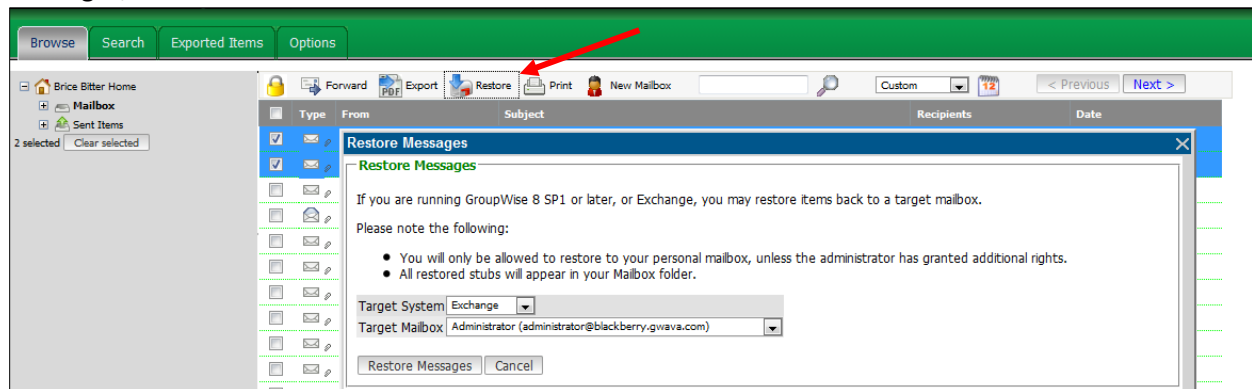


Restore

Retain can restore messages to target mailboxes in the mail system, for those users who have the appropriate rights to restore messages. If a user has the rights to restore messages their user interface will have the 'Restore' button visible on their toolbar. Restoring messages **requires** GroupWise 8, or Exchange.

The action that the restore function performs is dependent on what platform the messages are to be restored to. For GroupWise, when the restore function is called, the selected GroupWise message(s) will have a stub restored to the mail system. Because the stub connects to and references the full item in the Retain archive, this function is effectively a 'full restore' of the message or messages. If message(s) are restored to Exchange, a full copy of the entire message with attachments is restored to the mail system. GroupWise messages may be restored to either Exchange or GroupWise.

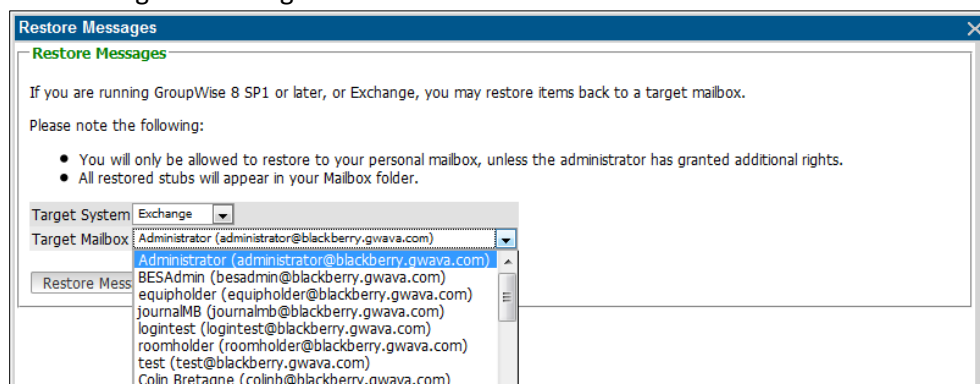
To restore a message or messages, log into the Retain archive, locate and select the desired message, or messages, and click on the 'Restore' button.



Select the desired target mail system for the restored messages. (This is usually the source mail system, though GroupWise messages may also be restored to the Exchange system.)



After the 'Restore' option has been selected, and if the user has rights to restore to any mailbox, the target mailbox must be specified, otherwise, only the user's mailbox will be available. When the 'Restore Messages' button is selected, Retain connects to the mail system and places a stub or the entire selected message in the target mailbox.

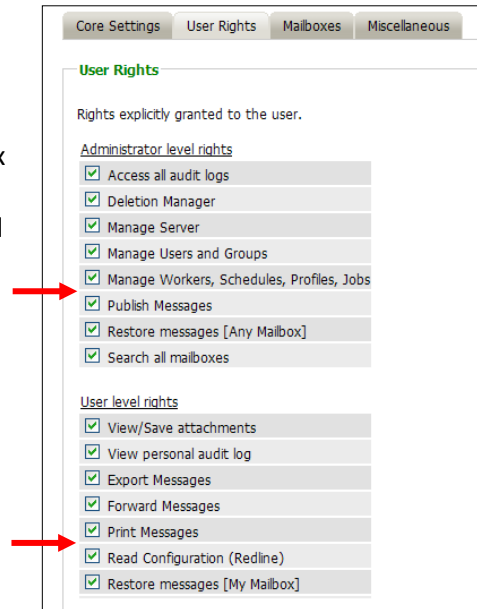


The Users Right that is required to use the Restore feature is called “Restore Messages”. Administrator level right allows users to restore messages to any mailbox, including mailboxes which did not have the message before.

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.


The Restore Stub feature requires:

- GroupWise 8 Client
- GroupWise 8 Post Office
- Retain Stubbing Server
- Stubbing correctly setup in ConsoleOne.
- Sufficient user rights to restore messages.



Search

This is the default **Search** screen.



In Search, you enter your criteria on the left, including the scope of your search. Multiple criteria are permitted and you can perform a precise search this way. Your search results appear in the main window. Users with more rights will have more tools and search options available. Only users with rights to the different options in the archive will have those options appear in their toolbar.

Retain will search across all mailboxes you have rights to. By default, a regular user only has rights to their own mailbox and therefore will only see results from that.

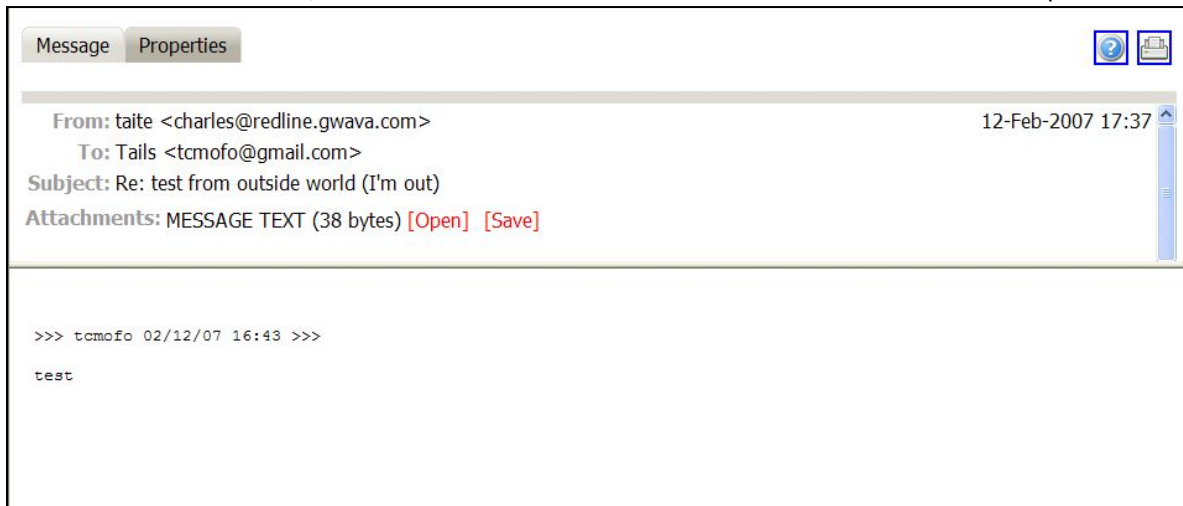
[See Search in More Detail](#)

Opening/Viewing a Message

Here is an example of a message opened in the Retain message window:



You will see all of the text, the headers and attachments as listed. Here is another example:



Notice how the message's metadata has all been saved: time/date, status of attachments and so on.

The user may print the message or save it to disk, if they possess the necessary [rights](#). You can also select the properties tab to see more information about the message:

Message Properties

From: taite <charles@redline.gwava.com> 12-Feb-2007 17:37
To: Tails <tcmofo@gmail.com>
Subject: Re: test from outside world (I'm out)
Attachments: MESSAGE TEXT (38 bytes) [Open] [Save]

Core

Item Type: Mail
Box Type: sent
Owner UID: 013BDEC0-0C5E-0000-AD6D-170031001600
Parent Node: 839
Read: no
Created: 12-Feb-2007 17:37
Delivered: 12-Feb-2007 17:37
Stored: 04-Jun-2007 12:27
Path: / taite / Sent Items / Re: test from outside world (I'm out)

Additional Properties


domain: gw2dom
priority: Standard
status: accepted,opened,read
postoffice: gw2po

The message's item type, status, relationship information, opened status, original location and so on are all stored. All of these items are valid based on the time the message was archived from GroupWise.

The properties section is where a lot of fields that are specific to the Mail system are saved.

Search in More Detail

The search window has some of the same top buttons as the browse window.



Search Window

The browse window allows the user to access a single mailbox at a time, access the folder tree of the mailbox and perform some basic searching and filtering. For more advanced searching options, and the ability to search multiple mailboxes at once use the Search Window.

Core Tab:


Search Terms – What To Look For

You specify up to 5 additional criteria here. All items are logically AND-ed together to narrow the search unless specified with a double pipe. Using the double pipe denotes an 'or' search criteria. (I.e. A search criteria with 'Retain || Vertigo' would result with messages containing the words 'retain' OR 'vertigo' in the specified field.) To use the Internet Header search term, you must

[enable the Internet Header options in the profile.](#)

You can search based on all items shown on the right with operators including:

- Contains (fuzzy)
- Contains (exact)
- Starts with
- Does not contain

You must have at least one search item. You can add and delete them using the  buttons to the left.

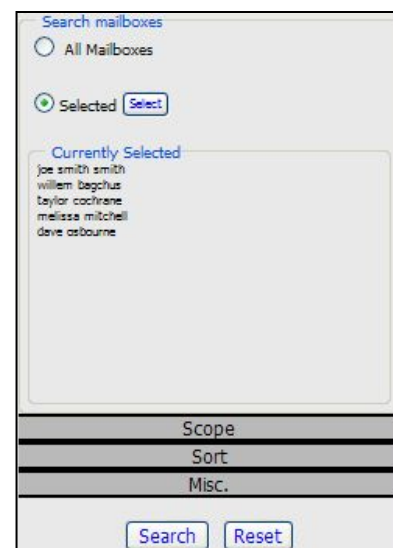
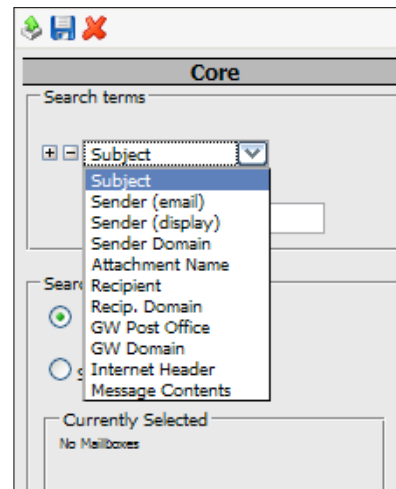
Search Mailboxes – Where to Look

Click the radio button to search through all mailboxes (that you have rights to) or click the other radio button to search through selected mailboxes.


You choose which mailbox, or mailboxes to search through using the mailbox selection tool. ([See below](#)) Mailboxes that you have selected for the search are listed in the Currently Selected panel.

Start Your Search

To initiate a search, click the search button. The results will appear on the right, and will be saved as a persistent query. Every time you execute search, even between login sessions, the same query will be executed. If you click on the Reset button, the current query is cleared.



The toolbar on the right provides similar functionality to the browse window toolbar, and allows the user to forward messages from the search results, print, etc. - assuming the user has rights to do so.



RESET – clears the query entirely.

Selecting Mailboxes


When you click on SELECT, you choose which mailboxes to search through using the selection tool. Use the check boxes to select mailboxes and click “Add Selected” to add them to the list in “Currently Selected”. You may remove mailboxes from the “Currently Selected” list by clicking on the red X next to them.

You will see the mailboxes made available to you via GroupWise proxy rights (if the function is enabled in Retain), via the “Search All Mailboxes” right or via mailboxes you have been granted access to specifically.

First	Last	PostOffice	Mailbox	Dept.	Type
<input type="checkbox"/>	Willem	Bagchus	gw2po	Willem	U
<input type="checkbox"/>		Greg	gw2po	Lynn	U
<input type="checkbox"/>	Colin	James	gw2po	Colin	U
<input type="checkbox"/>		Miss	gw2po	Hanigan	U
<input type="checkbox"/>	carl		gw2po	Heinz	U

Loading, Saving and Deleting Searches

The user may save searches for future use, and reload them whenever they are useful. These search queries may also be deleted freely.



The icons (below the tabs) from left to right mean:

- Load
- Save
- Delete

You can save up to ten queries and you may delete queries you no longer want.

Scope Tab:

This tab contains functions to further narrow what you're searching through.

Item Type

If no boxes are checked, you will search through every mail system item type. Otherwise, Retain will search through only those types of items you have checked off.

Item Source

If no boxes are checked, you will search through every item source. Otherwise, Retain will search through only those items whose source matches what you have selected.

Attachment Size

When searching through messages, you can tell Retain to search through all messages regardless of the size of the attachments or whether to restrict your search to only those items whose attachments are within the size range you select.

Sort Tab:

Choose how you want Retain to sort your search results.

You may sort your results by the date the items were created in the mail system, by the Subject, Sender, Sender's Domain, the Recipient's e-mail address, or display name.

Your sort may be up to three levels deep.

Misc. Tab:

Choose your Date Range here. It is important to note WHAT date is being filtered here. It is NOT the same as the date range selector on the results window. The results window date range filters based on the date the item was created in the mail system.

In this case, it is for filtering appointments and tasks. So you are filtering based on the start/end dates of the item, regardless of when it was created in the mail system. For example, you might have created an appointment for yourself that will happen in three months. The creation date is three months away from the start/finish dates.

The screenshot shows the 'Sort' tab selected in the interface. It features a 'Sort Results' section with three dropdown menus: 'First By:' set to 'Creation Date', 'Then By:' set to 'Sender Domain', and 'Then By:' set to 'Recip. (display)'. Below this is a 'Misc.' section containing 'Search' and 'Reset' buttons.

The screenshot shows the 'Misc.' tab selected in the interface. It features a section titled '(Appointment, Task)' with two date range selectors. The first is 'Start Date' with a dropdown set to '(any date)' and a 'Range' button. The second is 'End/Complete Date' with a dropdown set to '(any date)' and a 'Range' button. At the bottom are 'Search' and 'Reset' buttons.

Litigation hold


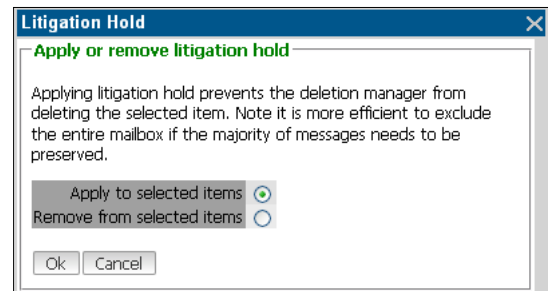
Litigation hold is a feature which counteracts the deletion manager. To use the Litigation hold feature, users **MUST** have the litigation hold right granted to them in the Retain Server.

The deletion manager is how mail is expired out of the archive, and removed when no longer needed due to age or other circumstance. Litigation hold tells the system to ignore the remove command from the deletion manager; messages or entire user mailboxes with litigation hold applied will not have any mail removed from the system until the hold is removed.



To apply the Litigation hold to an item, select the item then select the padlock at the top of the browse window. The following window will be displayed:

Choose whether to apply or remove the hold from the selected items. The entire use mailbox may be excluded from a deletion job in the deletion manager.




Items that have a litigation hold active on them are marked by a small padlock in the interface. If a hold needs to be removed, selecting the item then the litigation hold button at the top will bring up the option to remove the hold from the selected items.

Shared Queries


Queries may be shared with other users in the system. Shared queries are still subject to the individual rights of any user that the query is shared with. A shared query shows the full results to the query but does not grant access to the messages. If a user is restricted from viewing the mail from another user, but has a shared query, they will be unable to access the messages outside their rights.

Any search may be saved and any saved query may be shared with other users. When a shared query is created, the creator decides which users to share the query with. Once a query has been saved by using the disk icon at the top of the search window, it can become a shared query.

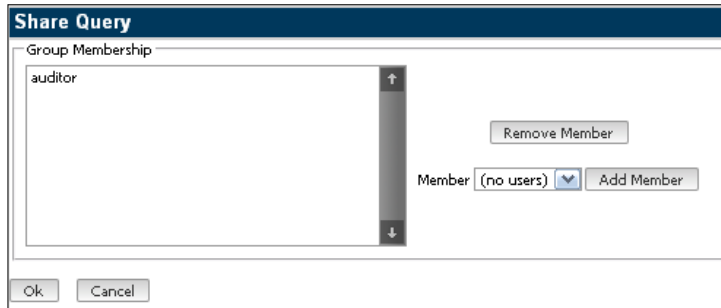
Creating a shared query takes two main steps. A query must be shared, and then accessed by the user it was shared with. To share an existing query, select the sharing icon at the top of the search window. This icon looks like red and blue users. The icon is between the save, and the close icons.



After selecting the share query icon, a window displaying all saved queries is displayed.



There are no queries shared in this window. The queries are listed, displaying name and description. If a query is to be shared, select the gray icon at the beginning of the saved query.




After selecting the shared query icon, the query creator must now select which users to share the query with.

Use the drop-down menu and 'Add Member, Remove Member' buttons to populate the group the query is to be shared with.

Once the group has been selected, click 'Ok' to be returned to the Share Query window. If the share query icon is now colored instead of gray, then the query has been successfully shared, and the users in the shared group now have access to the query.

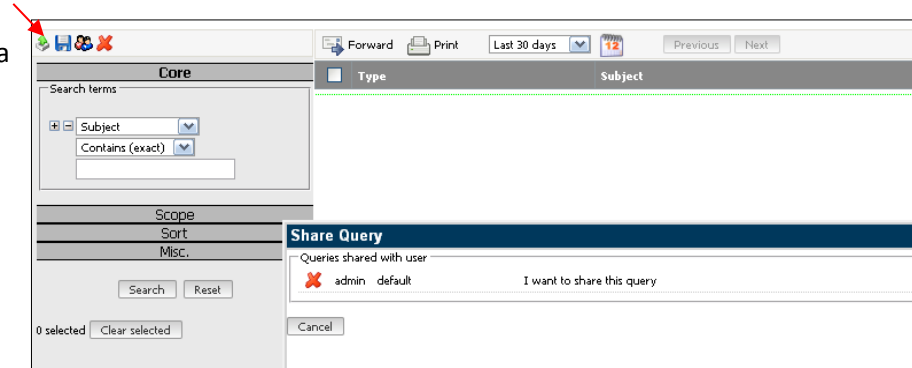
There is no limit to the amount of users that a query can be shared with. Adding users to an existing shared query group is accomplished by selecting the shared query icon and editing the existing list. The user which created and initially shared the query retains the rights to remove users from the shared query group membership or cancel the share altogether.



Accessing Shared Query

To access a shared query, a user which has membership in a shared query group logs into Retain and selects the Search tab. In the Search window, the user selects the Load icon at the top left of the search window. All saved queries will be

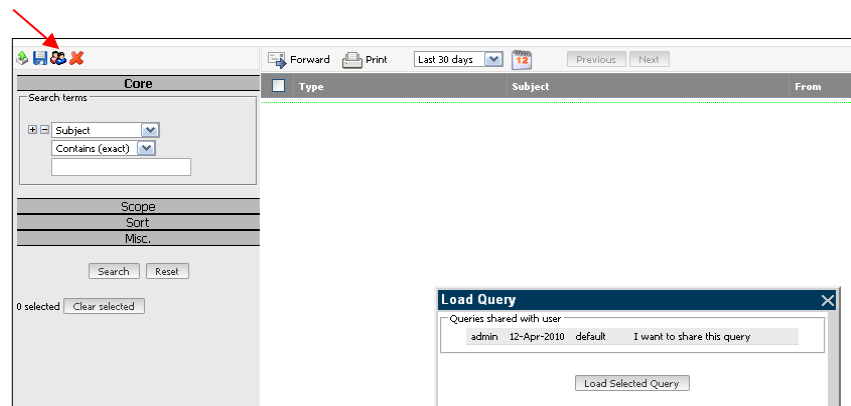
displayed, and the query that was shared will be available to be viewed. Select the desired query and the 'Load Selected Query' button to load the search criteria. Clicking 'Search' loads the messages fitting the shared criteria.




Removing a Shared Query

Users have the option to remove a shared query from their saved query list.

If a user desires to remove a shared query from their query list, it is accomplished through the 'Shared Query' window. Select the 'Shared Query' icon to view all queries that are shared, and that the user is sharing.



Selecting the red 'X' next to the offending query in the shared query window will remove the query and remove the user from the shared query group.



This action cannot be reversed, but the query may be re-shared.

Search Window

Options Tab

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.

The screenshot shows the 'Options' tab in the GWAVA Retain interface. At the top, there are four tabs: 'Browse', 'Search', 'Exported Items', and 'Options'. The 'Options' tab is selected. Below the tabs, there is a section titled 'Options' with a user icon. Underneath, there are four sub-tabs: 'Core Settings', 'User Rights', 'Mailboxes', and 'Miscellaneous'. The 'Core Settings' sub-tab is selected. Below the sub-tabs, the 'User ID' is displayed as 'admin'. The 'Core Settings' section contains a table with the following information:

Comment	
Authentication Method	Offline Authentication (forced) (Use exclusively)
Primary UID (links your rights)	df8616e9-c942-42a8-919e-6a4fb7e2ccc0 (none)
Group Membership	(none)
Account can expire	<input type="checkbox"/>
Change Internal Password	Old <input type="text"/> New <input type="text"/> Confirm <input type="text"/>

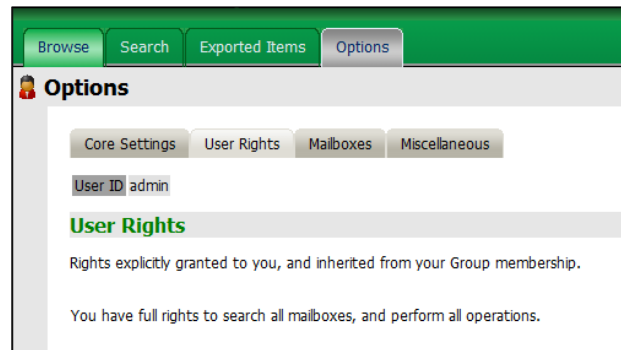
Core Settings

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

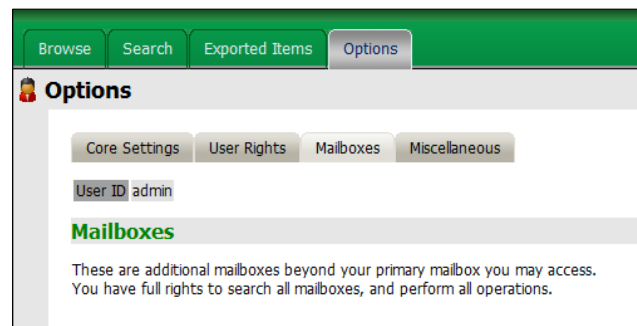
User Rights Tab

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.



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.



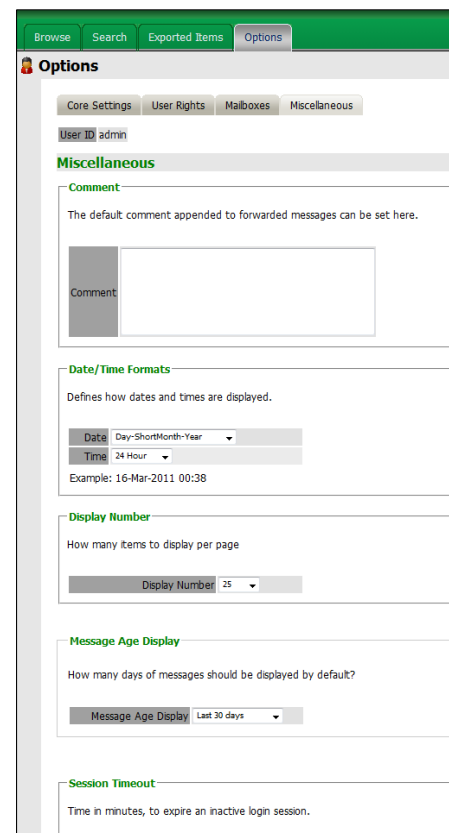
Miscellaneous Tab

This contains your display and session options.

These work just as described in the [user display](#) options in the Administration screen.

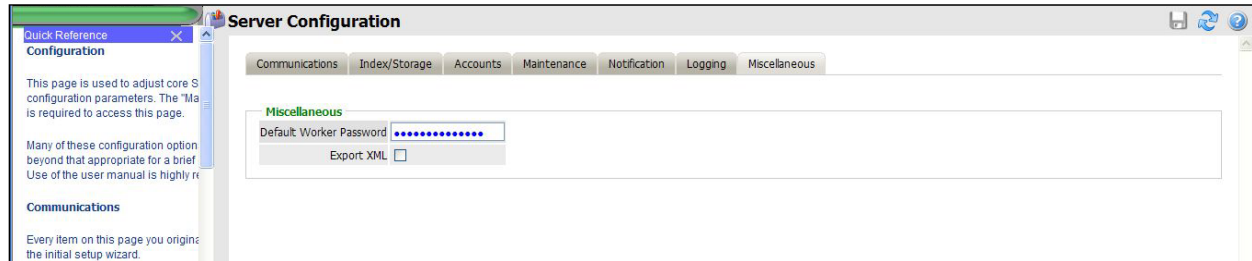
Options are:

- **Comment** Default comment for forwarding messages.
- **Date/Time Format** How you want your dates and times to be displayed.
- **Display Number** How many items to display per page.
- **Message Age Display** Default date filter for searching. It may be changed on the fly.
- **Session Timeout** Inactive session timeout. Can be between 5 and 60 minutes.



Using the Retain Help System

Retain has context-sensitive help screens that appear in the left column when you click the help icon. (



To view help, click the help icon on the right (the blue question mark circle icon) and the help will appear on the left.

Scroll through and read what you need. To hide the help and make the menus re-appear, click the X at the upper right of the help screen.

All topics in the main window are covered in help.

About and Diagnostics Page

There is a special page in Retain containing credits for those who helped with the project.

There are also valuable diagnostics on this page. When you tune the memory usage of the Retain components, you may want to refer to these diagnostics and statistics.

To get this page from anywhere in the administrative console of Retain, click on the GWAVA logo in the top left corner, or the 'about' link in the top right.

There are two parts to this screen. The credits part is at the top.



The diagnostics part is below.

Operational statistics are listed such as uptime, component version 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.

uptime	0d : 1h : 9m
currentTime	Tue Aug 02 22:51:48 MDT 2011
timezone	Mountain Standard Time
retainName	SLES11x64
jvmVersion	20.1-b02
jreVersion	1.6.0_26-b03
javaHome	/opt/beginfinite/retain/java6/jdk1.6.0_26/jre
tomcatHome	/opt/beginfinite/retain/tomcat
indexDoc	45363
logHome	/opt/beginfinite/retain/tomcat/logs
osArchitecture	amd64
osName	Linux
userName	tomcat
tmpDir	/var/cache/retain-tomcat5/work/Catalina/localhost/RetainServer
tomcatVersion	Apache Tomcat/5.5.27
tomcatFreeMemory	193878
tomcatMaxMemory	932096
tomcatTotalMemory	250752
indexEngine	lucene
indexHost	
indexPort	10000
indexProtocol	http
retainHome	/opt/beginfinite/retain/RetainServer
retainEdition	Standard
retainBuild	250
retainAvailability	alpha
configSchema	16
messageStoreSchema	3
indexingSchema	1
derbyVersion	10.5.3.0
indexingType	lucene
indexingVersion	3.0.3
hibernateVersion	3.3.2.GA

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

Diagnostic Buttons

In usual operation, none of these buttons will be used without instruction of support. The diagnostic buttons found on the diagnostics page are used in conjunction with support and can be dangerous for the system if used without understanding or instruction. The utilities contained 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.

Core info, JAR List, Jobs, Schedules, Profiles, Workers

The **Core info, JAR List, Jobs, Schedules, Profiles, and Workers** buttons initiate a download of the configuration .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 configuration.

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.

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.

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.

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.

Mailbox Merger / Splitter

This utility lets you move a mailbox to belong to part of the virtual mailbox of another entity. It should be used with care, and knowledge. Backups recommended.

Merge

Source Mailbox:
admin (po.dom)

The virtual mailbox associated with this item has the following mailboxes connected to it:

- admin (po.dom), which has these mailbox types connected:
 - gw (4AC5FF80-11DB-0000-B1DB-F15CE5F2B5EF)

DestinationMailboxGroup:
New Entity

This uid doesn't have ANYTHING associated with it

After moving, if the original entity container is now orphaned, remove it

☐

Merge

Batch Load from a File

See [mergesample.xml](#) for the format.

XML File:

Merge

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 performed once.

Manually merging and splitting mailboxes is repeatable and is also the only way to undo a merge, automatic 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

This utility lets you move a mailbox to belong to part of the virtual mailbox of another entity. It should be used with care, and knowledge. Backups recommended.

Merge

Source Mailbox:
Administrator (LION.lion.gwava.com)

The virtual mailbox associated with this item has the following mailboxes connected to it:

- Administrator (LION.lion.gwava.com)
- migration (LION.lion.gwava.com)
- migration3 (LION.lion.gwava.com)
- migration4 (RHINO.lion.gwava.com)
- migration5 (LION.lion.gwava.com)
- David Andersen (PO1.Domain)
- David Anderson (LION.lion.gwava.com)
- Jason Bailey (PO1.Domain)
- Jason Bailey (LION.lion.gwava.com)
- Russ Bateman (RHINO.lion.gwava.com)
- Russ Bateman (PO1.Domain)
- Michael Bell (LION.lion.gwava.com)
- Ben Bishop (LION.lion.gwava.com)
- Ben Bishop (PO1.Domain)
- Brice Bitter (LION.lion.gwava.com)
- Brice Bitter (PO1.Domain)
- Colin Bretagne (LION.lion.gwava.com)
- Tony Caras (RHINO.lion.gwava.com)
- Tony Caras (PO1.Domain)
- Marlow Draney (RHINO.lion.gwava.com)

DestinationMailboxGroup:
New Entity

This uid doesn't have ANYTHING associated with it

After moving, if the original entity container is now orphaned, remove it

☐

Merge

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 previously merged accounts.

GW to Exchange Scope

The GroupWise to Exchange Scope utilizes mailbox mappings to direct data from the mailbox and is required for the Migration.

The scope is used to specify which users are to be included in a migration job. If a user's mailbox data is desired to be migrated from GroupWise and placed into Exchange, they must be included in a scope file used by the migration.

From this point you can create a new scope file or modify an existing scope file. Please select one of the two options to continue.

Select 'New Scope File' unless a scope file already exists and is to be modified. After an uploaded file is loaded and modified, a new file will be created.

From the new scope page, users which have been merged are displayed below. For any users to appear, their accounts must first be merged.

Select the mailboxes mappings that this scope file should contain. Mailbox mappings can be created automatically or using the MergeUD tool.

<input type="checkbox"/>	Source	Destination
<input type="checkbox"/>	BenB Ben Bishop benb@retainesting.gwava.com	benb Ben Bishop benb@lon.gwava.com
<input type="checkbox"/>	BriceB Brice Bitter briceb@retainesting.gwava.com	briceb Brice Bitter briceb@lon.gwava.com
<input type="checkbox"/>	Candace Candace Riedelbach candace@retainesting.gwava.com	candace Candace Riedelbach candace@lon.gwava.com
<input type="checkbox"/>	DaronP Daron Parcell daronp@retainesting.gwava.com	daronp Daron Parcell daronp@lon.gwava.com
<input type="checkbox"/>	DavidM David Morgan davidm@retainesting.gwava.com	davidm David Morgan davidm@lon.gwava.com
<input type="checkbox"/>	EricM Eric Maughan ericm@retainesting.gwava.com	ericm Eric Maughan ericm@lon.gwava.com
<input type="checkbox"/>	Ernie Ernie Riedelbach ernie@retainesting.gwava.com	ernie Ernie Riedelbach ernie@lon.gwava.com

Select the desired users from the list by selecting the checkbox next to the name. A global select option is available at the top.

Once the desired users have been selected, click on the 'Generate Scope File' button to create and download the scope file.

GW Exchange Migrator

The GroupWise to Exchange Migrator is enabled via an additional license file. Because of the complexities of the Exchange environment and possible configuration options with Retain, it is recommended that you consult with your reseller or GWAVA Sales Engineer who can provide the license as well as assistance to ensure a successful system migration.

Warning ✕

You can really badly break a mail system with this tool.

By clicking "Agree" below you understand the risks associated with this utility and that use is governed by the end user license agreement.

The process to migrate mail from GroupWise to Exchange begins with an existing Exchange system containing existing accounts equivalent to those found in GroupWise. The GW to Exchange migrator will not create accounts, it must have destination accounts to place the mail, and be told which types of mail, and which accounts to put it into.

Migrating data from GroupWise to Exchange is not a straight-forward endeavor, as Exchange contains default settings which, at default, block messages from 'untrusted' source, messages with specific attachment types or attachments which are too large. (Default limit for Exchange is 10 mb.) Setting appropriate trusted message types, sources, and size settings in Exchange are the responsibility of the system administrator. GWAVA Inc. is not responsible for any data that is not migrated due to incorrect settings.

To migrate data from an existing GroupWise system into an Exchange system, several steps must be performed in correct order.

1. Configure and cache the address book for both GroupWise and Exchange modules.
2. Merge the virtual Retain accounts either [manually](#) or via the automatic merge through [mailbox mapping](#)
3. **Immediately** run an [Address book refresh](#) from the modules page. (If the address book is not re-cached *immediately* after an automatic merge, the merge will fail and addresses will have to be manually separated and re-connected. Forcing a failed automatic merge can be accomplished by contacting GWAVA Support.)
4. Create a migration [scope profile](#)
5. Open the migration tool and fill in the pertinent information
6. Run the migration

To initiate a migration, select the 'GW to Exchange Migrator' from the diagnostic page. (If your system license does not enable Migration, you will be notified; contact your Retain sales representative.)

The migration tool requires a scope, which functions as a profile for the migration specifying included mailboxes, types of mail to be migrated, and the applicable date range. When an appointment is migrated from GroupWise into Exchange, an alarm can be set in the Exchange system to notify the user of the active appointment. The optional start date creates alarms only for appointments after specified date.

Migrate Store

☐ Enable Null Mode

Default Sender:

Start Date:

End Date:

Scope File:

Duplicates:

Message Types:

- ☒ Mail
- ☒ Appointment
- ☒ Note
- ☒ Task
- ☒ Phone

Enable Alarms: ☒ Set Alarms for applicable items

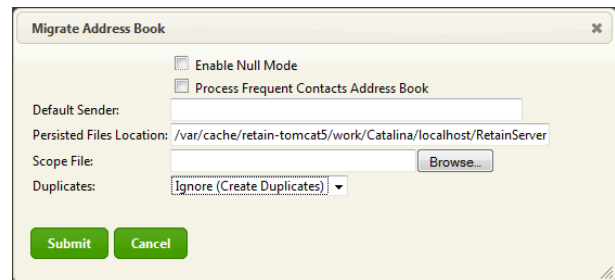
Start Date:

Alarm Time (Minutes):

When using the 'Migrate Store' option, only users specified in the scope will be migrated. Browse to the saved scope file, set the start and end date, (blank migrates everything), the message types desired, and then click 'submit' to initiate the migration. If a specific user is necessary for mail system operation, specify a default sender name.

Migrate Address Book migrates the users' personal address book for the users specified in the scope file.

Be sure to manage the duplicates accordingly, as the migration may be performed multiple times.

The screenshot shows a 'Migrate Address Book' dialog box with a title bar and a close button. It contains several configuration options: two checkboxes at the top, 'Enable Null Mode' and 'Process Frequent Contacts Address Book'; a 'Default Sender:' text field; a 'Persisted Files Location:' text field with the path '/var/cache/retain-tomcat5/work/Catalina/localhost/RetainServer'; a 'Scope File:' text field with a 'Browse...' button; a 'Duplicates:' dropdown menu currently set to 'Ignore (Create Duplicates)'; and two green buttons at the bottom, 'Submit' and 'Cancel'.

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.

System Properties

The System Properties utility displays basic system information and settings about the Retain Server, host environment, and functioning parts.

Appendix A: Backing up Retain

There are a few areas of Retain that are important to backup.

- ASConfig.cfg
- License
- Indices
- Archive
- SQL Database

The storage directory and location holds all of these files, and the backed up versions of those that should not be backed up while in use. (To find your storage location, see the [Storage tab](#) under Server Configuration.)

Retain automatically creates a backup of certain files, you may specify the settings for the Lucene index backup under [Maintenance](#). Retain will automatically backup ASConfig.cfg and Lucene Indices. (Exalead indices are **not** backed-up by Retain, as Exalead is a completely separate system and should be backed up individually.) These backups are located in the storage location, under the backup directory. The Archive and the License are kept in the storage location as well, though Retain does not create backups of these, so these require a manual backup.

The sensitive location to backup for Retain is:

<your storage location>/backup
<your storage location>/license
<your storage location>/Archive

Steps to back up the SQL database are dependent on the SQL database. See the recommendations of your SQL vendor to find instructions and backup procedures.

Fine tuning GWAVA Retain

Tomcat Memory Usage Tuning Overview

WARNING: *Failure to tune Tomcat's memory usage will guarantee failure of the Retain components and may lead to severe and irreparable data loss and/or corruption!*

The Tomcat servlet container is configured with parameters to indicate how much memory is available for allocation to web applications. Regardless of the amount of physical memory available in the server machine, Tomcat will never exceed these values.

Tomcat is preconfigured with extremely conservative memory allocation parameters— usually from 64 MB to 256MB RAM parameters. This is by design, to avoid taking away too much memory away from other server processes. Once memory is allocated by Tomcat, it is NEVER returned to the general server memory pool.

Two Critical Issues

However, Tomcat's conservative settings present two critical issues:

- The memory allocation is shared among all web applications running under the Tomcat instance. You **cannot** specify a specific amount to be reserved for one specific web application. This is the major reason it is strongly recommended that Retain components (especially the Server component) run on a dedicated Tomcat instance on a dedicated server -- Other web applications such as WebAccess, iManager, iFolder, etc., may consume too much of the shared resources.
- The preconfigured memory parameters may be sufficient for the Retain Worker components. They are definitely not sufficient for the Retain Server component.

Failure to alter these parameters will guarantee a slow Retain Server, or severe unexplainable crashes and/or irreparable data loss. Recovery of information can potentially be impossible.

Hence, it is crucial that the administrator carefully examine and alter the memory allocation parameters on each server.

General Recommendations

The following are our general recommendations for the maximum memory allocation parameter. The initial memory allocation parameter may be anywhere from 50%-100% of these values. These parameters are discussed in more detail in the next section.

Retain Server Component

A minimum of 1 GB of RAM is recommended. For larger systems, 2-3 GB RAM may be fully warranted. Remember, if this RAM is in use by Tomcat, it is unavailable to other server processes – leave sufficient physical RAM for these. Of course, these others server processes are ideally minimal. Also, remember this memory is shared with other web applications.

Retain Worker Component(s)

The Retain Worker component(s) have been designed to be light weight. Nonetheless, a minimum of 64 MB RAM is recommended, and 160 MB is strongly recommended, to give room for future functionality.

Remember, if this RAM is in use by Tomcat, it is unavailable to other server processes – leave sufficient physical RAM for these. Also, remember this memory is shared with other web applications.

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 usage 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 sizeable 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.**

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 128k-160k. Increase conservatively, in 64 KB chunks, if you see out-of-stack-space errors in the logs.

How to change memory values

Linux Embedded Tomcat Memory Tuning

For Embedded Tomcat which ships with Retain, Edit the following configuration file:

```
<<Tomcatconfigurationfilej2ee location>>
```

```
<</etc/opt/beginfinite/retain/tomcat/j2ee>>
```

Edit the following line with the appropriate parameters for your system:

```
CATALINA_OPTS="-Xms256m -Xmx1024m -Xss160k"
```

This line sets the initial memory pool at 256MB, the maximum at 1024MB, and the stack size at 160KB.

By default, the embedded Tomcat is pre-tuned for basic functions, with 1024MB as the default for the Retain Server, and 256MB for the Retain Worker. THIS WILL NOT BE ENOUGH IN A FULL PRODUCTION ENVIRONMENT. Please tune to fit your system needs.

Windows Embedded Tomcat memory tuning:

- 1) Go to the System Tray icon for Tomcat. (If using the embedded Retain Tomcat, launch tomcat5w.exe located in C:\Program Files\Apache Software Foundation\Tomcat 5.5\bin. Skip step 2.)
- 2) Choose configure.
- 3) Go to the Java tab.
- 4) Set the Minimum Memory (always used), Maximum Memory, and Stack Size. (Retain Tomcat is pre-tuned to 1GB for the Server, and 256MB for the worker, and 160K stack size.)

Checking memory:

After logging into Retain Server, click the bug icon at the upper left. Among other stats, you'll see the tomcat Memory statistics. These are for all web apps, not just Retain Server.

UPGRADED SYSTEMS

If you upgraded, it is STRONGLY recommended to return any previous memory tuning back to normal for the original Tomcat while using the embedded Tomcat that comes with Retain. If you are not using the standard Tomcat for any other web resource, simply uninstall or completely disable the standard Tomcat installation.

Windows:

The embedded Tomcat installation overwrites the standard Tomcat installation, and any memory tuning you have previously completed no longer is in effect.

Appendix B – Troubleshooting for Common GroupWise Problems

General GroupWise Error troubleshooting can be effectively done by checking the GroupWise documentation. For any error code not found below, download and check the GroupWise documentation. (<http://www.novell.com/documentation/groupwise.html> select your GroupWise version and have a look into the Troubleshooting guide for the Error Messages.)

Or, the GroupWise online documentation is usually found at
www.novell.com/documentation/<groupwiseversion>/
i.e.

<http://www.novell.com/documentation/gw8/>

<http://www.novell.com/documentation/gw7/>

From the troubleshooting menu, select 'Error Messages'

Some common error codes are:

EA04

EA04: The response is too large, SOAP protocol error, or the POA received too much data at once. This error message is largely invalid and gets thrown due to a bug in the current (7.0.2) GroupWise Post Office Agent. Retain has been tested with multiple attachments, some even tens of megabytes in size, that it does get the entire attachment.

To change the setting, edit the POA startup file and add the `/soapsizelimit` variable. Please consult the Novell GroupWise online documentation for the syntax.

Example on NetWare: `/soapsizelimit-2048`

The real solution will be in GroupWise 7 SP3. One workaround is to set the chunking size limit to 2048. We will only take 1024 chunks, and this setting is sufficiently high to fool the POA into not throwing the message.

D712/D714

D712/D714 The administrator has disabled/expired the entire post office or a specific user database. Retain cannot log in.

D714 The Database is temporarily disabled. GroupWise believes a MOVE is in progress. The MOVE must complete before Retain is allowed to log in.

C05D

C05D Item store is missing from Database. A mandatory database file is missing. This implies one of two possibilities:

- 1) Significant database corruption in the user database, which may be fixed with GWCHECK, or
- 2) This account has never been logged into before. Until an account is logged into at least once from the GroupWise client, Retain cannot access it.

Other notes regarding GroupWise 7, SP 3 (7.0.3):

Retain cannot access mailboxes with either of the following set


- Limited client license

Novell will be addressing the limited client license issue in GroupWise 7 SP 3.

Appendix C – The Tools Folder

There is a TOOLS folder included in the Retain download package.

The tools are grouped into categories, diagnostic and Retain archive tools and platform specific:



GroupWise tools

GroupWise Archive Migration Tool

This migration tool is designed to import native GroupWise archives directly into Retain.

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.

Archive Migration Tool

The Archive Migration Tool is for taking your existing native GroupWise archives and sending them to your Retain archive.

The directory contains 5 files:

- 1) The migration tool itself.
- 2) The migration tool configuration utility.
- 3) Two support DLL's and a readme file.

System Requirements

- A Windows XP workstation running GroupWise client 7.0.1 and up.
- Access to the archives you want to migrate.

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.

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 GroupWise 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 specified.

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.

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.

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 workstation 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.
- 3) 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 settings 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.

Server Connection

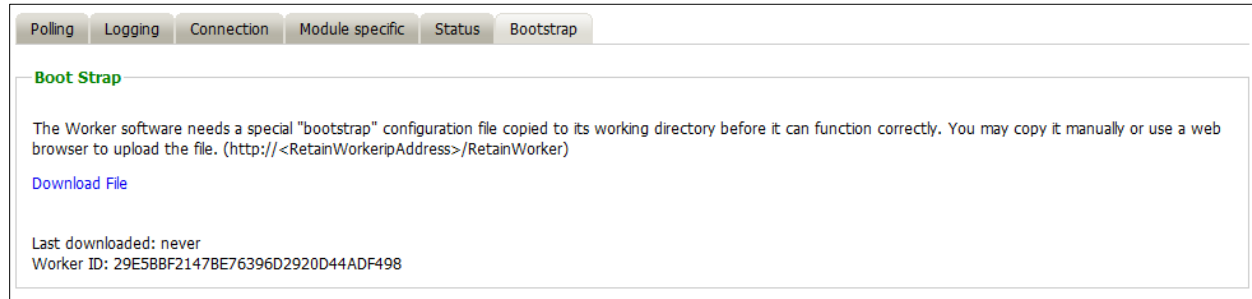
The worker will need to know how to connect to the Server. It is highly recommended that you use a DNS host name for the Host Name. It's important to note that you are specifying the URL from the point of view of the **Worker contacting the Server**, which may or may not match the URL you'd use from your browser to reach the Worker or Server.

For administrators using the Reload interface, note the worker password assigned must be the same for each worker that contacts the Reload Server. As part of setting of the Reload Server, you must know this password.

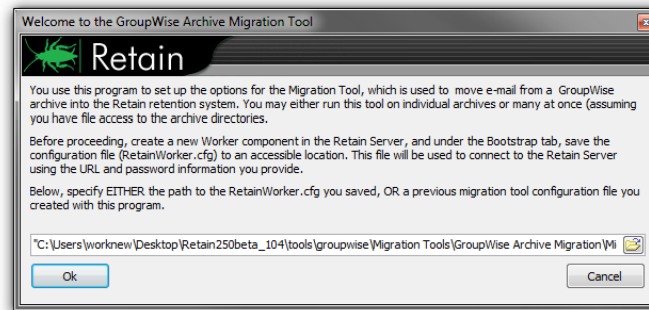
Server Protocol	http
Server Host Name	192.168.1.106
Server Port	80
Server Path	/RetainServer
Worker Password

Note: Changing any of these parameters will require recopying the worker bootstrap file!

- 4) Once you've saved changes, you will be able to save the bootstrap file. Save it in the directory of the migration tool.



- 5) 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.



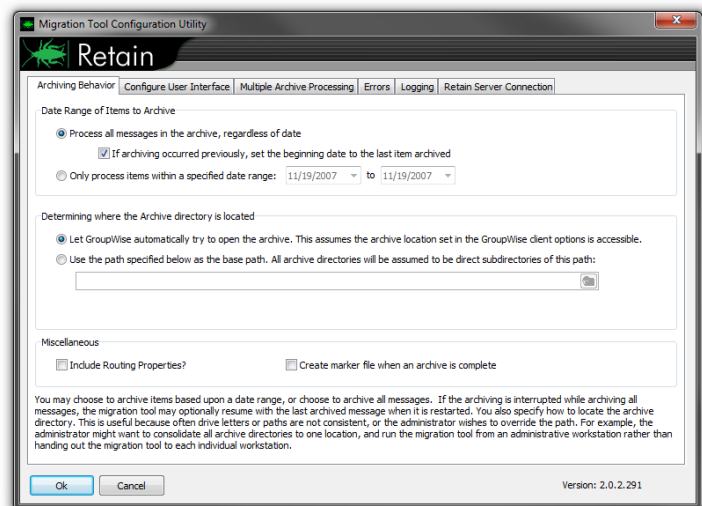
(This window will also accept saved configurations from previous sessions.)

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.




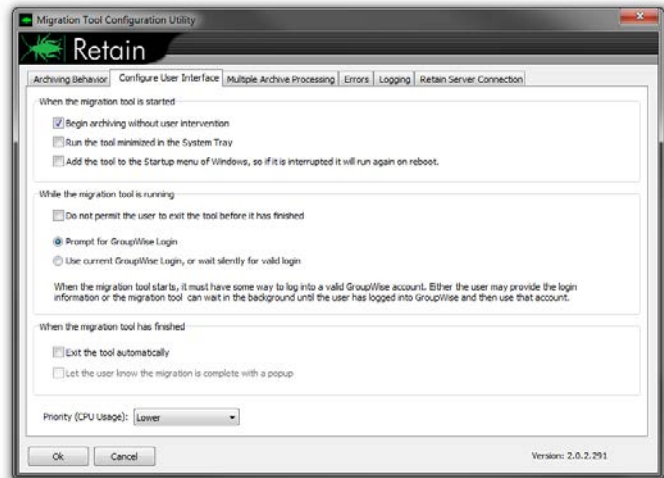
Configure User Interface

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.

The CPU priority usage is also set here.

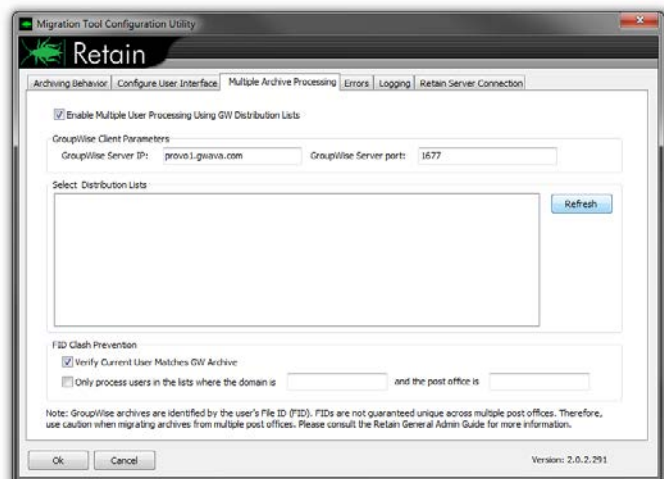


Multiple Archive Processing

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 Multiple User section only appears in **multi-user mode**. (If you click “scan multiple user archives”.) It displays the available distribution lists. 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 verifying archives or only processing users which can be verified by a specified Domain and Post Office, (specified in DNS:port or IP address:port).



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.

Logging


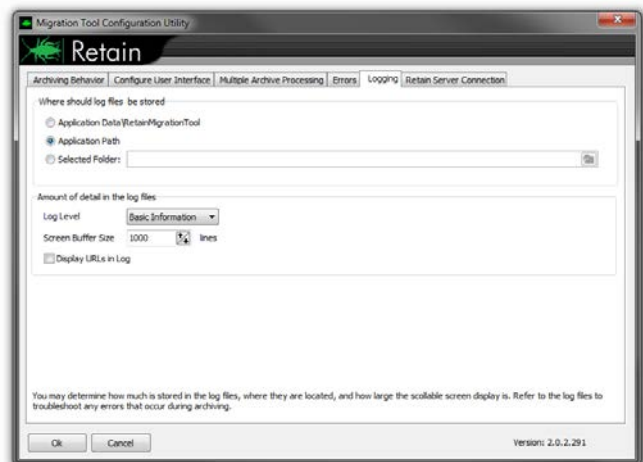
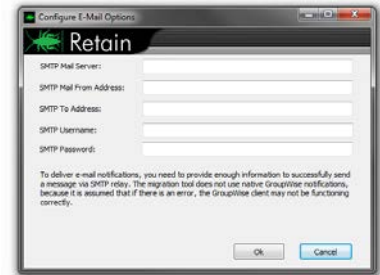
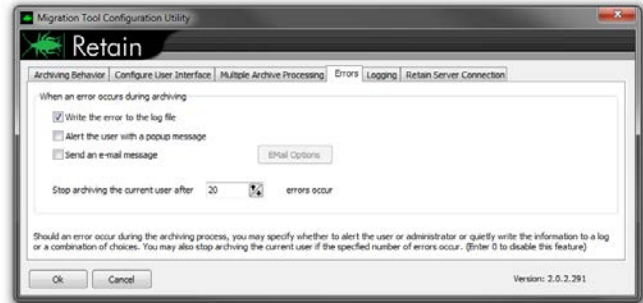
The location, logging levels, and buffer of the logging screen are all configured on this tab.

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.

Retain Server Connection

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.




Once the configuration has been setup as desired, click 'Ok' and select the location to save the changes into a 'MigrationTool.cfg' file.



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 this tool as often as you need to set up the migration job to your liking.

When the tool runs, you will see a screen like this:



You will see any errors in this screen and when the job is complete, it will say so.

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.

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.

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 manually or it can be run automatically.

The Archive Plug-in

This is a simple C3PO plug-in for Win32 clients that allows you to

- add a Retain menu item for easy access to web archives
- hide current Archive options.

It can be run interactively or via command line. Run the program for the command line switches.

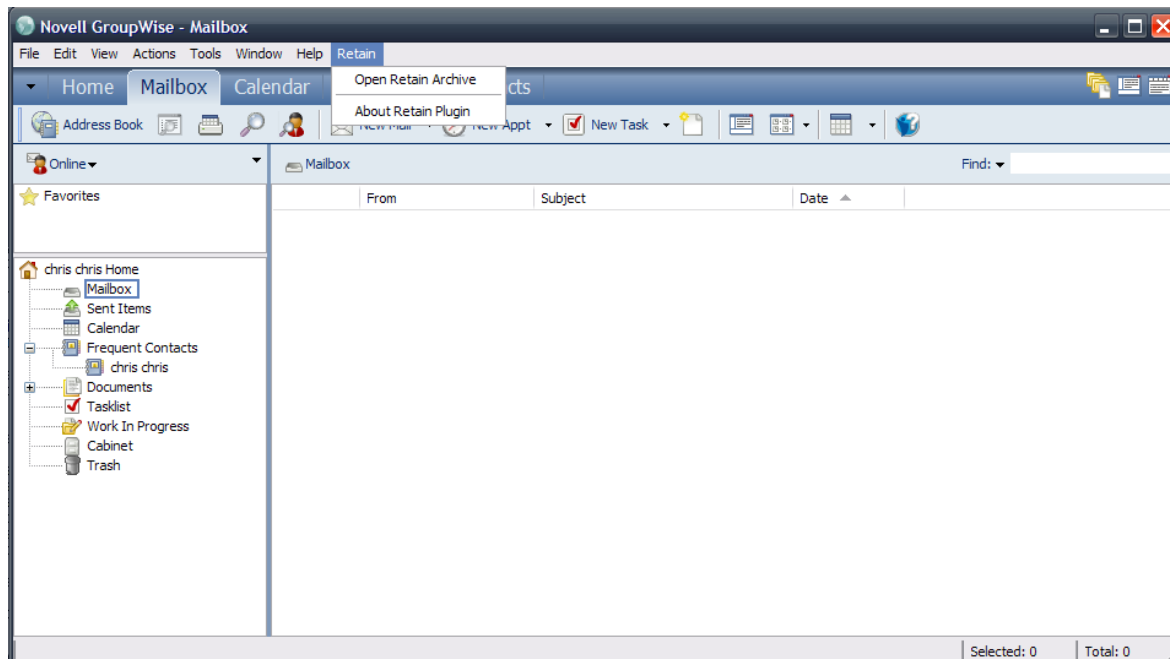


Blob Extraction Tool

This command line utility provides a quick way to extract the original contents of any compressed/encrypted blob in the system. This is for diagnostic and disaster recovery purposes. Simply type BlobExtractionTool at the command line for the usage options.

Single Sign On

The C3PO Plugin for Win32 now supports single sign on. To utilize the single sign on feature, the Win32 GroupWise client must have the Retain plugin installed. The Retain Plugin is located in the downloaded install source. (.../<Retain>/tools/Win32 Client Archive Plugin/Retain Archive Plugin.1.1.0.119.exe) The installer must be run on the machine where Single Sign On will be used. Single Sign requires Retain Server 1.8x. Single Sign On integrates a menu option into the GroupWise Win32 client providing easy access to online archives. The Installation requires the specified address of the RetainServer.



Once selected, the Open Retain Archive option will sign the user into Retain using the GroupWise authentication active in the GroupWise client, providing a transparent login. As such, it is not recommended for use on public workstations.

The program may be run from command line, the options are:

- /url=URL (Provide URL to the Retain Server)
- /hidemenu (Hide standard GroupWise archive)
- /SILENT (Run without user intervention – requires /url argument.)

WebAccess

(Same as Single Sign On above)

Requires:

- GroupWise 8 or later
- (Installer only) Java 1.5
- Windows or Linux

NOTE: If you are running WASP2™ you must update to the latest version of WASP (Version 2.0.5) in order to continue using both WASP and the Single Sign On Retain Module on the same WebAccess server. If an earlier version of WASP is in use, it will be broken with the addition of the incompatible Retain WebAccess Single Sign On plugin.

The WebAccess module .jar file is located in the Retain install sources. After Retain has been downloaded and extracted, the .jar file can be found in the .../Retain180/tools/WebAccess plugin directory.

To install the module, run the .jar java executable file in an X session.

i.e.

On Linux:

```
java -jarRetain_setup_20091005.jar
```

On Windows:

double-click the file.

For NetWare or remote systems SEE THE FOLLOWING LINK:

<http://support2.GWAVA.com/kb/?View=entry&EntryID=1608>

Installing to a remote system **requires** direct file access.

The installer file may be run from a workstation fulfilling the requirements to install on a remote machine. To install remotely, WebAccess source directories must be exported and mounted to the workstation.

i.e.

(For Linux):

/srv/www/tomcat5/base/webapps must be exported and mounted via NFS or a SAMBA share

/var/opt/novell/ must be exported from server and mounted to /var/opt/novell on the workstation.

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.

Windows:

Two folders need to be created on the GroupWise server. From where GroupWise was installed, browse to:

...\GroupWise\webaccess\retain\templates\webacc

Create two new folders here:

mobile

simple

Restart Tomcat and the plugin installation is complete.

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.

Restart Tomcat

`rcretain-tomcat restart`

Once Tomcat has been restarted the plugin installation is complete.

Once installed, WebAccess will add a new tab to the User client window called "Retain", which will sign the user into Retain when selected.

For a text-only or shell install on Linux:

1. Copy the "WebAccess plugin" directory (or contents thereof) to the webserver where the WebAccess Application (the Tomcat servlet, not GWINTER) is installed.

2. Flag the .sh files as executable

`chmod +x *.sh`

3. Run desired install:

`./installtext.sh` (For text mode only install)


`./installgui.sh` (If running X-windows system and want a graphical install)

(The folders described above must still be created on the GroupWise server.)

"WASP" and "WASP2" are trademarks of GWAVA Inc., 100 Alexis Nihon, Suite 500, Saint Laurent, Quebec H4M 2P1

Exchange Tools

The Exchange tools folder contains the tools appropriate for the Exchange mail system: personal archive migration, and single sign on plugins.




Exchange PST Migrator

The PSTMigrator is a Java utility used to input mail in the PST archives into the Retain system. The PST migrator interfaces with the Retain system through a worker configuration file to contact the Retain server and input the mail found in the PST archive to the Retain system.

The PSTMigrator should be run on the local workstation where the Outlook PST archive file is located, and requires Java to be locally installed.

The PSTMigrator requires a Worker configuration and bootstrap file to know where and how to communicate with the Retain Server.

1. In the Retain Server management console, create a new worker, it would be appropriate to name it 'migrator' or 'pst migrator'. A general 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 PSTMigrator have an open connection to the Retain server. When the settings are correct, save the new worker.

The screenshot shows the 'Worker' configuration page in the Retain Server management console. The worker name is 'migrator'. The 'Server Connection' tab is selected. The page contains the following fields and text:

Worker: migrator

Server Connection

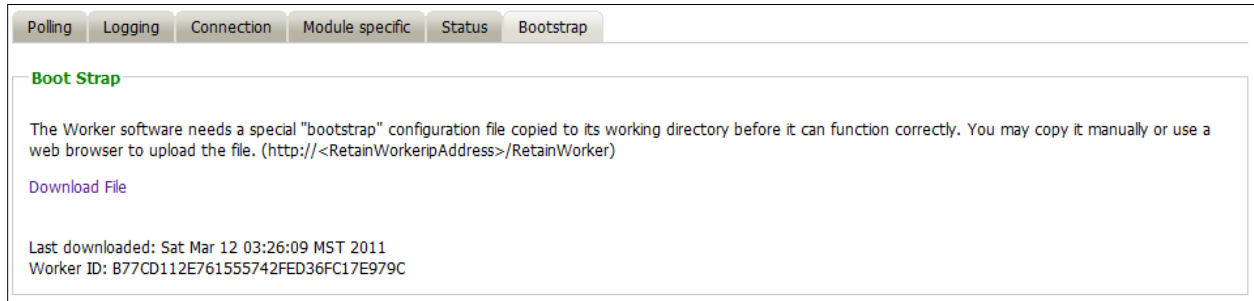
The worker will need to know how to connect to the Server. It is highly recommended that you use a DNS host name for the Host Name. It's important to note that you are specifying the URL from the point of view of the **Worker contacting the Server**, which may or may not match the URL you'd use from your browser to reach the Worker or Server.

For administrators using the Reload interface, note the worker password assigned must be the same for each worker that contacts the Reload Server. As part of setting of the Reload Server, you must know this password.


Server Protocol	http
Server Host Name	192.168.1.106
Server Port	48080
Server Path	/RetainServer
Worker Password

Note: Changing any of these parameters will require recopying the worker bootstrap file!

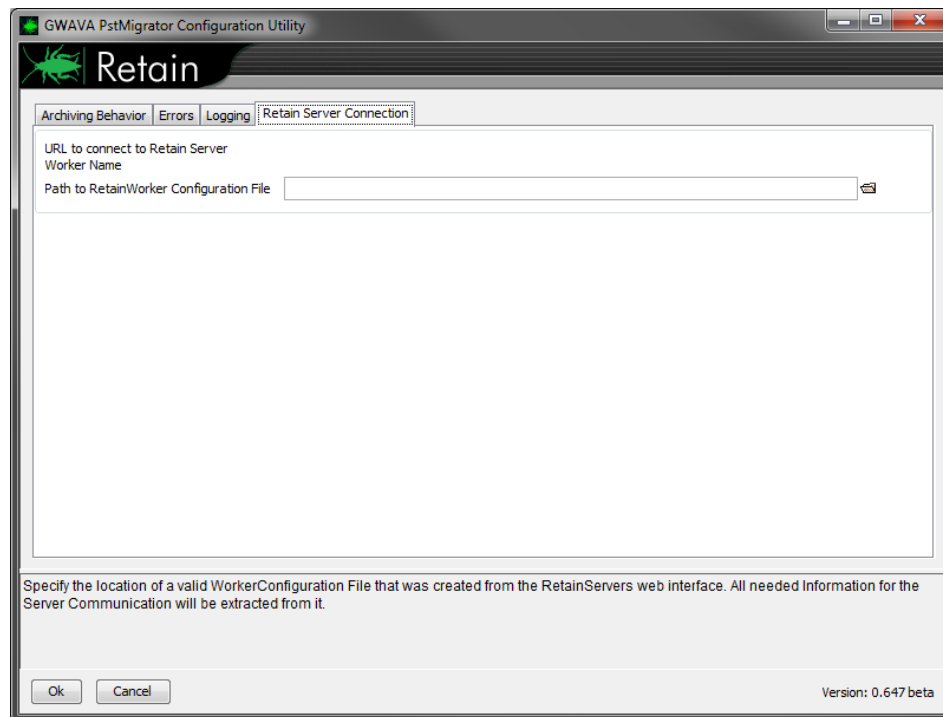
2. Download the bootstrap file.



3. Run the PSTMigrator utility on the Windows workstation containing the pst archive file. If Java is installed, the PSTMigrator.jar file can be run simply by double-clicking on the file. When initially run, the PSTMigrator is placed in the taskbar, and must be selected to be opened.



4. Upload the Worker configuration to the PSTMigrator.
Right click on the PSTMigrator icon in the taskbar, and select 'configuration'.
The initial screen of the migrator is opened to the Retain Server Connection tab, where the bootstrap file is specified.




Browse to the location of the worker configuration downloaded in step 2, and select the configuration.

5. Configuration logging, error reporting, and specify the pst archive file.

Logging

Log files will contain as much information as the administrator desires, from a normal or basic level, to diagnostic.


Select the path for the logging files, whether that is a custom file path, application path (stored in the same folder as the application), or with the application data, (in Retain).



Errors

If an error is encountered, what the migrator utility does with that error is determined by the options configured here. The error may be logged, recorded in a popup message, sent in an email notification, or all three options at the same time.

An error limit may also be specified which determines how many errors the migrator can encounter before ceasing the migration job.



Archiving Behavior

The Archiving behavior tab determines what is archived. If the pst archive file has been partially migrated at a previous time, then the PSTMigrator provides the option to only archive the parts which have not been previously archived.

The PSTMigrator also has the option to only migrate mail according to a set data range. The PSTMigrator will work on all archives listed in the Archive File window. To add an archive file to the migration list, select 'add' and browse to and select the desired pst archive file. After selecting the archive file, the migrator will attempt to access to the archive.



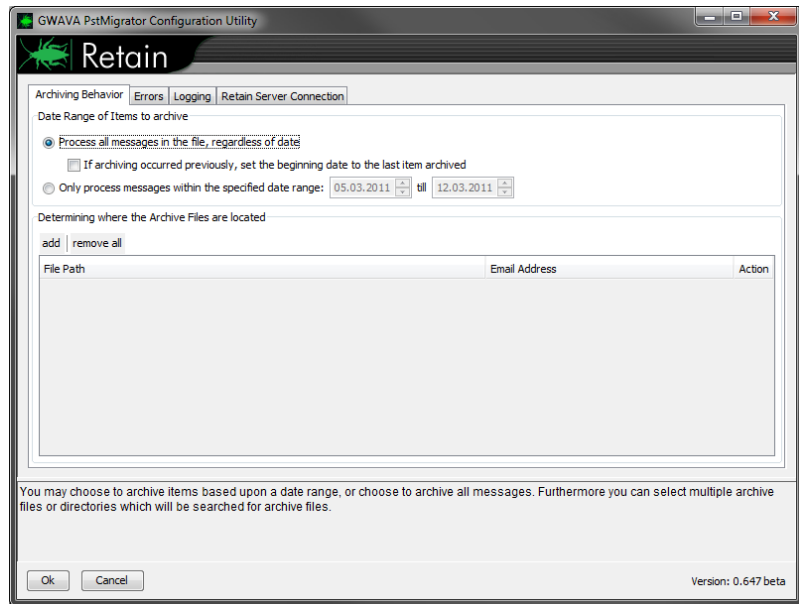
The PSTMigrator may require the email address of the source account in order to achieve access to the archive. Provide the source address.

6. Migrate the pst archive.

Selecting 'Ok' does not initiate the migrator, it only sets the configuration. To start a migration job, right click on the taskbar icon and select 'Run now'.

7. Wait for the migration to finish.

The migrator will run as a background service and complete the entire list of archives. If configuration needs to change between archives, migrate the archives separately with configuration changes in-between.



Single Sign On

Single sign on for Exchange comes in two different versions: one which plugs into the Outlook client, and one that works with WebAccess on the Exchange server. Select the version which fills the appropriate need.

Single sign on creates a new entry in the user interface which opens a new window directly into the user's personal Retain archive. This eliminates the need to log into Retain separately from the User's normal mail client.

Exchange Server WebAccess install

The WebAccess plugin installation is simple; run the installer on all applicable Exchange servers and follow the prompts.

The installer prompts for install location and for the URL of the Retain server.

IIS must be restarted. The installer offers the ability to restart IIS automatically.

Outlook Plugin

(NOTE: Trial versions of Outlook have plugins disabled; the Retain outlook plugin will not work on trial versions.)

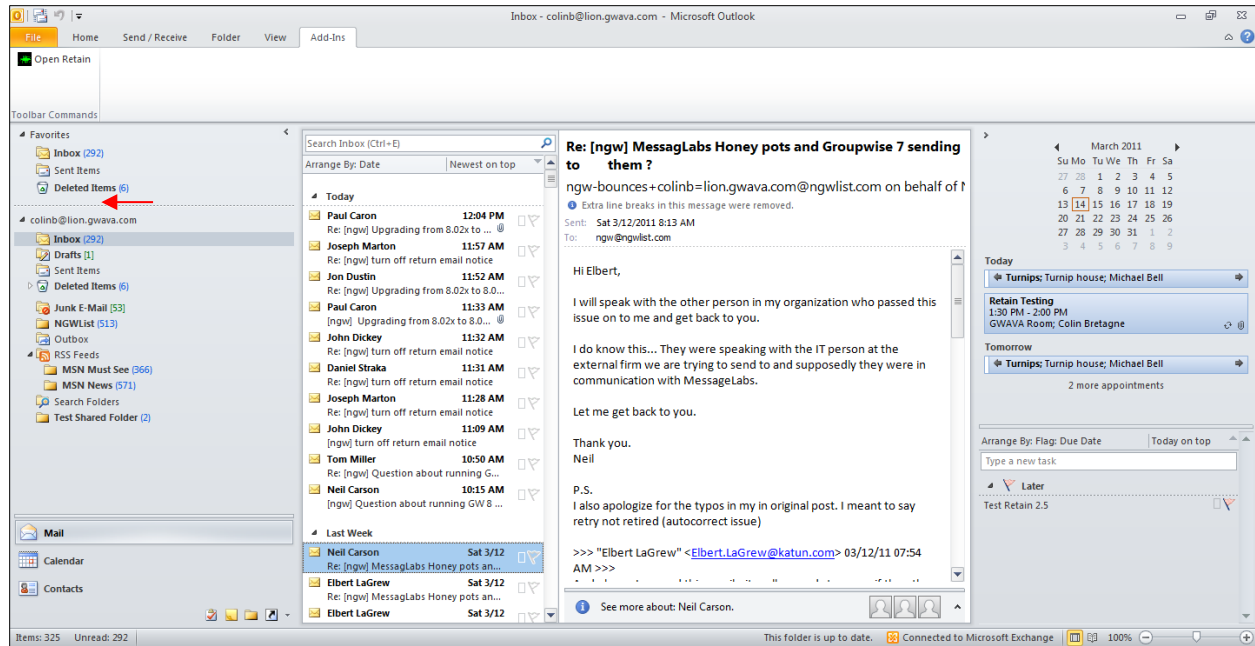
The Outlook plugins, both the 2007 and 2010 versions, are essentially the same thing; Single Sign-on plugins for Retain which allows Outlook users to connect to the Retain archive in simply by connecting to their accounts in the Exchange system. Be sure to install the appropriate one for your system.

The module requires Outlook configured and currently running, authenticated to and part of Active Directory. For the plugin to work correctly, **administrator rights are required during install**.

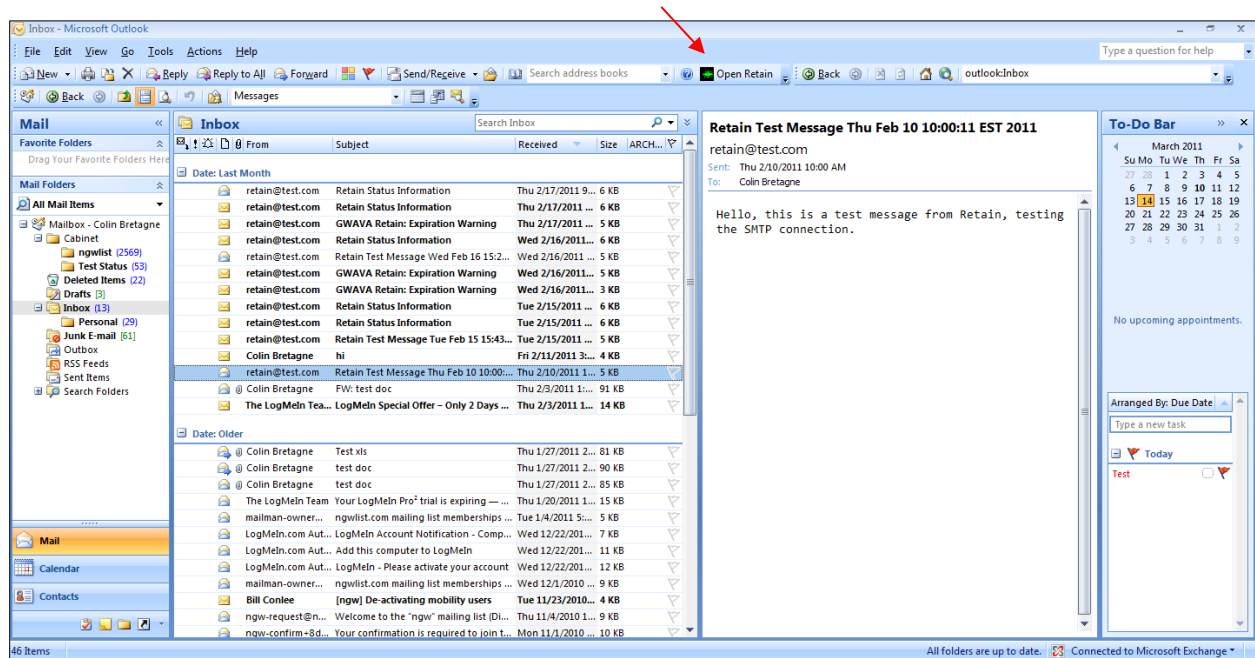
Installation is simple, run the setup.exe installer and follow the prompts. Running the setup.exe installer will download and install all necessary and appropriate files and run the .msi as appropriate. An Internet connection to the target workstation is required to download and install the necessary files, (32-bit or 64-bit as appropriate).

Uninstallation is achieved through the Control Panel | Uninstall or Change a Program utility in Windows.

Accessing the plugin is different in each version of Outlook. The Retain plugin under Outlook 2010 is found under 'Add-ins', where Retain has a toolbar button 'Open Retain'.



In Outlook 2007, the option adds a toolbar button. To access the archive, select the 'Open Retain' button from the toolbar.



Recover Admin

This is an emergency tool you can use if you happen to lose your admin account/password.

This utility will let you create a new admin.

Consult the readme file for further information on its use.

Appendix D – Reload Integration

GWAVA Reload and GWAVA 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?

1. Reload is very good at moving data efficiently from point A to point B.
 - a. It copies your post office data in its original form.
 - b. 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.
 - c. By having the backed up data available in its original form, it can serve as a data source for Retain.
 - d. Reload's backups are available the moment the backup job is complete.
2. Retain moves a lot of data and needs strong network links to do so rapidly.
 - a. 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.
 - b. 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.
 - c. By integrating Reload and Retain, you can centralize your archives and ensure good backups and achieve a single data pull.

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 working.

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 tremendously 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:

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.

This client installed a Reload server in each location. The servers backed up the local post office in addition to making a backup of the Reload server in the other location. Thus, each Reload box effectively had backups of both servers.

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 connection 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%.

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.

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. (I.e.

`./RetainInstall.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.


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.

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 possible.)

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](#)).



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.

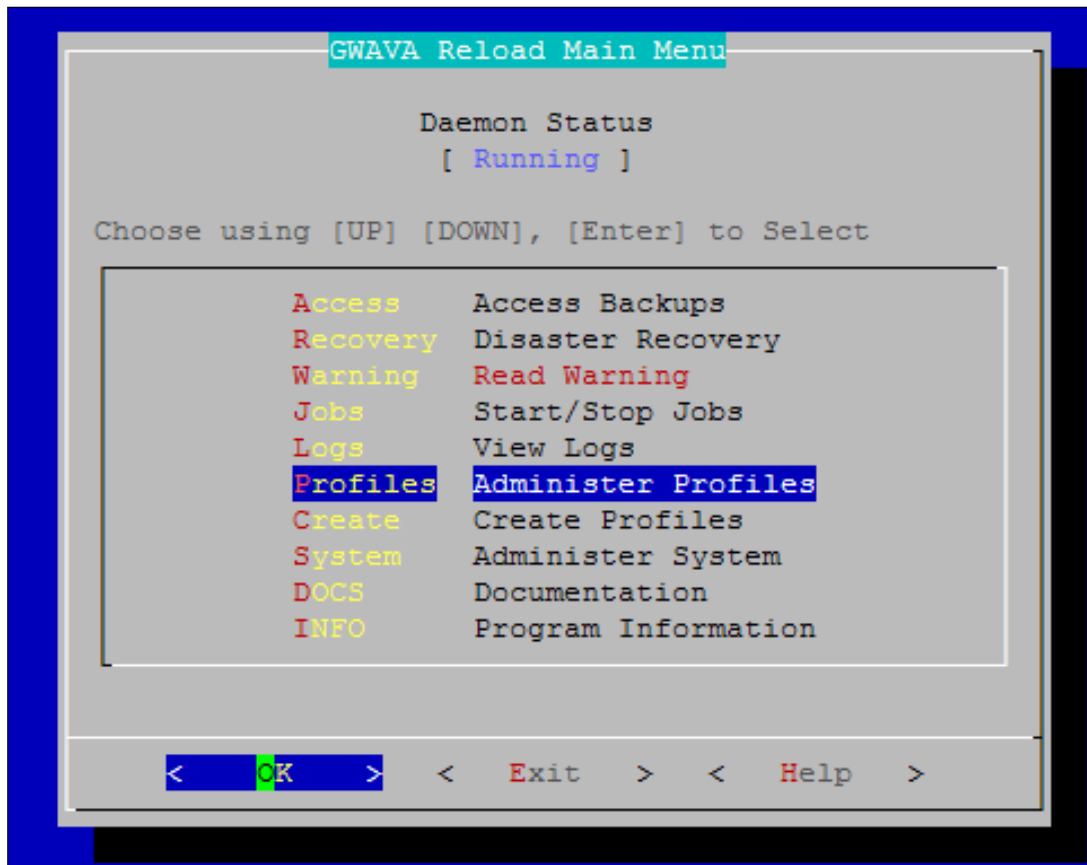
First, Reload must be set up so that the backups are available. There is a special feature in Reload for this. It calls up a post office agent that stays up all the time and it only goes down long enough to change to the most recent backup. So it will always be there with very small interruptions as 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.

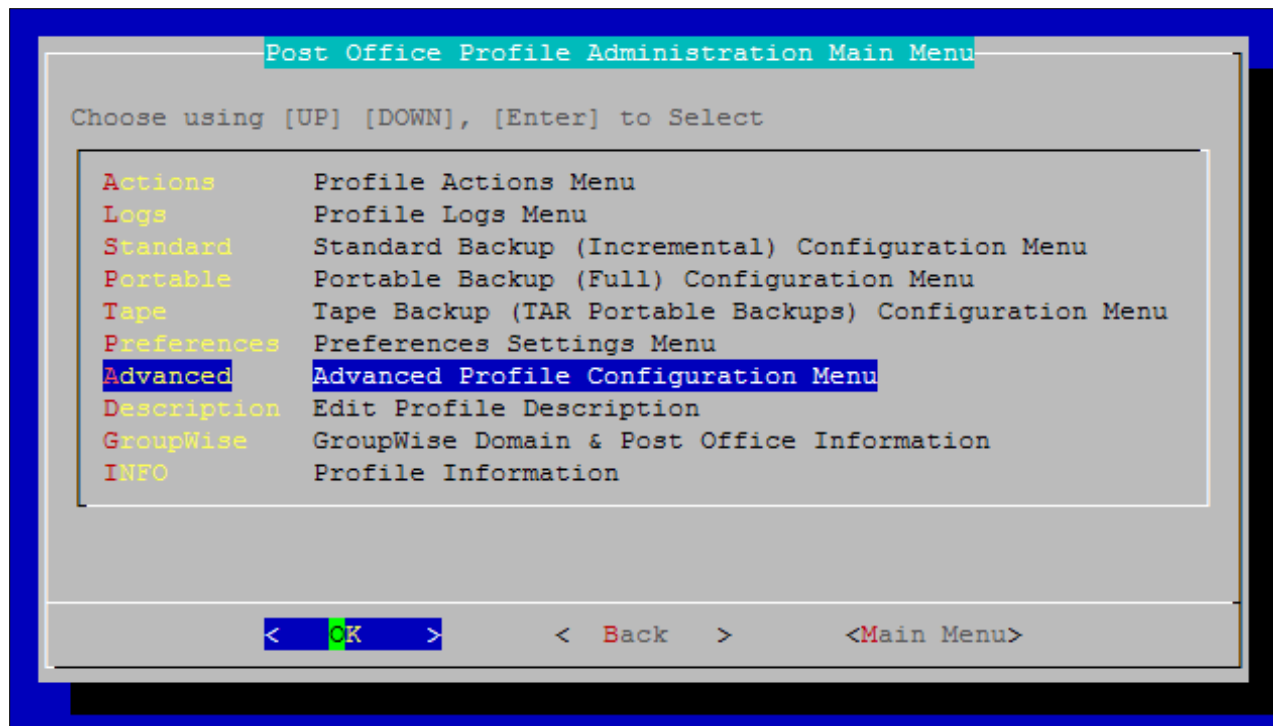
- 1) Start up Reload's Administration menu.



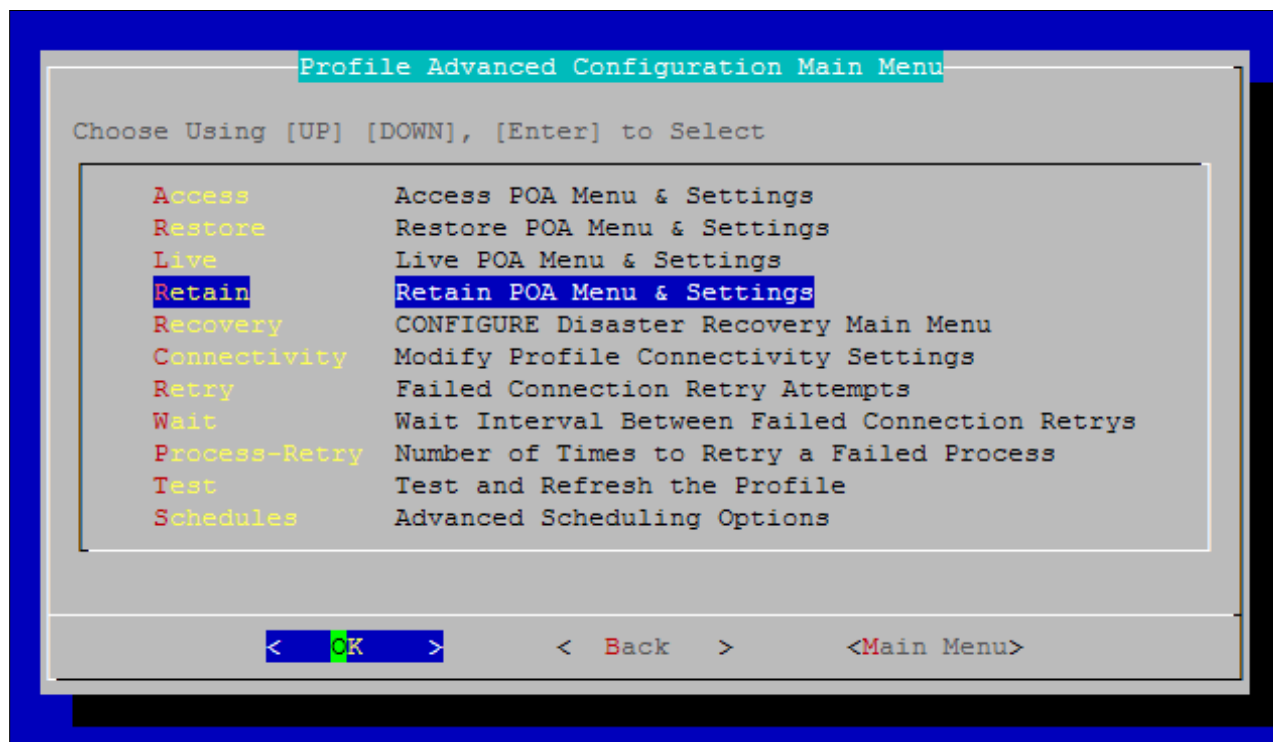
- 2) From the main menu, choose Profiles – Administer Profiles.

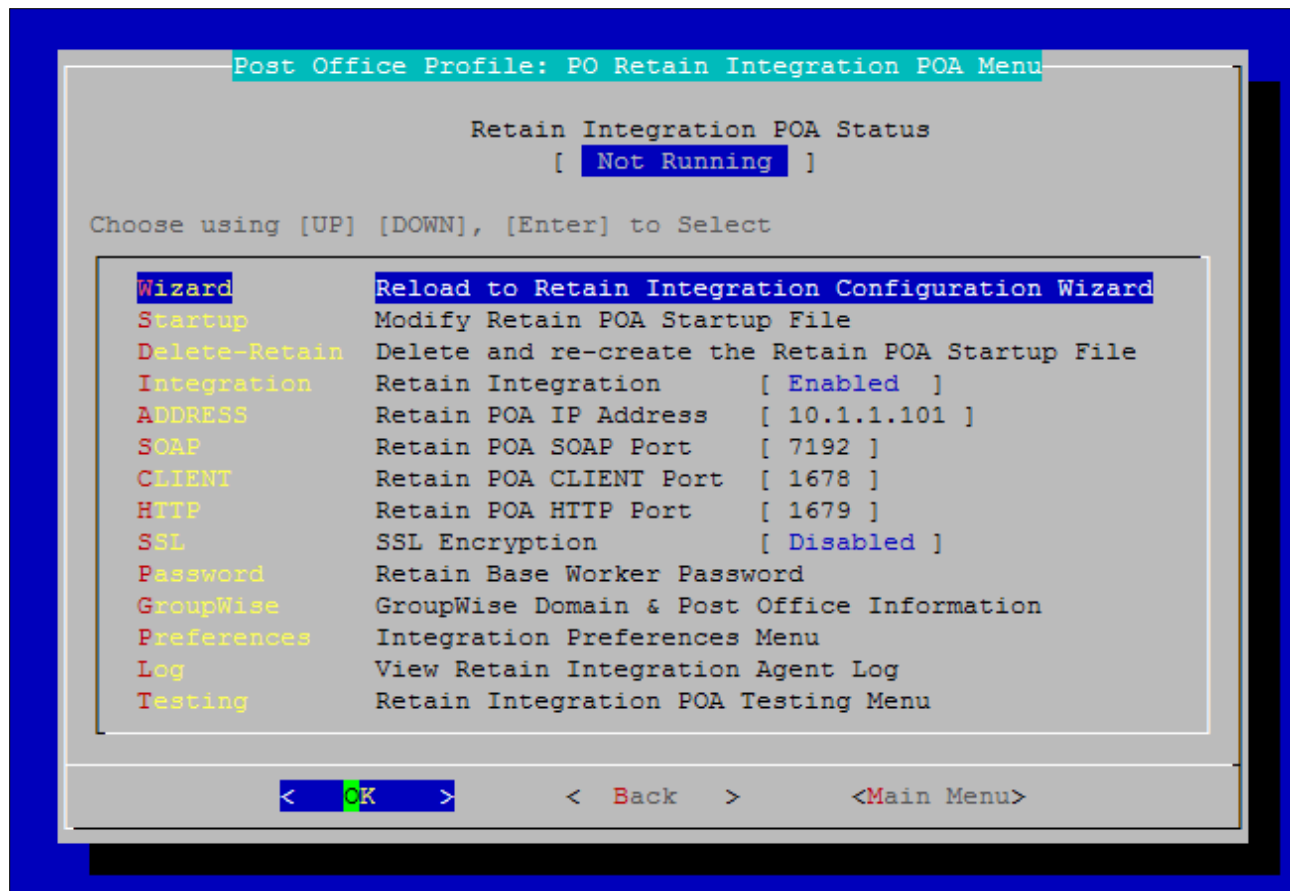


- 3) Choose Advanced Profile Configuration Menu



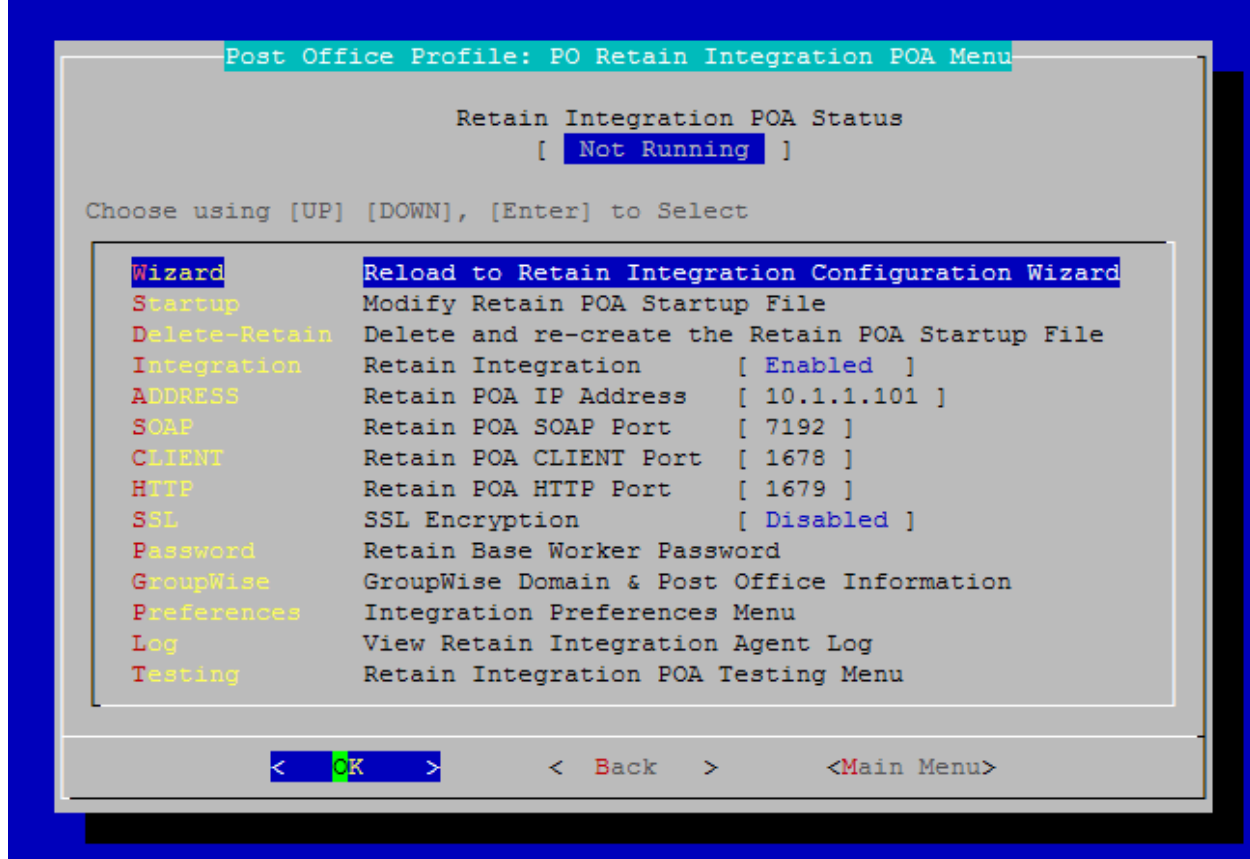
4) Choose Retain POA Menu & Settings





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

GroupWise Post Office Name

In the field below indicate the name of the GroupWise Post Office that this Reload Profile represents.

gwpo

< Save >

GroupWise Domain Name

In the field below indicate the name of the GroupWise Domain that owns the GroupWise Post Office:

gwpo

gwdom

< Save >

2) Enter the name of the post office and domain.

Retain Worker Authentication Key

In the field below indicate a key (string of characters) that the Retain Worker can use to get access to the Reload to Retain integration for this profile.

NOTE: Use only letters and numbers for this key. The case of the letters does matter.

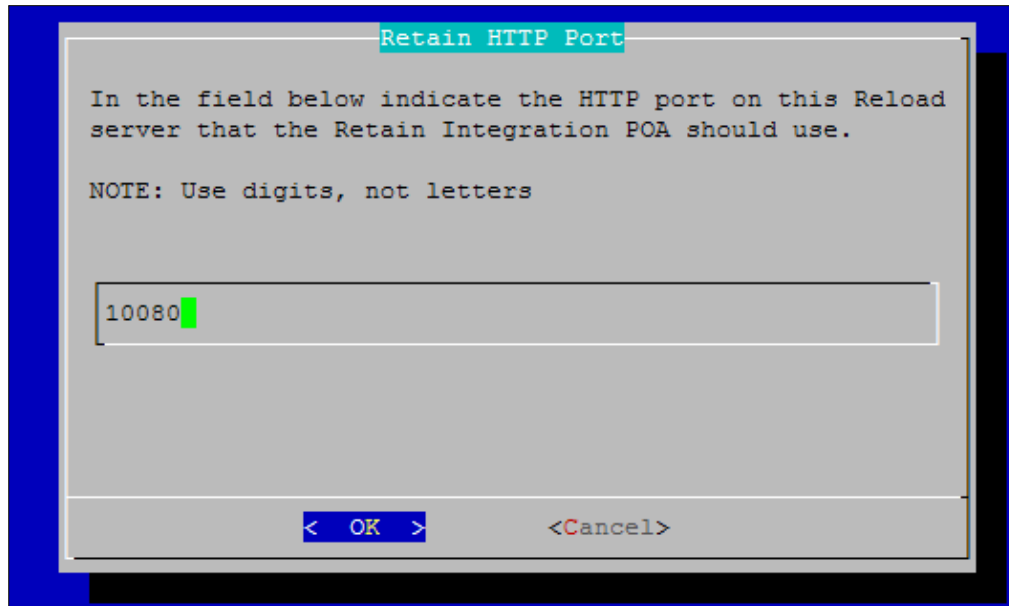
1234

< OK > <Cancel>

3) Choose an authentication key that Retain will use to access this POA. **This must match the password you assigned to the Retain Worker.** (See the [Retain Settings](#))

The image contains two screenshots of a software interface. The top screenshot is a dialog box titled "Retain POA SOAP Port". It contains the text: "In the field below indicate the SOAP port on this Reload server that the Retain Integration POA should use." followed by "NOTE: Use digits, not letters". Below the text is a single-line text input field. At the bottom are two buttons: "< OK >" and "<Cancel>". The bottom screenshot is a dialog box titled "Retain POA CLIENT Port". It contains the text: "In the field below indicate the CLIENT port on this Reload server that the Retain Integration POA should use." followed by "NOTE: Use digits, not letters". Below the text is a single-line text input field. At the bottom are two buttons: "< OK >" and "<Cancel>".

- 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 different or unique port, so there is no conflict with your original POA. If your Reload installation is on a separate 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.

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.

Profiles

default
reload

Add Profile
Edit Profile
Remove Profile

No jobs are associated with this Profile.

Profile reload saved.

Profile Name: reload

Core Settings | Scope | Miscellaneous | Advanced

Date Range to Scan

All Messages (ignore date)

Duplicate Check

Retain only stores a single instance of each message and attachment. Defining how the Worker determines new items (so it may skip sending them to the Server) is an important performance factor. Using the retention or purge flag choices are the fastest choices, if these are options. Never publish all messages unless you are priming the system for the first time.

Publish all messages newer than last stored message (fast)

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.

Set Storage Flags

If you are using either the Purge or Retention features in GroupWise, you probably want these to be advanced automatically as items are stored, so users may delete messages in their mailbox that have been stored by Retain. The Item Store flag is of similar function, stored in Retain itself, but cannot prevent users from deleting item. It is most useful in conjunction with GWAVA Reload or in use with multiple overlapping jobs.

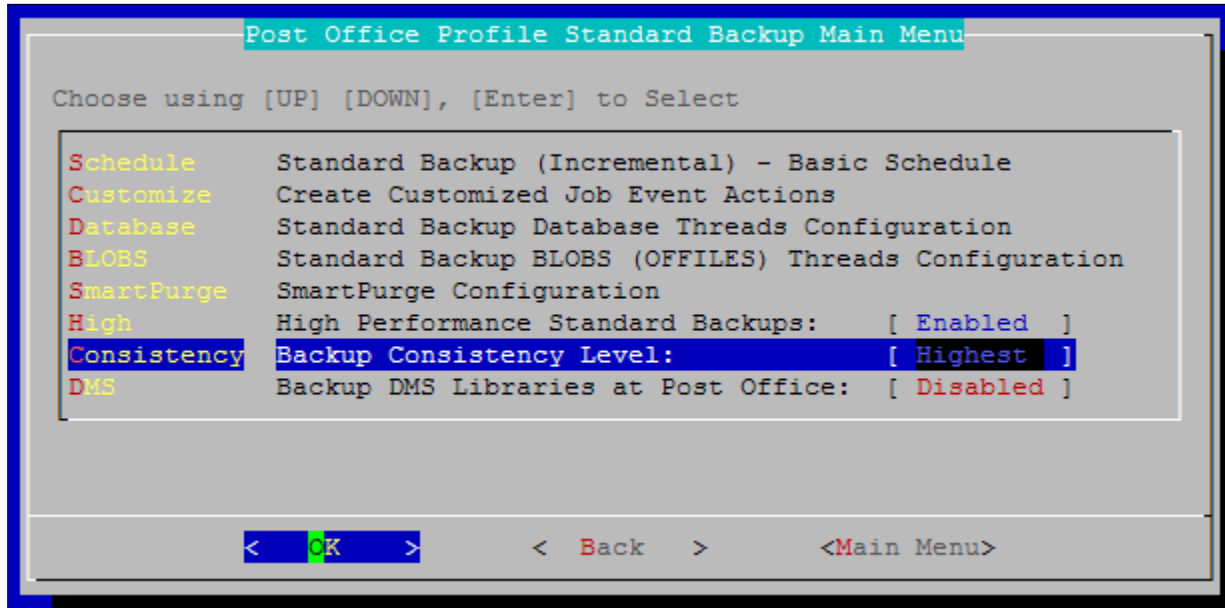
☐ Retention Flag ☐ Purge Flag ☒ Item Store Flag

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 Profile) > 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.

Appendix E – Migrating Retain to a new Server or Platform

Migrating Retain between different platforms is possible, though time consuming. Retain can also be migrated between 32-bit and 64-bit systems. Migrating or moving the Retain Server can be complicated, because the different Retain components that connect to the Retain server by IP address, file paths mapped in the system, and the unique server identification which allows access to the data storage are all tied to the original Retain Server. The important variables for the Retain Server are found in the ASConfig.cfg file. The connecting components will need new bootstrap files. In a migration, all these variables must either be corrected, or copied to the new system for the migration to be a success.

There are **no guarantees** when migrating a server. **BE SURE TO CREATE A FULL BACKUP WHEREVER POSSIBLE BEFORE STARTING THE PROCEDURE.** **SQL databases and data storage systems are not supported or administered by GWAVA Inc., and are the responsibility of the local system administrator and the customer.**

Regardless of the source and destination platform, the main steps will be the same. The steps for migrating a Retain Server are:

1. Create a backup of the server, data store, and database.
2. IF the SQL database is housed on the source Retain Server, copy or migrate the SQL database to its final destination in accordance with recommended practices for the respective SQL system. Setup the SQL server and get it running on the destination machine.
3. Download and install Retain. After Retain has been installed, shutdown Tomcat.
4. Copy the data store to the new destination. The data store consists of more than just the storage directory. The data store contains the configuration database, indexes, backups, license, and the archive, (if not on a SAN). If migrating to Linux, the ownership and execute rights to the data store **MUST** be changed. If ownership execute rights are not set correctly, the Retain Server will not function correctly, and will generate inexplicable errors. Commands are:

```
chown -R tomcat:www <storage_directory>
chmod -R u+rx <storage_directory>/*
```
5. Open the ASConfig.cfg file from the old server with the text editor of choice. (Text editor must be able to view .xml files.) The following file paths which were mapped for the original server must be corrected for the new server:

```
<serverID>unique_server_ID</serverID>
```

It is **ESSENTIAL** that the ServerID setting be identical to the original server ID, or the migration will fail.

```
<basePath>/retain</basePath>
<archivePath>/retain/storage</archivePath>
<xmlPath>/retain/xml</xmlPath>
<indexPath>/retain/index</indexPath>
<backupPath>/retain/backup</backupPath>
<licensePath>/retain/license</licensePath>
<EBDBPath>/retain/ebdb</EBDBPath>
```

ie. If migrating from Windows to Linux, the <basePath> would need to be changed from:
<basePath>C:\retain</basePath> to <basePath>/retain</basePath>

6. When the ASConfig.cfg file has been modified correctly, copy the modified file into the destination server. ../RetainServer/WEB-INF/cfg
Once the file has been copied into the new server, start Tomcat.
7. Verify that theRetainServer is working. If it is not working, view errors and double-check the settings in the ASConfig.cfg file.
8. IF the final destination server has a different IP address than the original RetainServer had, all workers and stubbing servers must be reconfigured by re-uploading new bootstrap files for each one. If this step is not completed, the workers and stubbing server will not have contact with the Retain Server, and will not be able to function. See '[correcting the bootstrap](#)' in the worker section.

The process should be done with extreme care and attention to detail, if not under the supervision and by the guidance of Retain Support. Completing ALL steps for your system is extremely important. If there is a problem call support and/or revert to the backup made in step one. Correctly following these steps will result in a fully functioning system on a new platform.

Appendix F – Switching the Indexing Engine

Switching between indexing engines is possible, though frequent changing is not recommended as it is a time consuming process. This guide assumes that the system will be switched from Lucene to the more robust Exalead indexing engine, though the process is nearly identical for the switch back to Lucene, it simply requires selecting 'lucene' instead of 'exalead' when the options are provided, as the positions of Lucene and Exalead are reversed when switching from Exalead to Lucene.

Switching the indexer to Exalead

The Exalead indexing engine is not required, and the system can be switched from Lucene to Exalead at any time. However, the index engines are not compatible with each other and have no way to export the indexes between them, causing a need to re-index the entire mail archive each time the index engine is changed to ensure full coverage. Re-indexing the archive will take substantial amounts of time, and switching the indexing engine should only be run during a 'quiet time' when there are no jobs scheduled to run.


Exalead allows much more flexibility and options in searching and scaling. However, Exalead does require a separate physical machine with robust hardware. For a full list of Exalead requirements and install instructions, see the Exalead Install document. Switching to the Exalead indexing engine will cause some server downtime.

Before you begin the migration to Exalead for the indexing engine, you must first have Exalead installed and working. Log into the Retain Server management console as Administrator, and ensure that:

- The Retain System has been backed-up.
- Exalead server has been Installed, setup, and is accessible .
- No deletion manager jobs should be running, or set to run.
- Indexer status is set to '0'. On Retain Server status page, the "Total emails in memory queue awaiting indexing" should show '0'. If this is not set to '0', and no job is running, contact support before performing the index engine migration.

From this point on, actual changes will be made in your Retain Server system. To change the index engine, perform the following steps

1. As admin, log into the Retain Server management console and browse to the **Configuration | Server Configuration** page and select the **Index** tab and click on the 'Migration Steps' link to spawn the migration steps window.



Retain Welcome, admin | About | Status & Updates | Logout

Server Configuration

Communications Index **Storage** Accounts Maintenance Notification Logging Miscellaneous

Indexing Engine

At the core of Retain's searching functionality is the indexing engine. Retain supports two indexing engines currently, Exalead and Lucene. It is possible to migrate from one to the other, but this may require significant effort. It is better to choose an indexing engine initially that fits your needs.

Here is a brief summary of the Indexing Engines provided. There is more detail in the user manual:

- Lucene - Fast and easy to use. Requires no additional software or hardware installation. Well suited for small scale environments that have only modest searching and storage requirements. Some features, such as extended queries, enhanced publishing and exporting, etc are not available using the Lucene engine. Not clusterable or high availability friendly.
- Exalead - Fast, extremely powerful, and very scalable. Can be clustered with multiple index and search nodes. All Retain advanced features available. Requires one or more dedicated 64 bit servers (quad core recommended), with at least 16GB RAM. Requires an installation process. Recommended for medium to large customers, as well as organizations that need advanced searching, exporting, availability functions.

If you do decide to switch indexing engines, be sure to read the user manual, and follow the [Migration Steps](#) carefully. The Migration Steps page contains a complete step by step guide, as well as special migration functionality.


Current Indexing Engine:

Indexing

Here, you adjust the different types of items to be indexed. You may add, delete, or edit items.

28 days remaining in eval

Open the migration link in a new window or tab.



Indexing Engine Migration

Summary

Migration from one indexing engine to another is a significant step, one not taken lightly. Inherently the process of migration requires reindexing all previous stored mail, a process which may take days or weeks depending on the size of your store.

Motivations for migrating include:

- Enhanced scalability and availability - for example, Exalead can be installed as a multiple node cluster solution. Lucene cannot.
- Enhanced functionality - We will be introducing features (such as advanced search functionality, enhanced publishing and exporting, etc) that require the Exalead engine. We expect most additional major features depending on the search engine might require Exalead.
- Greater simplicity - Lucene is a simple embedded solution, Exalead requires dedicated hardware.

We urge you to read the user manual and this page and consider your options carefully. Click [here](#) to print this page.

Steps for Migration

Backup

Backup your database, and optionally your indexes. These are both large datasets, but it would be wise to back them if possible as a safety precaution. There are steps that incur database changes. No index changes occur to the old index, so you could always back off to the old index engine by changing the Indexing Engine type in [Server Configuration](#). Please read [this section](#) for more information.

Note the Documents Indexed value

In the [Status and Updates](#) screen, note the total number of Documents indexed, under Server status. Write that number down. Note there is NOT a 1:1 correspondence between Lucene and Exalead in the number, but it can be useful to track the progress once you begin reindexing.


Set up Exalead Server (If switching to Exalead)

Set up your Exalead Server, using instructions in the user manual. Start it, wait about one minute, and then check if you can reach the dashboard (<http://ipaddress:basePort+1/api-ui>), so for example if your Exalead server were at 10.0.0.3, with baseport of 40000, <http://10.0.0.3:40001/api-ui>.

Test the Connection

This window has links and information on the migration as well as the migration settings further down the page. LEAVE THIS PAGE OPEN.

2. Scroll down the page to the link to the Indexer Reset page.



Indexing Engine Migration

DATABASE CHANGE OCCURS HERE AND IN FUTURE STEPS. YOU CANNOT STOP AT THIS POINT, AND DOWNTIME WILL SHORTLY BE INCURRED

All of your data needs to be reindexed. So we must reset the indexed status, which is a field in the EMail table of your database server to the value 0. (After being indexed, the value is set to >0 if fully indexed, and a negative value if only partially indexed. But non-zero in any case!). Then, the background indexer will gradually reindex all of these items.

To reset the indexing status, login as a user with the Manage Servers right (such as your initially created admin account), and run the **Indexer Reset** utility: <http://RetainServerAddress/RetainServer/Util/indexerReset.jsp>. Normally you want to allow it to reset your entire message store. This operation directly affects only the message store; the index is not written too. (although the background indexer will certainly start reindexing items relatively shortly).

Switch Indexing Engines

In **Server Configuration**, under the Indexing tab, change your configuration to the new engine, and configure as needed (For example baseport and host for Exalead.). Save your changes, and then **continue to the next step, without pause**

Set Migration Options

(Optional but recommended strongly)

By default, the new engine will handle all indexing, deletion, and searching tasks. Because reindexing takes a while, switching over to the new Indexing engine for your searches would result in NO SEARCH HITS initially. This may be unacceptable. Instead you may want to temporarily continue to use your old indexing engine for search results. Of course, it's important to realize the searches will get increasingly out of date, as new messages will be stored in the new index only. You may also optionally configure an email to be sent to a designated person when the indexing is complete. Again, keep in mind all new items will be indexed only by the NEW engine, and unavailable for search results.

To configure these options, fill out the following form. **Do not fill out and submit this form until you are ready to do so!**

Read the instructions below the form before filling it out and submitting it

Migration Parameters

☐ Migration is in progress?

☐ Activate an alternative searching server?

Searching Server should be


☐ Send an e-mail when migration is completed?

Comma delimited list of e-mail recipients:

Type "I understand". This is required to save the form, and is for your protection:

Open the link in a new page or tab. <http://RetainServerAddress/RetainServer/Util/indexerReset.jsp>

By default, the Indexer Reset page should look like the picture below. (Pay attention to the Unindexed Items (DB): number. It should read '0'. If it does not, do not proceed unless under the instruction of support.



Indexer Background Indexer is alive: true
Indexer currently RUNNING engine: lucene (You may have already switched engines, but this will not be reflected until restart)

Unindexed Items (DB): 0

This utility allows you to reset the indexed status of either the entire message store or a specific mailbox. It may take quite some time on a large database.

Do not run this utility without specific instructions to do so! Misuse can corrupt the index, lead to high CPU utilization, and be quite inconvenient for users, as until the index is rebuilt any items affected will not be searchable.

With Lucene, resetting the index is NOT enough. You must either delete the entire existing index (with tomcat unloaded) or delete all items belonging to a certain mailbox. Otherwise redundant information will be stored! See [luceneDelete.jsp](#) for more.

☐ Reset the entire message store. All items will slowly be reindexed by the background indexer

☒ Reset just a specific user's store. All items will slowly be reindexed by the Background indexer

User UUID to reset (case sensitive):

You may also add an additional restriction - that only email stored after a certain timestamp value is affected. Leave this entry as -1 to disable this feature.

Timestamp restriction:

You must type "I AGREE" below before hitting the submit button.

I understand the significant implications of what I am doing:

The next few steps need to be performed with little time delay between them. No more than a couple minutes. If not switched soon, the indexer reset switch will cause Lucene to re-index, instead of causing Exalead to index the message store, resulting in empty search results when switched to Exalead and slow responses due to Lucene re-indexing the message store. Read through the next steps and decide settings before proceeding. These steps must be completed in the order specified.

- From the Indexer Reset page, Select **Reset entire message store**. Type **"I AGREE"** in the dialog after reading the warnings. Select **Submit**.

Indexer Background Indexer is alive: true
Indexer currently RUNNING engine: lucene (You may have already switched engines, but this will not be reflected until restart)

Unindexed Items (DB): 0

This utility allows you to reset the indexed status of either the entire message store or a specific mailbox. It may take quite some time on a large database.

Do not run this utility without specific instructions to do so! Misuse can corrupt the index, lead to high CPU utilization, and be quite inconvenient for users, as until the index is rebuilt any items affected will not be searchable.

With Lucene, resetting the index is NOT enough. You must either delete the entire existing index (with tomcat unloaded) or delete all items belonging to a certain mailbox. Otherwise redundant information will be stored! See [luceneDelete.jsp](#) for more.

☒ Reset the entire message store. All items will slowly be reindexed by the background indexer
☐ Reset just a specific user's store. All items will slowly be reindexed by the Background indexer
User UUID to reset (case sensitive):

You may also add an additional restriction - that only email stored after a certain timestamp value is affected. Leave this entry as -1 to disable this feature.
Timestamp restriction:

You must type "I AGREE" below before hitting the submit button.

I understand the significant implications of what I am doing:

Proceed **without delay** to the next steps

- THIS STEP MUST BE DONE QUICKLY AFTER RESETTING THE INDEX. Go the server configuration in the Retain Server admin page, and switch indexer engine to 'exalead'.

Welcome, admin | About | Status & Updates | Logout

Server Configuration

Communications Index Storage Accounts Maintenance Notification Logging Miscellaneous

Indexing Engine

At the core of Retain's searching functionality is the indexing engine. Retain supports two indexing engines currently, Exalead and Lucene. It is possible to migrate from one to the other, but this may require significant effort. It is better to choose an indexing engine initially that fits your needs.

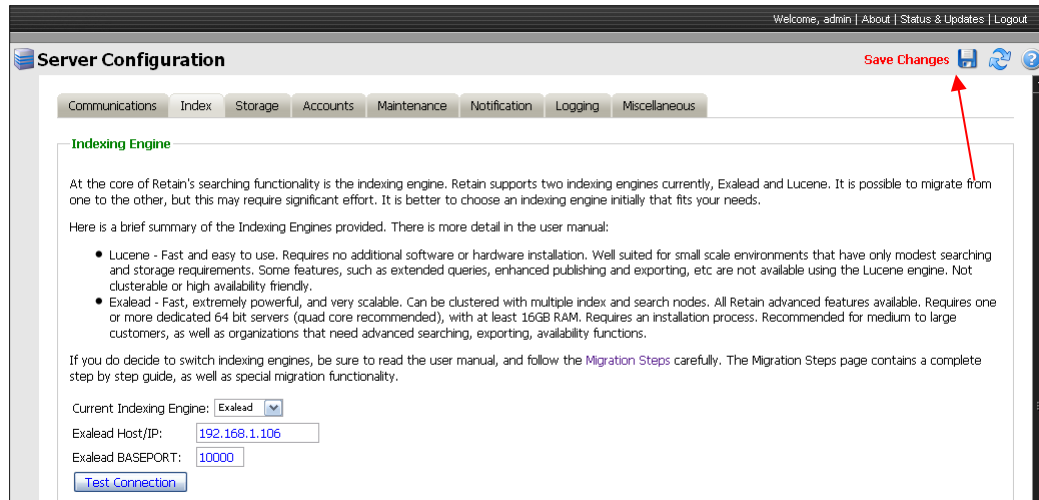
Here is a brief summary of the Indexing Engines provided. There is more detail in the user manual:

- Lucene - Fast and easy to use. Requires no additional software or hardware installation. Well suited for small scale environments that have only modest searching and storage requirements. Some features, such as extended queries, enhanced publishing and exporting, etc are not available using the Lucene engine. Not clusterable or high availability friendly.
- Exalead - Fast, extremely powerful, and very scalable. Can be clustered with multiple index and search nodes. All Retain advanced features available. Requires one or more dedicated 64 bit servers (quad core recommended), with at least 16GB RAM. Requires an installation process. Recommended for medium to large customers, as well as organizations that need advanced searching, exporting, availability functions.

If you do decide to switch indexing engines, be sure to read the user manual, and follow the [Migration Steps](#) carefully. The Migration Steps page contains a complete step by step guide, as well as special migration functionality.

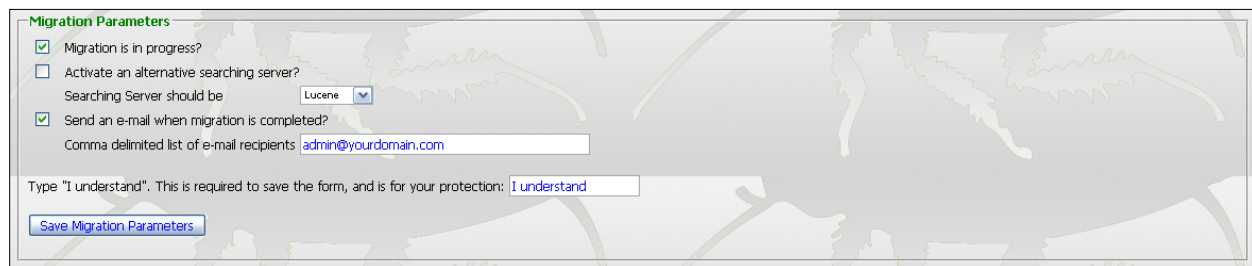
Current Indexing Engine:

After selecting Exalead as the indexing engine, input the connection information (IP address and base port, the default base port is 10000).



After all changes have been made, select the **Save Changes** button.

5. Go to the **Migration Steps** page and Select, migration in progress, and select to **Send an e-mail when migration is completed**. Specify a destination address.



6. Activate an alternative searching server if desired.

While Exalead is indexing the archive, there will be no mail in the index to be searched through. Depending on how much mail is in the archive, the indexing operation may take a very long time. To have the searching features of Retain available while the archive is processed, specify a separate searching server. In this case, it would be Lucene. The only downside to this operation is that any new mail added to the system during the indexing operation will not be added to the Lucene index, and will not be available to search until Exalead has finished and the search engine has been changed to Exalead.

7. After the settings are correct, type **"I understand"** and select **Save Migration Parameters**.
8. Restart Tomcat on the Retain Server system. When Tomcat starts back up, the settings will be cemented and the archive will be re-indexed into Exalead.

During the indexing process of Exalead, the statistics of the Retain Server Status and update page should rapidly increase. It is one way to monitor the progress of the migration. (The screen shot shown is before the migration is started.)

System Status

Jobs Workers **Server** Updates

Server

Index Migration in Progress	false
SQL Server	OK
Server	RetainServer (FXROPPHOTTI)
Stored EMails	0
Engine	standard
Server started at:	10-Mar-2010 14:00:11
Archived since server started (total/dupes/new)	0 / 0 / 0
Server errors	0

Installation time	08-Mar-2010 15:10:53
Enable Address Book Caching	Yes
Last Address Book Cache	08-Mar-2010 15:13:43
Enable Index Optimization	Yes
Last Index Optimization	Never

Total documents ever indexed	19897
Total emails in memory queue awaiting indexing	0
Total documents indexed so far since uptime	0
Total emails indexed so far since uptime	0
Total emails with indexing errors since uptime	0
Total emails with items skipped (*) since uptime	0
Last item indexed at	Never

9. When the migration complete email arrives, return to the Migration Steps page and unselect the **Migration is in progress** setting, type **I understand**, and **Save Migration Parameters**.

Migration Parameters

☐ Migration is in progress?

☐ Activate an alternative searching server?

Searching Server should be: Lucene

☒ Send an e-mail when migration is completed?

Comma delimited list of e-mail recipients: admin@yourdomain.com

Type "I understand". This is required to save the form, and is for your protection: I understand

[Save Migration Parameters](#)

Even though Lucene has been unselected, it will still exist with all the indexes until removed. Again, Re-indexing may take a long time depending on the hardware of the Exalead system, the size of the archive, and the speed of the connecting network. Allow for plenty of time for the indexing to be performed. Time may range from several days to several weeks for extremely large systems. Once it has been determined that Lucene is no longer wanted or needed as a backup for the system, to save system disk space, the Lucene index should be removed from the system. This is not a necessary step, though it does free up system disk space that currently holds duplicate data that now resides in the more robust Exalead system.

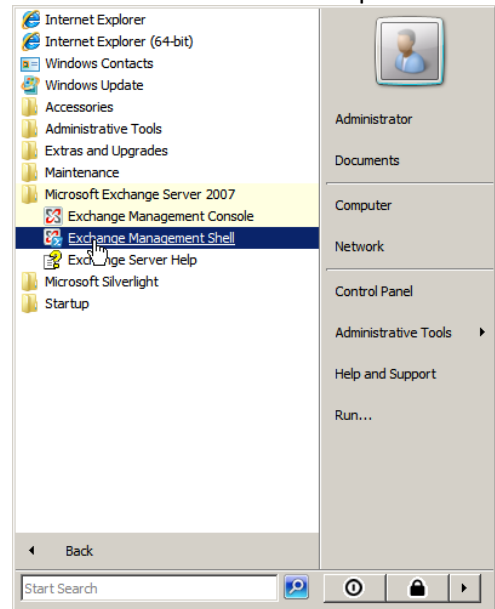
To remove the Lucene Indexes, login as admin to the Retain Server, open a new window or tab, and browse to http://Retain_Server_Address/RetainServer/Util/luceneDelete.jsp to remove the users one by one according to UUID. The lucene delete page is also linked from the **Indexer Reset** page. To manually delete the Lucene indexes, remove the contents of the <storage area>/index directory. There should also be a copy of the index directory in the backup folder of the storage area. This measure is not required and

Appendix G – Exchange Permissions required for Retain

Exchange settings

To connect with exchange, Retain needs a user with appropriate 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 administrator 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.

Grant Impersonation Permissions to the Retain user.

The commands required are different depending on the version of the Exchange Server. Exchange 2010 requires 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 2007 and 2010 servers, the different commands must be completed on one server of each type.

Exchange 2010 command:

For Exchange 2010, 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.

Exchange 2007 commands:

(‘Retain’ is the name of the Retain user, or Service Account below)

```
Get-ClientAccessServer | Add-AdPermission -UserRetain  
-ExtendedRights ms-Exch-EPI-Impersonation
```

```
Get-MailboxDatabase | Add-AdPermission -UserRetain  
-ExtendedRights ms-Exch-EPI-May-Impersonate
```

```
Get-MailboxServer | Add-ADPermission -userRetain -ExtendedRights  
Send-As, Receive-As, ms-Exch-Store-Admin
```

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 both 2010 and 2007.

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 2010 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
```

Exchange 2007 commands:

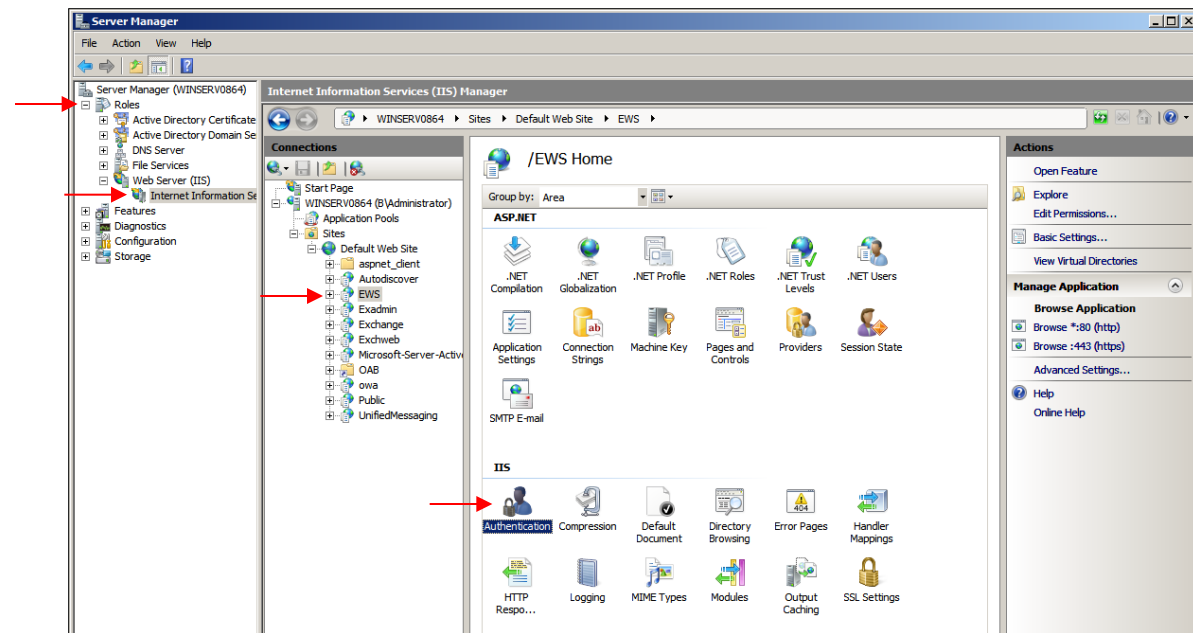
```
Get-MailboxDatabase | ForEach-Object {Get-Mailbox -ResultSize  
Unlimited -Database $_.DistinguishedName | Add-MailboxPermission  
-UserRetain -AccessRights FullAccess}
```

```
Get-MailboxDatabase | ForEach-Object {Add-ADPermission -Identity  
$_.DistinguishedName -UserRetain -ExtendedRights  
Receive-As, Send-As}
```

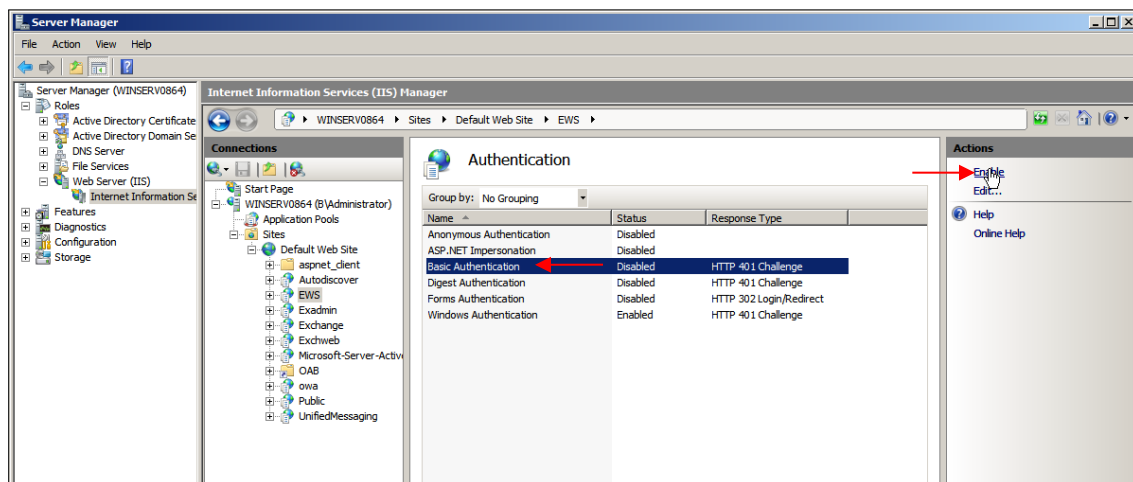
Authentication Methods

Retain requires Basic Authentication to be enabled on EACH CAS Exchange server in the system. Open “Server Manager” on Exchange server.

1. In left pane, expand “Roles”, expand “Web Server (IIS)”, select “Internet Information Services (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.

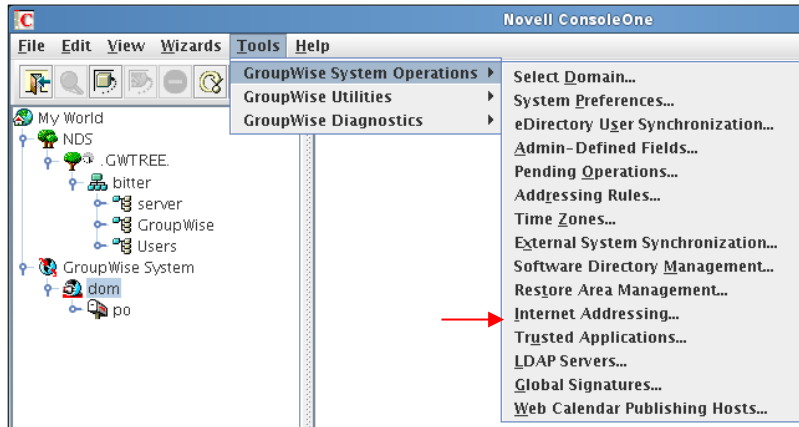
You can now close “Server Manager”.

Appendix H: Trusted Application Key Creation

There are two ways to create a Trusted Application key: through ConsoleOne in GroupWise 8, or via utilizing a Trusted Application key generator on a connected and authenticated win32 workstation.

Trusted Application Key creation through Console One for GroupWise 8




ConsoleOne with GroupWise 8 supports the native creation of a Trusted Application Key. To create a trusted application key for Retain, select the GroupWise system from the system tree and select **GroupWise System Operations | Trusted Applications...** from the 'Tools' menu.



The Configure Trusted Application window will open, displaying any and all trusted applications currently listed in the system.

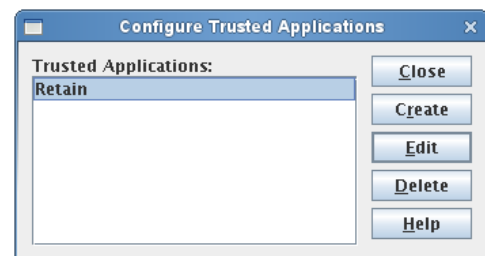
From the Configure Trusted Applications window, select 'Create'.

Enter the Name and description of the Trusted Application, in this case, 'Retain'. The name and description may be anything desired. Check the 'Provides Message Retention Service' and 'Archive Service Address' and provide the IP address or DNS name to the Retain server and the port: 48080



The location for

the key file and name may be anything desired, but the key file is saved in plain text, so add the .txt extension or open with an editor. The key may be copied and pasted into the Retain interface, or entered manually. The key file may be saved for records or discarded. The new Trusted Application Key is active and now appears to be edited, removed, or removed and re-created.



GroupWise Trusted Application Key generator

IMPORTANT Information about the Trusted Application Key generator:

- The Trusted Application Key generator must be run from a win32 or compatible system.
- **You must use the program included with Retain.** The program included with other GWAVA products won't work for Retain.
- Your workstation must have the Novell and GroupWise clients installed.
- You must be logged in as the ADMINISTRATOR of the PRIMARY DOMAIN.
- The Trusted App Generator requires access to the domain database; in Linux this requires a SAMBA share or equivalent.

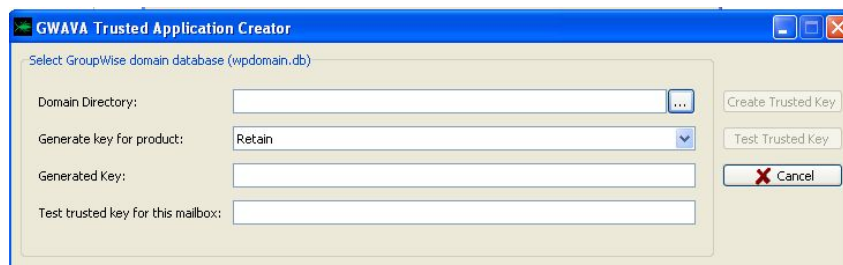
Note: The Trusted Application Key takes a few moments to propagate through the GroupWise system and therefore it may not be ready for testing or use immediately after you've generated it.

From the GroupWise Module page, select the link to download the Trusted Application Key generator, and select 'run', or run a previously downloaded and saved copy.

- 1) Click RUN to download and run the Trusted Application Key generator.




- 2) Follow the wizard. Click Next.



- 3) Find your GroupWise domain database. Click on the “...” button.

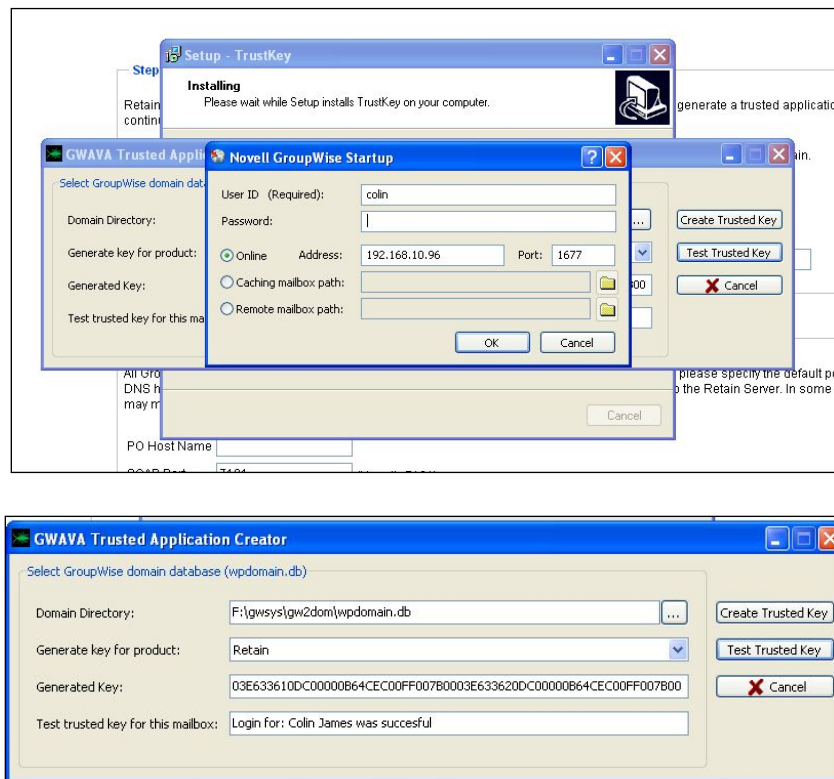


- 4) Once you have found your domain, click “Create Trusted Key”.



NOTE: Please allow a few moments for the Trusted Application Key to propagate through the GroupWise system before testing it. You may proceed without testing the key but if you want to test it, please allow the time for the key to propagate through the system.

- 5) When the key has been generated, enter the name of a valid user on your post office and test the key by clicking “Test Trusted Key”. You will be asked to log in to GroupWise for this user.



- 6) Once your login was successful, you know that both your trusted application key and the SOAP mechanism are working properly. The trusted application key is now copied to your clipboard so you can paste it to the Retain configuration page.